Learnable scattering filters extract more cosmological information from limited simulations.

# Learning Balanced Field Summaries of the Large-Scale Structure with the Neural Field Scattering Transform

Matthew Craigie,[1,*] Yuan-Sen Ting,[2,3] Rossana Ruggeri,[1,4] and Tamara M. Davis[1]

[1] *School of Mathematics and Physics, The University of Queensland, QLD 4072, Australia*
[2] *Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA*
[3] *Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH 43210, USA*
[4] *School of Physics and Chemistry, Queensland University of Technology, QLD 4000, Australia*

Tested on noiseless simulated lensing maps; real-survey performance remains untested.

We present a cosmology analysis of simulated weak lensing convergence maps using the Neural Field Scattering Transform (NFST) to constrain cosmological parameters. The NFST extends the Wavelet Scattering Transform (WST) by incorporating trainable neural field filters while preserving rotational and translational symmetries. This setup balances flexibility with robustness, ideal for learning in limited training data regimes. We apply the NFST to 500 simulations from the CosmoGrid suite, each providing a total of 1000 square degrees of noiseless weak lensing convergence maps. We use the resulting learned field compression to model the posterior over $\Omega_m$, $\sigma_8$, and $w$ in a $w$CDM cosmology. The NFST consistently outperforms the WST benchmark, achieving a 16% increase in the average posterior probability density assigned to test data. Further, the NFST improves direct parameter prediction precision on $\sigma_8$ by 6% and $w$ by 11%. We also introduce a new visualization technique to interpret the learned filters in physical space and show that the NFST adapts its feature extraction to capture task-specific information. These results establish the NFST as a promising tool for extracting maximal cosmological information from the non-Gaussian information in upcoming large-scale structure surveys, without requiring large simulated training datasets.

arXiv:2506.05090v3 [astro-ph.CO] 10 Jun 2025

## I. INTRODUCTION

The universe's large-scale structure encodes rich information about cosmological parameters and their evolution. Effectively capturing this information is essential to maximize scientific returns from upcoming missions by tightening cosmological parameter constraints, in an effort to reconcile growing discrepancies between observations of the universe and the standard ΛCDM model of cosmology [1–5].

Extracting cosmological information from the large-scale structure requires effective compression of the various uninformative degrees of freedom in the field. For this compression, cosmology inference methods often use summary statistics: compact sets of statistics that capture selected information from the field. Traditional cosmology inference approaches develop analytic models that make predictions for how these statistics depend on cosmology, and compare these predictions to observations. The leading approach in modern large-scale structure survey analysis involves using two-point statistics, which are widely used across galaxy clustering [6], weak lensing [7–9], and joint clustering-lensing analyses [10–12].

However, these two-point methods cannot access the non-Gaussian portion of the large-scale structure, which is rich in cosmologically relevant information. Capturing this non-Gaussianity requires accessing higher-order information. This information is most directly accessed by 3-point and 4-point statistics [13–18], but modeling challenges and noise sensitivity limit their applicability to cosmology inference [19–21].

One powerful alternative is to use deep learning to capture field information. Deep learning models are highly flexible, meaning they can adapt to capture cosmologically relevant information. The leading models in vision-based deep learning tasks, Convolutional Neural Networks (CNNs), show promise for constraining cosmology with 2D weak lensing datasets, with applications to the Kilo-Degree Survey [22] and Dark Energy Survey [23]. Likewise, CNNs have proven successful for constraining cosmology with the 3D galaxy distribution [24].

However, CNNs generally require large datasets to generalize effectively due to their high flexibility. For cosmology simulations with sufficiently high precision to accurately model non-Gaussian field features, it can be a computational challenge to provide enough simulated volume to train effective CNN models. Instead, in this limited data regime, CNNs suffer from over-fitting and poor generalization. CNN models often address this issue through various forms of regularization. For example, [23] use an ensemble of models and varied noise realizations in their training simulations to mitigate over-fitting. [24] implement regularization in the form of dropout and an $\ell_2$ weight penalty, while [25] add a domain-adaptive penalty to the loss to improve generalization.

An alternative to these general regularization schemes is problem-specific, physically-informed regularization. Using assumptions about the data, known as inductive biases, models are constrained to learn more optimal solutions, helping prevent noise-dominated learned solutions that generalize poorly. The large-scale structure

* m.craigie@uq.edu.au



The comparison deliberately limits training data while keeping a large separate test set.



is well-suited for inductive biases due to firm assumptions from the cosmological principle: global rotational and translational symmetry.

The Wavelet Scattering Transform (WST) method incorporates these inductive biases in a CNN-like architecture, providing fixed convolution filters and processing with rotational and translational symmetry. The WST has emerged as a powerful alternative to two-point approaches, with wide application to cosmology inference [26–28]. However, while the WST's architecture is highly robust, its use of fixed wavelet filters means it is susceptible to missing additional useful information, since it lacks the flexibility to adapt to a specific field.

To introduce flexibility, [29] use a WST architecture with learnable wavelet parameters and see an improvement over the standard WST for image classification. Similarly, [30] develop Learnable Wavelet Scattering Networks that optimize wavelet filter parameters using genetic algorithms, improving fault diagnosis tasks in engineering domains.

Extending this flexibility further, in [31] we introduce the Neural Field Scattering Transform (NFST) which adds trainable convolution filters to the WST's architecture. Crucially, the NFST parametrizes the filters using a neural field, which promotes more smooth and stable convolution filters than pixel-wise kernels. The NFST's architectural inductive biases and filter smoothness enhance its data-efficiency, while its flexibility improves information capture. Together, these provide a balanced approach that improves over both the WST and a CNN model for detecting parity violation in large-scale structure-like data [31]. This previous success of the NFST motivates its application to different fields and training objectives.

In this work, we apply the NFST to constrain cosmology using weak lensing as a tracer of the large-scale structure. In section II we introduce the CosmoGrid weak lensing dataset, describe the summary statistics in greater detail and explain our methodology. In section III we show the results, and visualize the differences between a trained NFST and the WST. We place these results in context in IV, before concluding in section V.

## II. DATA & METHODOLOGY

### A. Data

For large-scale structure data, we use CosmoGrid, a suite of simulations designed for cosmology inference applications that includes 2500 cosmological N-body simulations run across varied cosmology parameters, with the primary grid simulations varying over $\Omega_m$, $\sigma_8$, $w$, $H_0$, $n_s$ and $\Omega_b$. CosmoGrid provides noiseless full-sky weak lensing convergence maps across 4 redshift slices. The maps are provided in healpix format with `nside=1024`, and have been sampled to correspond to Dark Energy Survey (DES) tomographic bins. The simulations incorporate baryonic feedback effects through a post-processing baryonification model that modifies the dark matter-only density distribution to account for baryonic physics effects on the matter field structure.

To reflect realistic constraints on simulation resources, we intentionally conduct our analysis in a limited data regime. In practical cosmology applications, high-fidelity simulations capture the rich non-Gaussian structure more accurately. However, increasing the fidelity of simulations often requires reducing the total number of realizations due to the heavy computational cost of accurately simulating small scale structure. Therefore, we aim to identify summary statistics that remain effective when trained on a small number of simulations. For this analysis, we choose a subset of 500 cosmologies for training (400) and validation (100), selected at random. The remaining 2000 are used as a test set to benchmark each summary statistic's performance.

We project the Healpix maps into flat $10° \times 10°$ patches using a Cartesian projection. This patch size means projection distortions are limited to sub-percent in area, and distortions manifest independent of cosmology and should therefore have minimal effect in a cosmology analysis. The resulting fields have a pixel size corresponding to $\sim 4$ arcmin, approximately matching the resolution limit of the healpix maps.

Each cosmology contributes ten map patches, totaling about 1,000 square degrees.

Although CosmoGrid provides whole-sky convergence maps for each cosmology, we choose to use 10 of the 100 deg$^2$ patches, to approximate the survey size of a $\sim 1000$ deg$^2$ photometric survey such as the Kilo-Degree Survey (KiDS, [32]). Increasing area (a possibility with future surveys such as Euclid [33] and LSST [34]) would have the effect of reducing the noise on each summary statistic. However, given that this data uses a perfect, systematic-free and noiseless reconstruction of the weak lensing field, the resulting constraint precision should only be interpreted for relative comparison purposes, not as an indication of the potential recovered precision in a observational survey.

### B. Summary Statistics

We test the ability of three summary statistics to extract cosmologically useful information from the weak lensing fields. Each summary statistic outputs a different set of coefficients that contain information about physical aspects of the field.

#### 1. *The Wavelet Scattering Transform*

The standard Wavelet Scattering Transform (WST) extracts information from homogeneous fields by applying convolution filtering across a set of complex wavelet filters with different scales and orientations. The method convolves the input field $\delta$ with filters $\psi_{jl}$ (indexed by scale $j$ and angle $l$) and computes scattering

cosmologies

400 train | 100 validation | 2,000 test

NFST keeps the scattering architecture and lets smooth filters adapt to the task.



coefficients via spatial averaging:

$$S^1_{jl} = \langle |\delta \star \psi_{jl}| \rangle \tag{1}$$

$$S^2_{j_1 l_1 j_2 l_2} = \langle ||\delta \star \psi_{j_1 l_1}| \star \psi_{j_2 l_2}| \rangle \tag{2}$$

where $\star$ denotes convolution and $\langle\rangle$ denotes spatial averaging.

For an alternative intuition, we can consider convolution filtering in the Fourier domain, where it is simply the product of the Fourier input field $\hat{\delta}$ and the Fourier filter $\hat{\psi}_{jl}$. In this view, the complex wavelets are instead smooth, directional band pass filters. The first pass of convolution isolates structure of the first filter's scale, while the second pass of convolution captures how the first structure clusters on the second filter's scale.

By averaging in the spatial domain, the first order coefficients measure the amount of structure of scale $j$, while second order coefficients measure the clustering (of scale $j_2$) of structures (of scale $j_1$). Second-order coefficients with $j_1 < j_2$ are not computed, with the physical intuition that structures cannot cluster on scales smaller than themselves. For first order coefficients, we average over angle index $l$ to preserve rotational symmetry. For second order coefficients, we apply an angular averaging scheme that preserves rotational symmetry while retaining cosmologically relevant relative angle information by averaging over angular differences $\Delta_\ell = |l_1 - l_2|$ rather than individual angles. Through these physically-motivated symmetry assumptions, the WST produces a compact set of coefficients representing the clustering information in the field.

For our tests, we use $J = 3$ scales (maximum scale $\sim 30$ arcmin) and $L = 6$ angles, giving 4 $\Delta_\ell$ coefficients for second-order terms. Although going to larger scales would likely tighten cosmology constraints, this information starts to be dominated by linear structure, which can be completely characterized by the power spectrum, and capturing this information is therefore not the goal of this test.

### 2. The Neural Field Scattering Transform

The Neural Field Scattering Transform (NFST) extends the WST by replacing fixed Morlet wavelet filters with trainable neural field filters [31]. The filters are parameterized as $\hat{\psi}_{jl}(\vec{k}) = F_{jl}(\vec{k}, j)$, where $F_{jl}$ is a neural network that maps Fourier-space coordinates $\vec{k}$ to filter values. Filters learn directly in Fourier space since in practice the convolutions occur in this domain for computational efficiency. Different scale filters use the same neural network, but providing the scale index $j$ as an additional input enables some filter variation across scales. We apply Fourier space truncation such that the $j$th filter is cut at frequencies above $N/2^j$ (where $N$ is the field size), ensuring each filter processes different scales while providing computational speedup. The neural field parameterization allows trivial dilation and rotation operations to readily construct a full filter set, and additionally promotes filter smoothness in Fourier space leading to more stable information extraction.

The NFST uses the same scattering architecture as the WST but with optimized filters that adapt to specific tasks. This combination of robust symmetry-exploiting architecture with flexible trainable components excels in scenarios with limited data but high complexity requirements. We provide full implementation details in [31]. For this work, we use a neural field network with 64 neurons, initialized to Morlet wavelets. The NFST model requires more computational time to train its filters than a standard WST. However, efficient GPU implementations and Fourier space truncation speedups means the NFST converges in $< 30$ minutes for this dataset on a single GPU. The result of this additional training time is a set of coefficients with equivalent scale and angle properties as the WST that are richer in task-specific information.

### 3. Convolutional Neural Network Compression

Scale and relative-angle structure remain part of the representation.

Convolutional Neural Networks (CNNs) are a popular class of deep learning architectures that extract spatial features from 2D images by convolving the input with a series of learnable filters. Unlike the WST and NFST, which incorporate symmetry constraints explicitly, CNNs learn their features purely from data, allowing them greater flexibility but at the cost of reduced robustness.

CNNs are distinct from the WST and NFST summary statistics in that they do not have explicitly defined output coefficients. However, CNNs usually construct compact internal representations of the field, which act as pseudo-summary statistics. To provide an equivalent comparison, we construct the CNN to have an output dimension equivalent to the number of coefficients in the WST and NFST models, effectively producing a set of CNN coefficients.

The CNN architecture begins with a three-stage convolutional block, consisting of a 3×3 convolution with 32 output channels, stride 1, and padding 1, followed by ReLU activation, 2×2 max pooling, and dropout regularization. The resulting feature maps are flattened and passed through a two-layer fully connected latent embedding, comprising a ReLU-activated linear layer followed by dropout and another linear projection to an output vector that matches the dimension of the WST and NFST outputs, of size 16.

This CNN architecture represents a standard baseline that is appropriate for the data volume and complexity of our cosmological inference task. While more sophisticated CNN architectures exist, they would likely suffer from overfitting given our dataset size. By using a relatively simple CNN architecture and channel-based and feature-based dropout regularization, we make the CNN more robust to over-fitting, thereby providing a fairer comparison to the WST and NFST models.

WST: fixed filters → scattering coefficients

NFST: learned filters → scattering coefficients

Freeze the learned summaries before fitting parameter predictions and posteriors.

For each statistic, we average computed coefficients over the ten field patches per cosmology, resulting in one set of 16 coefficients per redshift bin, for a total of four sets of coefficients. We concatenate these and normalize across cosmologies to provide the final set of 64 coefficients per cosmology for each summary statistic.

Four redshift bins supply 16 coefficients each: 64 inputs in total.

Together, these three summary statistics provide a spectrum of balance between informative power and incorporated symmetries, providing increasing flexibility and decreasing robustness in the order they are listed. In these tests, we do not compare to an explicit power spectrum summary statistic, since several previous studies have already demonstrated that the WST's capabilities exceed the power spectrum for capturing information in the non-linear large scale structure (e.g. [27, 35]).

### C. Cosmology Inference

We benchmark the performance of the three summary statistic models for predicting cosmology parameters from the large-scale structure. We train each summary statistic model to predict $\Omega_m$, $\sigma_8$ and $w$ from the CosmoGrid weak lensing datasets, demonstrating their ability to constrain $w$CDM cosmology. The other cosmology and astrophysical parameters varied in CosmoGrid are unconstrained by the model, and are therefore automatically marginalized over.

The full posterior modeling process for each summary statistic comprises three stages:

1. Summary Statistic Pre-Training: Use raw lensing fields to predict cosmology parameters with the NFST and CNN models. Then, freeze the models' trainable parameters such that they behave deterministic summary statistics rather than trainable models, and compute the corresponding set of coefficients for each.

2. Direct Parameter Prediction: Use explicit coefficients to directly predict cosmology parameters.

3. Posterior Modeling: Use predicted cosmology parameters to construct a posterior probability density model.

#### 1. *Summary Statistic Pre-Training*

The NFST and CNN models contain trainable parameters that are tuned throughout training, while the WST model is fixed and deterministic. To treat each summary statistic equally in the downstream analysis, we choose to pre-train the NFST and CNN summary statistics and freeze their trainable components. Each model then acts as a fixed function that produces an equally-sized set of coefficients, rather than an end-to-end trainable model.

Another advantage of this approach is the ability to use a different training setup for each phase. In the pre-training phase, we train each model for fewer epochs but at a higher learning rate to quickly learn useful feature representations. Then, in the downstream training phase, we train the final prediction model for more epochs at a lower learning rate to fine-tune the cosmological parameter predictions. We find that this two-stage approach consistently achieves higher performance with significantly lower compute time requirements compared to end-to-end training.

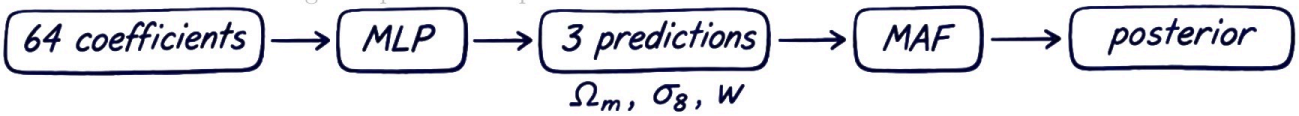
For pre-training, we connect each model's output coefficients to a basic Multi-Layer Perceptron (MLP) neural network with an output for each cosmology parameter. We then train the models to predict cosmology by minimizing the mean-squared error loss between predictions and target cosmology parameters. The MLP model consists of two linear layers with 512 nodes and ReLU activation, and a final linear layer feeding to an unconstrained, 3-dimensional output corresponding to the cosmology parameters. We pre-train for 300 and 500 training epochs (iterations over the full training dataset) for the NFST and CNN models respectively, with an Adam optimizer using a learning rate of $10^{-3}$. For each model, the chosen number of training epochs is sufficient to achieve convergence, as determined from the train and validation loss curves. For each, we use early stopping to ensure generalizable solutions. After pre-training, we freeze the trainable components in each model and apply the models to each weak lensing field, computing their corresponding set of 64 coefficients each.

#### 2. *Direct Parameter Prediction*

We next train a regression model to predict the respective cosmology parameters from each of the summary statistics' coefficient sets, again by minimizing the mean-squared error loss. We use the same MLP model and training setup as with pre-training, but instead train for 2000 epochs with a learning rate of $10^{-4}$.

After training, we test performance by computing the Root Mean Squared Error (RMSE) between predicted cosmology and true cosmology in the test set, providing a measure of prediction accuracy in physical units. A lower RMSE indicates that the summary statistic can provide tighter constraints on that cosmology parameter.

Deep learning models have inherent stochasticity which introduces variability to test performance. This variability is amplified in scenarios with limited training data, where over-fitting is more prominent. Since we are interested in assessing the summary statistics' performance, we repeat each test 10 times to confirm that any improvement is systematic and not a stochastic artifact.



NFST improves the tested sigma-eight and dark-energy predictions, with no gain in matter density.



*3. Posterior Modeling*

While predicting cosmology parameters directly is a useful indicator of summary statistic constraining power, it fails to capture the uncertainty in predictions and any degeneracies between parameters. To capture these, we instead use Neural Posterior Estimation (NPE). NPE is an approach where a neural model learns to approximate the posterior distribution of parameters given observed data, without requiring explicit likelihood evaluation. NPE trains a neural density estimator to directly model $P(\theta|x)$ by learning from pairs of simulated parameters $\theta$ and corresponding observations $x$. It works by training a conditional density estimator to learn the mapping from data to posterior distributions. Once trained, we can sample from the NPE model to provide full posterior estimates for new observations, capturing both parameter uncertainties and degeneracies that point estimates cannot reveal [36]. As observations $x$ we use the predicted cosmology parameters from the previous direct parameter estimation. Using these highly-compressed parameter predictions removes the compression task from the density model, which we find results in more accurate posteriors than using the raw coefficients directly.

To fit the NPE model, we use the `sbi` python package [37] and opt to use a Masked Autoregressive Flow (MAF, [38]) as a density model. For this setup, we perform a complete grid search over `hidden_features` from [16, 32, 64, 128] and `num_transforms` from [2, 4, 6, 8, 10, 12], finding that `hidden_features=64` and `num_transforms=8` produce the most accurate posterior. We emphasize that we are interested in the relative performance between the models under different input coefficients, so optimizing the posterior model further is secondary to ensuring consistency between tests. Within `sbi`, we use an uninformative flat prior over $[-5, 5]$ for each of the normalized parameters. CosmoGrid's finite cosmology parameter space sampling imposes an effective prior, but this is equivalent across all summary statistics, and all test cosmologies fall within this prior.

We assess the effectiveness of each summary statistics' learned posterior by measuring their quality over the test set. As a metric, we compute the mean log probability across all data and parameter pairs in the test set. Higher values indicate the model assigns greater probability density to the true parameter values, with differences reflecting the relative quality of parameter estimation. This means that for higher values, the corresponding summary statistic provides a more accurate and trustworthy estimate of the cosmology parameters and their confidence. The 2000 data-parameter pairs form a sufficiently large test set to use mean log probability for assessing performance [39]. Again, we report the results for 10 repeats to separate systematic improvement from training stochasticity.

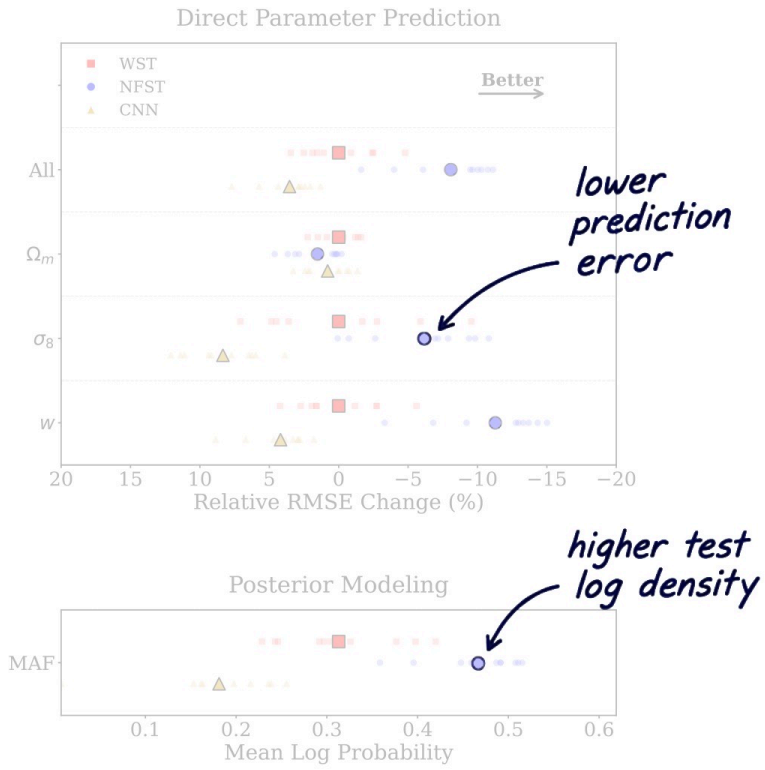



FIG. 1. Comparison of the WST (red squares), NFST (blue circles) and CNN (golden triangles) summary statistics for constraining cosmology. Large shapes with black outlines indicate the average performance across 10 training runs, while smaller shapes indicate individual training runs. **Top:** Relative summary statistic performance for direct cosmology parameter prediction. We measure performance as the Root Mean-Squared Error (RMSE), relative to the average RMSE of the WST. **Bottom:** Mean log probability of a full posterior model that uses a Masked Autoregressive Flow (MAF) density estimator.

## III. RESULTS

### A. Cosmology Inference Models

To visualize the ability of the different summary statistics to capture cosmological information, in the top panel of Figure 1 we show the corresponding change in RMSE as percentages relative to the average RMSE of the WST model. To extend this to posterior estimation, in the bottom panel of Figure 1 we show the mean log probability for the posterior model built from each summary statistic. In both cases, we show the average performance across all ten repeats, as well as the individual runs to illustrate the model convergence variability. The model performance is often spread in a skewed or bimodal manner, but the average still provides a reasonable summary of the performance tendency of each model.

In the direct parameter estimation test, the NFST consistently improves over the the WST. The NFST av-

Large outlined markers average ten runs; small markers show individual runs.

The learned Fourier filters broaden the range of angular modes being combined.



erages a 6% decrease in prediction RMSE for $\sigma_8$, and a 11% decrease in prediction RMSE for $w$. This improvement extends to posterior modeling, with consistently higher mean log probability over the unseen test data. Across the ten runs, the NFST averages a mean log probability of 0.47, compared to 0.31 for the WST. This corresponds to a 17% increase in the average probability density assigned to each parameter-cosmology pair in the test dataset, which is a substantial improvement in posterior quality. These results show that the NFST recovers additional information compared to the WST from the small scales of the weak lensing convergence field. We see no improvement in $\Omega_m$, indicating that the WST already captures the majority of the field's information related to matter density.

In contrast, the CNN summary statistic shows a consistent decrease in performance in all three cosmology parameters. Likewise, for posterior modeling, the CNN generally provides less accurate posterior constraints over the test set, equivalent to a 14% decrease in average assigned probability density over the test set. This result indicates that the CNN struggles to capture information in this limited data regime, resulting in overfitting and poor generalization to test data.

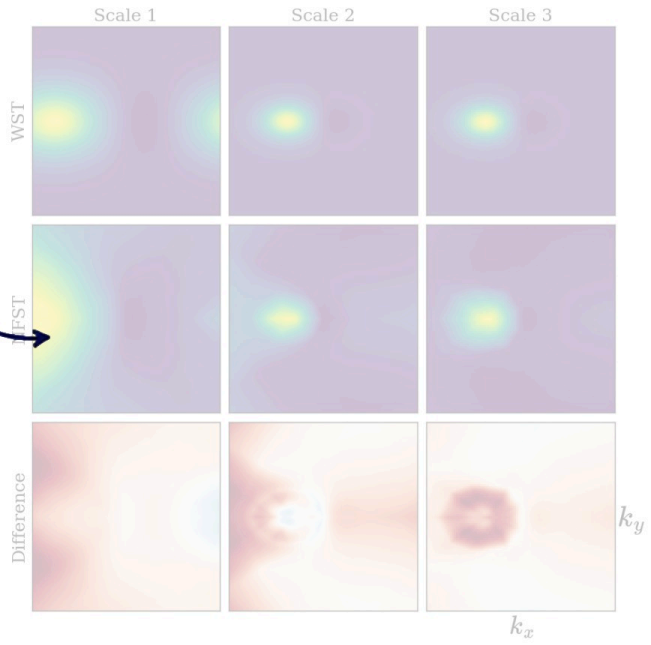



FIG. 2. The first angular filter of the WST (top row) and trained NFST (middle row) shown in Fourier space (with coordinates $k_x$ and $k_y$), across the three scales used in the cosmology prediction. The middle of each image is the origin, with frequency increasing outwards. In the WST and NFST rows, brighter color indicates higher filter values. In the difference plot, red regions indicate an increase in the NFST's filters relative to the WST, while blue indicates a decrease. Scale 2 and Scale 3 appear visually similar despite representing different scales because the Scale 3 filter's high frequency modes are truncated for computational efficiency.

### B. Insights from NFST Filter Visualization

A strength of the NFST coefficients is their direct physical meaning, with each coefficient representing the strength of features in the field with corresponding scale $j$ and angle $l$ (or relative angle $\Delta_l$ for second order coefficients) of the filters. In physical units, the three scales correspond to approximately 8, 16 and 32 arcmin respectively, which manifest as different physical scales for each redshift bin.

In Figure 2 we visualize the differences between the learned filters and the standard WST's Morlet filters, to better understand the information that the NFST captures.

The WST's Morlet filters' structure in Fourier space is a Gaussian bump centered on a nonzero frequency, with a smaller negative Gaussian bump subtracted at zero frequency to ensure the filter integrates to unity. The angle of this bump determines the associated angle in the filter. As a result, the Morlet filter acts as a smooth, directional band-pass filter. In each successive scale, the filter's highest frequency Fourier modes are removed, so successive filters capture larger scales. The first Morlet filter pushes up against the high-frequency edge, allowing it to capture the most small-scale information for each scale. To accommodate this, the Gaussian bump wraps to the other side to keep the wavelet smooth in configuration space.

The NFST's first filter, representing the smallest scales, exhibits a clear widening of the frequency pass band. This broader structure enables the downstream scattering transform operations to average over a greater number of high-frequency modes, likely improving the signal-to-noise ratio of the resulting coefficients. However, this comes at the cost of reduced precision on directional information. The preference for wide filters suggests that maximizing cosmological constraining power benefits more from averaging across broader frequency ranges than from maintaining precise angular resolution. The second and third filters both show a similar widening effect, though less pronounced. The third filter also displays a reshaped band-pass bump, shifting towards a sharper cutoff along both scale and angles.

### C. Insights from NFST Coefficient Visualization

Although studying the NFST filters can provide insights into the first order coefficients and general areas of improvement, second order coefficients are more complex to understand. To help provide insights into the higher-order information that these coefficients access, we develop a method to directly visualize the information that the summary statistic captures. We then apply this to the NFST to visualize where the improvement over the WST originates from in physical, field-level space.

These are Fourier-space filters, not maps of galaxy density.

Optimize a synthetic field to reveal what a chosen coefficient responds to.



### 1. Summary Statistic-Based Field Maximization

This new method for visualizing summary statistics relies on maximizing their value in a field. We expand upon work from [40], which visualizes the information that an arbitrary summary statistic captures from a target field. It works by minimizing the difference between the statistics evaluated over the target field and over a learned field. The minimization uses gradient descent, and can therefore leverage efficient deep learning framework computations and GPU acceleration. Building on this framework, we develop a new method that instead maximizes an individual summary coefficient over a learned field, while suppressing all other summary coefficients. This provides a clear visual representation of the information probed by that specific summary statistic.

Let $F(\theta)$ be the output field generated by a neural model with parameters $\theta$, and let $S(\cdot)$ be a vector-valued summary statistic function producing $k$ normalized components. We define normalization relative to a fixed baseline summary $S_{\text{base}}$, computed from a reference distribution, such that:

$$\hat{S}_i(\theta) = \frac{S(F(\theta))_i}{S_{i,\text{base}} + \varepsilon} \tag{3}$$

where $\varepsilon$ is a small constant ensuring numerical stability. For a neutral baseline, we use uncorrelated standard normal noise. This normalization helps keep targets within the same order of magnitude, whereas each component of the summary statistic could be at any value.

Let $I_C$ be the indices corresponding to the coefficients we wish to constrain towards 1 (thereby suppressing their characteristics in the field to the noise baseline), and $I_B$ for those we wish to boost beyond 1 (thereby enhancing their characteristics). We define the constraint potential $P_C(\theta)$ and the boost potential $P_B(\theta)$ as

$$P_C(\theta) = \left\langle \left| \hat{S}_i(\theta) - 1 \right| \right\rangle_{i \in I_C} \tag{4}$$

and

$$P_B(\theta) = \left\langle \operatorname{sgn}\left(\hat{S}_j(\theta) - 1\right) \cdot \left| \hat{S}_j(\theta) - 1 \right| \right\rangle_{j \in I_B} \tag{5}$$

where $\langle \cdot \rangle$ indicates an average over the respective set of indices. The use of the signed difference in $P_B$ prevents pushing towards highly negative values throughout optimization.

The total loss is defined as a weighted difference:

$$\mathcal{L}(\theta) = \lambda_C \cdot P_C(\theta) - P_B(\theta) \tag{6}$$

where $\lambda_C \geq 0$ is a scalar parameter controlling the relative weight on the constraint potential. This loss penalizes components in $I_C$ for deviating from the neutral reference value of 1, and rewards components in $I_B$ for increasing above 1 and penalizes them for falling below it. Setting $\lambda_C = 0$ promotes unconstrained maximization of $I_B$, which often results in non-convergent loss, but can still lead to useful visualizations after a moderate training period.

Additionally, supplying a specific reference field and setting $I_C$ such that all indices are constrained is approximately equivalent to the method described in [40], providing a visualization of how a given summary statistic captures the information in a specific field.

We also provide a modification to the trainable field. In [40], $F(\theta)$ is directly parameterized at the pixel level, such that each pixel is a trainable parameter. While using this approach in our new maximization framework proves effective for some summary statistics, for others it can result in convergence stuck at local minima, providing an incomplete representation of summary statistic properties.

The field is optimized for interpretation; this is not a new observed sky map.

To improve the optimization, we propose to parameterize $F(\theta)$ more flexibly with a dual-domain approach that operates in both real and Fourier space:

$$F(\theta) = w_x \cdot [\mathcal{C}_x(z_x) + f_x] + w_k \cdot \mathcal{F}^{-1}[\mathcal{C}_k(z_k) + f_k] \tag{7}$$

where $\mathcal{C}_x$ and $\mathcal{C}_k$ are CNN models (we use U-Net architectures) that operate on learned latent spaces $z_x$ and $z_k$ respectively, $f_x$ and $f_k$ are additional learned fields in real and Fourier space, $w_x$ and $w_k$ are learnable weights controlling the relative contributions, and $\mathcal{F}^{-1}$ denotes the inverse Fourier transform.

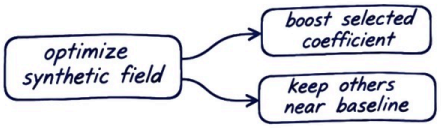
This dual parameterization provides complementary advantages: the real-space component $\mathcal{C}_x(z_x) + f_x$ excels at capturing localized features and sharp structures, while the Fourier-space component $\mathcal{F}^{-1}[\mathcal{C}_k(z_k)+f_k]$ naturally captures global patterns, periodic structures, and long-range correlations. The learnable weights $w_x$ and $w_k$ allow the optimization to automatically balance between these complementary representations based on the specific requirements of the target summary statistic.

This new parameterization is far more expressive than direct pixel-level optimization and has greatly improved convergence properties. The CNN formulation provides significant flexibility in both domains and can be made arbitrarily complex to provide better field statistic visualizations, at the cost of greater computation time.

Our method shares conceptual similarities with activation maximization techniques (e.g. [41–43]), as both optimize a trainable input to enhance specific target outputs. While activation maximization seeks to amplify activations of particular neurons or layers exclusively within a neural network, our method instead generalizes beyond neural networks to visualize arbitrary summary statistics.

### 2. NFST Coefficient Visualization

Using this new framework, we compare the NFST coefficient with the most influence on cosmology to its



The advantage is promising, but noise, survey effects and three-dimensional data remain tests ahead.



WST counterpart. To identify the most influential coefficients, we use the Shapley Additive Explanation (SHAP) framework, which explains the importance of a model's inputs (summary statistic coefficients) to its corresponding outputs (predicted cosmology parameter) [44].

The coefficient with the highest overall SHAP importance is the second order coefficient that corresponds to $j_1 = 1$, $j_2 = 2$ and $\Delta_\ell = 0$, which physically represents how much $\sim 4$ Mpc sub-structures tend to cluster on scales of $\sim 16$ Mpc, where the sub-structures and their clustering align in the same direction. This coefficient has the highest SHAP value for both $\sigma_8$ and $w$. Interestingly, this coefficient represents the coupling between the two smallest scales probed by the NFST, indicating that the most relevant scales to cosmology are those that are deepest into the non-linear regime.

SHAP ranks influence on this model; it does not prove a unique physical cause.

To visualize the information held by this coefficient, we show in Figure 3 the coefficient maximized for both the NFST and WST using the gradient descent technique discussed above. The visualization clearly shows that the NFST's learned filter capture a less rigid structure, reflecting the NFST's reduced angular precision in both the first and second order coefficients. This visualization presents a configuration space representation of the finding that the NFST aggregates information over more angular modes of information than its WST counterpart, complementing the Fourier space representation identified from the NFST's filters. The less rigid structures captured by the NFST more accurately reflect the chaotic and disorganized nature of the large-scale structure, and bears a stronger resemblance to filamentary structure.

Beyond this specific application to the NFST, we demonstrate that this new method generally provides a useful and physically interpretable visualization of field summary statistics. Even for highly complex statistics with computations in Fourier space like the NFST, this approach can show which information the statistic is probing in physical space.

## IV. DISCUSSION

The NFST's performance reveals a fundamental trade-off in cosmological inference: while maximum flexibility (as in CNNs) allows capturing arbitrary patterns, it requires prohibitively large training datasets in cosmology where simulations are computationally expensive. Conversely, fixed summary statistics like the WST impose strong prior assumptions that may discard cosmologically relevant information. The NFST demonstrates that incorporating physics-motivated constraints (translational and rotational symmetries) while allowing smooth task-specific filter adaptation provides more effective regularization than general deep learning techniques in data-limited regimes.

The learned filters' broader frequency pass bands suggest that cosmological information at small and intermediate scales is better extracted through statistical averaging across wide angular bands rather than precise spatial localization. This is a strong indication that fine angular resolution should not be prioritized in wavelet-based cosmology inference, instead referencing the decrease in statistical variance that comes through averaging over more modes.

Systematics such as survey masking and distance uncertainties can have a strong influence on cosmology inference accuracy with deep learning models [45]. Furthermore, when moving towards observational datasets, the addition of noise can dampen non-Gaussian information, potentially reducing the information gain from the NFST. This contamination may exacerbate the need for wider filters to constrain statistical noise. The full effects of systematics and noise on the NFST's training procedure and inference accuracy warrants further investigation.

Survey noise and masking may reduce the gains seen in these idealized maps.

While this work focuses on understanding how the NFST's captured information contrasts with the WST, future work can also explore a full interpretability analysis to identify the specific redshift bins and structural information most relevant to predicting each cosmology parameter. Such an analysis can provide unique insights into the connections between the large-scale structure and cosmology parameters.

Future work can also explore the further application of the NFST to 3D large-scale structure, such as the full 3D galaxy distribution. Applications of CNNs to 3D LSS data have shown some success, but they rely on heavy regularization to avoid overfitting [24]. Following its demonstrated robustness to over-fitting in 2D scenarios, the NFST may also extend well to 3D. With its physically-motivated inductive biases and smooth neural field filter representation, the NFST may provide the necessary advantage to flexibly capture 3D information from the field, without requiring excessive simulated training datasets. Such an extension to the NFST would require adaptation to match the line-of-sight asymmetry in observational large-scale structure, but could prove highly beneficial to improving cosmology constraints or better understanding the non-Gaussian large-scale structure.

## V. CONCLUSION

We demonstrate the use of the Neural Field Scattering Transform (NFST) as a summary statistic for large-scale structure. The NFST's balance between flexibility and robustness lets it achieve more accurate and precise cosmology constraints for cosmological inference compared to previous methods. Through trainable filters, the NFST provides an improvement in informative capacity relative to a standard Wavelet Scattering Transform (WST) model. For direct parameter prediction, the NFST provides a 6% improvement in prediction RMSE for $\sigma_8$ and an 11% improvement for $w$. For posterior modeling, the NFST provides a 17% improvement

Optimized fields show the spatial patterns rewarded by each scattering coefficient.



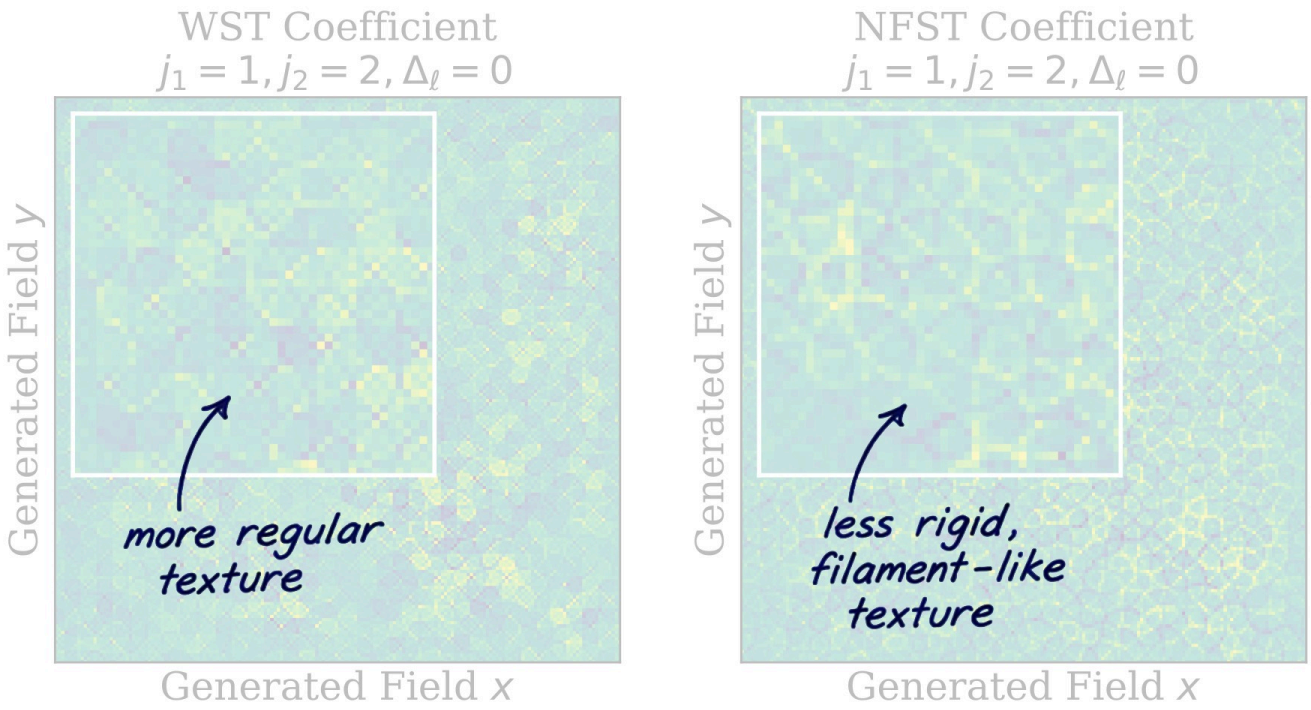



FIG. 3. Generated fields representing the information probed by a selected second order scattering coefficient for the WST (left) and NFST (right) models. The white insets show a zoomed-in portion of the field. The selected coefficient is the most cosmologically-informative coefficient for the NFST. The generated fields show that the NFST extracts more filament-like information to better constrain $\sigma_8$ and $w$. We produce each field using the summary statistic visualization technique introduced in this work.

in posterior quality.

To maintain robust generalization even with limited training data, the NFST uses physically-motivated inductive biases, setting it apart from less physical Convolutional Neural Network (CNN) models. For direct parameter prediction, the NFST provides a 15% and 17% improvement in prediction RMSE for $\sigma_8$ and $w$ respectively, and a 33% improvement over the CNN for posterior modeling.

We also show the direct interpretability of the NFST by comparing it to the WST at the filter-level. We see that for the lensing field, Fourier-space filters with a broader angular spread are more effective for constraining cosmology. Physically, this suggests that for cosmology inference, directional information is less important than aggregation for reducing noise.

To enable visualization at the coefficient level, we introduce a new summary statistic visualization technique. We use this to demonstrate how the NFST differs from the WST, and confirm in physical space that the NFST captures broader angular information, which more closely resembles filamentary structure.

This new technique can also be applied beyond the NFST, to visualize other summary statistics. By enabling interpretability with this new method, we can verify the trustworthiness of deep learning-based statistics and provide physical insights into the data.

More broadly, the NFST's ability to adapt to the specific lensing field suggests a domain-adaptability that could successfully capture information from other homogeneous fields beyond the cosmological large-scale structure. This work also marks the NFST's first application to cosmology inference, highlighting its task-adaptability beyond earlier demonstrations for parity violation detection.

In future, we will apply the NFST to more realistic simulated data, testing its robustness to simulation artifacts and observational systematics, and ultimately apply it to weak lensing and clustering datasets from next-generation surveys. We also plan to extend the NFST to 3D, to efficiently and robustly capture richer 3D clustering information.

Overall, these findings demonstrate that flexible and robust models such as the NFST, combined with deep learning interpretability techniques, can improve the cosmological analysis process and deepen our understanding of large-scale structure.

## VI. ACKNOWLEDGEMENTS

M.C. acknowledges the support of an Australian Government Research Training Program (RTP) Scholarship. M.C., R.R., T.M.D., acknowledge the support of an Australian Research Council Australian Laureate Fellowship (FL180100168) funded by the Aus-

These are synthetic diagnostic fields, not observed cosmic structures.

tralian Government. Y.S.T is supported by the National Science Foundation under Grant No. AST-2406729. R.R. acknowledges financial support from the Australian Research Council through DECRA Fellowship DE240100816. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231.

[1] M. Dainotti, B. D. Simone, G. Montani, T. Schiavone, and G. Lambiase, "The hubble constant tension: current status and future perspectives through new cosmological probes," (2023), arXiv:2301.10572 [astro-ph.CO].

[2] R. Camilleri, T. M. Davis, M. Vincenzi, P. Shah, J. Frieman, R. Kessler, P. Armstrong, D. Brout, A. Carr, R. Chen, L. Galbany, K. Glazebrook, S. R. Hinton, J. Lee, C. Lidman, A. Möller, B. Popovic, H. Qu, M. Sako, D. Scolnic, M. Smith, M. Sullivan, B. O. Sánchez, G. Taylor, M. Toy, P. Wiseman, T. M. C. Abbott, M. Aguena, S. Allam, O. Alves, J. Annis, S. Avila, D. Bacon, E. Bertin, S. Bocquet, D. Brooks, D. L. Burke, A. Carnero Rosell, J. Carretero, F. J. Castander, L. N. da Costa, M. E. S. Pereira, S. Desai, H. T. Diehl, P. Doel, C. Doux, S. Everett, I. Ferrero, B. Flaugher, P. Fosalba, J. García-Bellido, M. Gatti, E. Gaztanaga, G. Giannini, D. Gruen, D. L. Hollowood, K. Honscheid, D. J. James, K. Kuehn, O. Lahav, S. Lee, G. F. Lewis, J. L. Marshall, J. Mena-Fernández, R. Miquel, J. Muir, J. Myles, R. L. C. Ogando, A. Pieres, A. A. P. Malagón, A. Porredon, M. Rodriguez-Monroy, E. Sanchez, D. Sanchez Cid, M. Schubnell, I. Sevilla-Noarbe, E. Suchyta, M. E. C. Swanson, G. Tarle, A. R. Walker, N. Weaverdyck, and DES Collaboration, Mon. Not. R. Astron. Soc. **533**, 2615 (2024), arXiv:2406.05048 [astro-ph.CO].

[3] DES Collaboration, T. M. C. Abbott, M. Acevedo, M. Aguena, A. Alarcon, S. Allam, O. Alves, A. Amon, F. Andrade-Oliveira, J. Annis, P. Armstrong, J. Asorey, S. Avila, D. Bacon, B. A. Bassett, K. Bechtol, P. H. Bernardinelli, G. M. Bernstein, E. Bertin, J. Blazek, S. Bocquet, D. Brooks, D. Brout, E. Buckley-Geer, D. L. Burke, H. Camacho, R. Camilleri, A. Campos, A. Carnero Rosell, D. Carollo, A. Carr, J. Carretero, F. J. Castander, R. Cawthon, C. Chang, R. Chen, A. Choi, C. Conselice, M. Costanzi, L. N. da Costa, M. Crocce, T. M. Davis, D. L. DePoy, S. Desai, H. T. Diehl, M. Dixon, S. Dodelson, P. Doel, C. Doux, A. Drlica-Wagner, J. Elvin-Poole, S. Everett, I. Ferrero, A. Ferté, B. Flaugher, R. J. Foley, P. Fosalba, D. Friedel, J. Frieman, C. Frohmaier, L. Galbany, J. García-Bellido, M. Gatti, E. Gaztanaga, G. Giannini, K. Glazebrook, O. Graur, D. Gruen, R. A. Gruendl, G. Gutierrez, W. G. Hartley, K. Herner, S. R. Hinton, D. L. Hollowood, K. Honscheid, D. Huterer, B. Jain, D. J. James, N. Jeffrey, E. Kasai, L. Kelsey, S. Kent, R. Kessler, A. G. Kim, R. P. Kirshner, E. Kovacs, K. Kuehn, O. Lahav, J. Lee, S. Lee, G. F. Lewis, T. S. Li, C. Lidman, H. Lin, U. Malik, J. L. Marshall, P. Martini, J. Mena-Fernández, F. Menanteau, R. Miquel, J. J. Mohr, J. Mould, J. Muir, A. Möller, E. Neilsen, R. C. Nichol, P. Nugent, R. L. C. Ogando, A. Palmese, Y. C. Pan, M. Paterno, W. J. Percival, M. E. S. Pereira, A. Pieres, A. A. P. Malagón, B. Popovic, A. Porredon, J. Prat, H. Qu, M. Raveri, M. Rodríguez-Monroy, A. K. Romer, A. Roodman, B. Rose, M. Sako, E. Sanchez, D. Sanchez Cid, M. Schubnell, D. Scolnic, I. Sevilla-Noarbe, P. Shah, J. A. Smith, M. Smith, M. Soares-Santos, E. Suchyta, M. Sullivan, N. Suntzeff, M. E. C. Swanson, B. O. Sánchez, G. Tarle, G. Taylor, D. Thomas, C. To, M. Toy, M. A. Troxel, B. E. Tucker, D. L. Tucker, S. A. Uddin, M. Vincenzi, A. R. Walker, N. Weaverdyck, R. H. Wechsler, J. Weller, W. Wester, P. Wiseman, M. Yamamoto, F. Yuan, B. Zhang, and Y. Zhang, Astrophys. J. Lett. **973**, L14 (2024), arXiv:2401.02929 [astro-ph.CO].

[4] V. Poulin, T. L. Smith, and T. Karwal, "The ups and downs of early dark energy solutions to the hubble tension: a review of models, hints and constraints circa 2023," (2023), arXiv:2302.09032 [astro-ph.CO].

[5] A. G. Adame, J. Aguilar, S. Ahlen, S. Alam, D. M. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud, S. Avila, A. Aviles, H. Awan, B. Bahr-Kalus, S. Bailey, C. Baltay, A. Bault, J. Behera, S. BenZvi, A. Bera, F. Beutler, D. Bianchi, C. Blake, R. Blum, S. Brieden, A. Brodzeller, D. Brooks, E. Buckley-Geer, E. Burtin, R. Calderon, R. Canning, A. Carnero Rosell, R. Cereskaite, J. L. Cervantes-Cota, S. Chabanier, E. Chaussidon, J. Chaves-Montero, S. Chen, X. Chen, T. Claybaugh, S. Cole, A. Cuceu, T. M. Davis, K. Dawson, A. de la Macorra, A. de Mattia, N. Deiosso, A. Dey, B. Dey, Z. Ding, P. Doel, J. Edelstein, S. Eftekharzadeh, D. J. Eisenstein, A. Elliott, P. Fagrelius, K. Fanning, S. Ferraro, J. Ereza, N. Findlay, B. Flaugher, A. Font-Ribera, D. Forero-Sánchez, J. E. Forero-Romero, C. S. Frenk, C. Garcia-Quintero, E. Gaztañaga, H. Gil-Marín, S. Gontcho a Gontcho, A. X. Gonzalez-Morales, V. Gonzalez-Perez, C. Gordon, D. Green, D. Gruen, R. Gsponer, G. Gutierrez, J. Guy, B. Hadzhiyska, C. Hahn, M. M. S. Hanif, H. K. Herrera-Alcantar, K. Honscheid, C. Howlett, D. Huterer, V. Iršič, M. Ishak, S. Juneau, N. G. Karaçaylı, R. Kehoe, S. Kent, D. Kirkby, A. Kremin, A. Krolewski, Y. Lai, T. W. Lan, M. Landriau, D. Lang, J. Lasker, J. M. Le Goff, L. Le Guillou, A. Leauthaud, M. E. Levi, T. S. Li, E. Linder, K. Lodha, C. Magneville, M. Manera, D. Margala, P. Martini, M. Maus, P. McDonald, L. Medina-Varela, A. Meisner, J. Mena-Fernández, R. Miquel, J. Moon, S. Moore, J. Moustakas, E. Mueller, A. Muñoz-Gutiérrez, A. D. Myers, S. Nadathur, L. Napolitano, R. Neveux, J. A. Newman, N. M. Nguyen, J. Nie, G. Niz, H. E. Noriega, N. Padmanabhan, E. Paillas, N. Palanque-Delabrouille, J. Pan, S. Penmetsa, W. J. Percival, M. M. Pieri, M. Pinon, C. Poppett, A. Porredon, F. Prada, A. Pérez-Fernández, I. Pérez-Ràfols, D. Rabinowitz, A. Raichoor, C. Ramírez-Pérez, S. Ramirez-Solano, M. Rashkovetskyi, C. Ravoux, M. Rezaie, J. Rich, A. Rocher, C. Rockosi, N. A. Roe, A. Rosado-Marin,

A. J. Ross, G. Rossi, R. Ruggeri, V. Ruhlmann-Kleider, L. Samushia, E. Sanchez, C. Saulder, E. F. Schlafly, D. Schlegel, M. Schubnell, H. Seo, A. Shafieloo, R. Sharples, J. Silber, A. Slosar, A. Smith, D. Sprayberry, T. Tan, G. Tarlé, P. Taylor, S. Trusov, L. A. Ureña-López, R. Vaisakh, D. Valcin, F. Valdes, M. Vargas-Magaña, L. Verde, M. Walther, B. Wang, M. S. Wang, B. A. Weaver, N. Weaverdyck, R. H. Wechsler, D. H. Weinberg, M. White, J. Yu, Y. Yu, S. Yuan, C. Yèche, E. A. Zaborowski, P. Zarrouk, H. Zhang, C. Zhao, R. Zhao, R. Zhou, and T. Zhuang, J. Cosmol. Astropart. Phys. **2025**, 021 (2025), arXiv:2404.03002 [astro-ph.CO].

[6] S. N. Masot, H. Gil-Marín, L. Verde, J. Aguilar, S. Ahlen, S. Bailey, S. BenZvi, D. Bianchi, D. Brooks, E. Buckley-Geer, A. C. Rosell, E. Chaussidon, T. Claybaugh, S. Cole, A. Cuceu, K. S. Dawson, A. de la Macorra, R. Demina, A. Dey, B. Dey, P. Doel, S. Ferraro, A. Font-Ribera, J. E. Forero-Romero, E. Gaztañaga, S. G. A. Gontcho, A. X. Gonzalez-Morales, G. Gutierrez, H. K. Herrera-Alcantar, K. Honscheid, C. Howlett, S. Juneau, R. Kehoe, D. Kirkby, T. Kisner, A. Kremin, C. Lamman, M. Landriau, L. L. Guillou, M. E. Levi, C. Magneville, M. Manera, A. Meisner, R. Miquel, J. Moustakas, A. Muñoz-Gutiérrez, A. D. Myers, S. Nadathur, G. Niz, H. E. Noriega, W. J. Percival, C. Poppett, F. Prada, I. Pérez-Ràfols, A. J. Ross, G. Rossi, L. Samushia, E. Sanchez, D. Schlegel, M. Schubnell, H. Seo, J. Silber, D. Sprayberry, G. Tarlé, M. Vargas-Magaña, B. A. Weaver, P. Zarrouk, R. Zhou, and H. Zou, "Full-shape analysis of the power spectrum and bispectrum of desi dr1 lrg and qso samples," (2025), arXiv:2503.09714 [astro-ph.CO].

[7] L. Secco, S. Samuroff, E. Krause, B. Jain, J. Blazek, M. Raveri, A. Campos, A. Amon, A. Chen, C. Doux, A. Choi, D. Gruen, G. Bernstein, C. Chang, J. DeRose, J. Myles, A. Ferté, P. Lemos, D. Huterer, J. Prat, M. Troxel, N. MacCrann, A. Liddle, T. Kacprzak, X. Fang, C. Sánchez, S. Pandey, S. Dodelson, P. Chintalapati, K. Hoffmann, A. Alarcon, O. Alves, F. Andrade-Oliveira, E. Baxter, K. Bechtol, M. Becker, A. Brandao-Souza, H. Camacho, A. Carnero Rosell, M. Carrasco Kind, R. Cawthon, J. Cordero, M. Crocce, C. Davis, E. Di Valentino, A. Drlica-Wagner, K. Eckert, T. Eifler, M. Elidaiana, F. Elsner, J. Elvin-Poole, S. Everett, P. Fosalba, O. Friedrich, M. Gatti, G. Giannini, R. Gruendl, I. Harrison, W. Hartley, K. Herner, H. Huang, E. Huff, M. Jarvis, N. Jeffrey, N. Kuropatkin, P.-F. Leget, J. Muir, J. Mccullough, A. Navarro Alsina, Y. Omori, Y. Park, A. Porredon, R. Rollins, A. Roodman, R. Rosenfeld, A. Ross, E. Rykoff, J. Sanchez, I. Sevilla-Noarbe, E. Sheldon, T. Shin, A. Troja, I. Tutusaus, T. Varga, N. Weaverdyck, R. Wechsler, B. Yanny, B. Yin, Y. Zhang, J. Zuntz, T. Abbott, M. Aguena, S. Allam, J. Annis, D. Bacon, E. Bertin, S. Bhargava, S. Bridle, D. Brooks, E. Buckley-Geer, D. Burke, J. Carretero, M. Costanzi, L. da Costa, J. De Vicente, H. Diehl, J. Dietrich, P. Doel, I. Ferrero, B. Flaugher, J. Frieman, J. García-Bellido, E. Gaztanaga, D. Gerdes, T. Giannantonio, J. Gschwend, G. Gutierrez, S. Hinton, D. Hollowood, K. Honscheid, B. Hoyle, D. James, T. Jeltema, K. Kuehn, O. Lahav, M. Lima, H. Lin, M. Maia, J. Marshall, P. Martini, P. Melchior, F. Menanteau, R. Miquel, J. Mohr, R. Morgan, R. Ogando, A. Palmese, F. Paz-Chinchón, D. Petravick, A. Pieres, A. Plazas Malagón, M. Rodriguez-Monroy, A. Romer, E. Sanchez, V. Scarpine, M. Schubnell, D. Scolnic, S. Serrano, M. Smith, M. Soares-Santos, E. Suchyta, M. Swanson, G. Tarle, D. Thomas, and C. To, Physical Review D **105** (2022), 10.1103/physrevd.105.023515.

[8] M. Asgari, C.-A. Lin, B. Joachimi, B. Giblin, C. Heymans, H. Hildebrandt, A. Kannawadi, B. Stölzner, T. Tröster, J. L. van den Busch, A. H. Wright, M. Bilicki, C. Blake, J. de Jong, A. Dvornik, T. Erben, F. Getman, H. Hoekstra, F. Köhlinger, K. Kuijken, L. Miller, M. Radovich, P. Schneider, H. Shan, and E. Valentijn, Astronomy & Astrophysics **645**, A104 (2021).

[9] T. Hamana, M. Shirasaki, S. Miyazaki, C. Hikage, M. Oguri, S. More, R. Armstrong, A. Leauthaud, R. Mandelbaum, H. Miyatake, A. J. Nishizawa, M. Simet, M. Takada, H. Aihara, J. Bosch, Y. Komiyama, R. Lupton, H. Murayama, M. A. Strauss, and M. Tanaka, Publications of the Astronomical Society of Japan **72** (2020), 10.1093/pasj/psz138.

[10] T. Abbott, M. Aguena, A. Alarcon, S. Allam, O. Alves, A. Amon, F. Andrade-Oliveira, J. Annis, S. Avila, D. Bacon, E. Baxter, K. Bechtol, M. Becker, G. Bernstein, S. Bhargava, S. Birrer, J. Blazek, A. Brandao-Souza, S. Bridle, D. Brooks, E. Buckley-Geer, D. Burke, H. Camacho, A. Campos, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, F. Castander, R. Cawthon, C. Chang, A. Chen, R. Chen, A. Choi, C. Conselice, J. Cordero, M. Costanzi, M. Crocce, L. da Costa, M. da Silva Pereira, C. Davis, T. Davis, J. De Vicente, J. DeRose, S. Desai, E. Di Valentino, H. Diehl, J. Dietrich, S. Dodelson, P. Doel, C. Doux, A. Drlica-Wagner, K. Eckert, T. Eifler, F. Elsner, J. Elvin-Poole, S. Everett, A. Evrard, X. Fang, A. Farahi, E. Fernandez, I. Ferrero, A. Ferté, P. Fosalba, O. Friedrich, J. Frieman, J. García-Bellido, M. Gatti, E. Gaztanaga, D. Gerdes, T. Giannantonio, G. Giannini, D. Gruen, R. Gruendl, J. Gschwend, G. Gutierrez, I. Harrison, W. Hartley, K. Herner, S. Hinton, D. Hollowood, K. Honscheid, B. Hoyle, E. Huff, D. Huterer, B. Jain, D. James, M. Jarvis, N. Jeffrey, T. Jeltema, A. Kovacs, E. Krause, R. Kron, K. Kuehn, N. Kuropatkin, O. Lahav, P.-F. Leget, P. Lemos, A. Liddle, C. Lidman, M. Lima, H. Lin, N. MacCrann, M. Maia, J. Marshall, P. Martini, J. McCullough, P. Melchior, J. Mena-Fernández, F. Menanteau, R. Miquel, J. Mohr, R. Morgan, J. Muir, J. Myles, S. Nadathur, A. Navarro-Alsina, R. Nichol, R. Ogando, Y. Omori, A. Palmese, S. Pandey, Y. Park, F. Paz-Chinchón, D. Petravick, A. Pieres, A. Plazas Malagón, A. Porredon, J. Prat, M. Raveri, M. Rodriguez-Monroy, R. Rollins, A. Romer, A. Roodman, R. Rosenfeld, A. Ross, E. Rykoff, S. Samuroff, C. Sánchez, E. Sanchez, J. Sanchez, D. Sanchez Cid, V. Scarpine, M. Schubnell, D. Scolnic, L. Secco, S. Serrano, I. Sevilla-Noarbe, E. Sheldon, T. Shin, M. Smith, M. Soares-Santos, E. Suchyta, M. Swanson, M. Tabbutt, G. Tarle, D. Thomas, C. To, A. Troja, M. Troxel, D. Tucker, I. Tutusaus, T. Varga, A. Walker, N. Weaverdyck, R. Wechsler, J. Weller, B. Yanny, B. Yin, Y. Zhang, and J. Zuntz, Physical Review D **105** (2022), 10.1103/physrevd.105.023520.

[11] 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, A. H. Wright, A. Amon, M. Bilicki, J. de Jong, M. Crocce, A. Dvornik, T. Erben, M. C. Fortuna, F. Getman, B. Giblin, K. Glazebrook, H. Hoekstra, S. Joudaki, A. Kannawadi, F. Köhlinger, C. Lidman, L. Miller, N. R. Napolitano, D. Parkinson, P. Schneider, H. Shan, E. A. Valentijn, G. Verdoes Kleijn, and C. Wolf, Astronomy & Astrophysics **646**, A140 (2021).

[12] S. More, S. Sugiyama, H. Miyatake, M. M. Rau, M. Shirasaki, X. Li, A. J. Nishizawa, K. Osato, T. Zhang, M. Takada, T. Hamana, R. Takahashi, R. Dalal, R. Mandelbaum, M. A. Strauss, Y. Kobayashi, T. Nishimichi, M. Oguri, W. Luo, A. Kannawadi, B.-C. Hsieh, R. Armstrong, J. Bosch, Y. Komiyama, R. H. Lupton, N. B. Lust, L. A. MacArthur, S. Miyazaki, H. Murayama, Y. Okura, P. A. Price, P. J. Tait, M. Tanaka, and S.-Y. Wang, "Hyper suprime-cam year 3 results: Measurements of clustering of sdss-boss galaxies, galaxy-galaxy lensing and cosmic shear," (2023), arXiv:2304.00703 [astro-ph.CO].

[13] H. Gil-Marín, J. Noreña, L. Verde, W. J. Percival, C. Wagner, M. Manera, and D. P. Schneider, Monthly Notices of the Royal Astronomical Society **451**, 539–580 (2015).

[14] H. Gil-Marín, W. J. Percival, L. Verde, J. R. Brownstein, C.-H. Chuang, F.-S. Kitaura, S. A. Rodríguez-Torres, and M. D. Olmstead, Monthly Notices of the Royal Astronomical Society **465**, 1757 (2016), https://academic.oup.com/mnras/article-pdf/465/2/1757/8364690/stw2679.pdf.

[15] Z. Slepian, D. J. Eisenstein, F. Beutler, C.-H. Chuang, A. J. Cuesta, J. Ge, H. Gil-Marín, S. Ho, F.-S. Kitaura, C. K. McBride, R. C. Nichol, W. J. Percival, S. Rodríguez-Torres, A. J. Ross, R. Scoccimarro, H.-J. Seo, J. Tinker, R. Tojeiro, and M. Vargas-Magaña, Monthly Notices of the Royal Astronomical Society **468**, 1070–1083 (2017).

[16] O. H. E. Philcox, J. Hou, and Z. Slepian, "A first detection of the connected 4-point correlation function of galaxies using the boss cmass sample," (2021), arXiv:2108.01670 [astro-ph.CO].

[17] M. M. Ivanov, O. H. Philcox, T. Nishimichi, M. Simonović, M. Takada, and M. Zaldarriaga, Physical Review D **105** (2022), 10.1103/physrevd.105.063512.

[18] M. S. Cagliari, M. Barberi-Squarotti, K. Pardede, E. Castorina, and G. D'Amico, "Bispectrum constraints on primordial non-gaussianities with the eboss dr16 quasars," (2025), arXiv:2502.14758 [astro-ph.CO].

[19] V. Yankelevich and C. Porciani, Mon. Not. R. Astron. Soc. **483**, 2078 (2019), arXiv:1807.07076 [astro-ph.CO].

[20] A. Oddo, E. Sefusatti, C. Porciani, P. Monaco, and A. G. Sánchez, Journal of Cosmology and Astroparticle Physics **2020**, 056–056 (2020).

[21] M. S. Wang, F. Beutler, J. Aguilar, S. Ahlen, D. Bianchi, D. Brooks, T. Claybaugh, A. de la Macorra, P. Doel, A. Font-Ribera, E. Gaztañaga, G. Gutierrez, K. Honscheid, C. Howlett, D. Kirkby, A. Lambert, M. Landriau, R. Miquel, G. Niz, F. Prada, I. Pérez-Ràfols, G. Rossi, E. Sanchez, D. Schlegel, M. Schubnell, D. Sprayberry, G. Tarlé, and B. A. Weaver, "Window convolution of the galaxy clustering bispectrum," (2024), arXiv:2411.14947 [astro-ph.CO].

[22] J. Fluri, T. Kacprzak, A. Lucchi, A. Refregier, A. Amara, T. Hofmann, and A. Schneider, Physical Review D **100** (2019), 10.1103/physrevd.100.063514.

[23] N. Jeffrey, L. Whiteway, M. Gatti, J. Williamson, J. Alsing, A. Porredon, J. Prat, C. Doux, B. Jain, C. Chang, T. Y. Cheng, T. Kacprzak, P. Lemos, A. Alarcon, A. Amon, K. Bechtol, M. R. Becker, G. M. Bernstein, A. Campos, A. C. Rosell, R. Chen, A. Choi, J. DeRose, A. Drlica-Wagner, K. Eckert, S. Everett, A. Ferté, D. Gruen, R. A. Gruendl, K. Herner, M. Jarvis, J. McCullough, J. Myles, A. Navarro-Alsina, S. Pandey, M. Raveri, R. P. Rollins, E. S. Rykoff, C. Sánchez, L. F. Secco, I. Sevilla-Noarbe, E. Sheldon, T. Shin, M. A. Troxel, I. Tutusaus, T. N. Varga, B. Yanny, B. Yin, J. Zuntz, M. Aguena, S. S. Allam, O. Alves, D. Bacon, S. Bocquet, D. Brooks, L. N. da Costa, T. M. Davis, J. D. Vicente, S. Desai, H. T. Diehl, I. Ferrero, J. Frieman, J. García-Bellido, E. Gaztanaga, G. Giannini, G. Gutierrez, S. R. Hinton, D. L. Hollowood, K. Honscheid, D. Huterer, D. J. James, O. Lahav, S. Lee, J. L. Marshall, J. Mena-Fernández, R. Miquel, A. Pieres, A. A. P. Malagón, A. Roodman, M. Sako, E. Sanchez, D. S. Cid, M. Smith, E. Suchyta, M. E. C. Swanson, G. Tarle, D. L. Tucker, N. Weaverdyck, J. Weller, P. Wiseman, and M. Yamamoto, "Dark energy survey year 3 results: likelihood-free, simulation-based $w$cdm inference with neural compression of weak-lensing map statistics," (2024), arXiv:2403.02314 [astro-ph.CO].

[24] P. Lemos, L. Parker, C. Hahn, S. Ho, M. Eickenberg, J. Hou, E. Massara, C. Modi, A. M. Dizgah, B. R.-S. Blancard, D. Spergel, and SimBIG Collaboration, Phys. Rev. D **109**, 083536 (2024).

[25] A. Roncoli, A. Ćiprijanović, M. Voetberg, F. Villaescusa-Navarro, and B. Nord, "Domain adaptive graph neural networks for constraining cosmological parameters across multiple data sets," (2024), arXiv:2311.01588 [astro-ph.CO].

[26] G. Valogiannis and C. Dvorkin, (2022), 10.48550/ARXIV.2204.13717.

[27] G. Valogiannis and C. Dvorkin, Physical Review D **105** (2022), 10.1103/physrevd.105.103534.

[28] G. Valogiannis, S. Yuan, and C. Dvorkin, "Precise cosmological constraints from boss galaxy clustering with a simulation-based emulator of the wavelet scattering transform," (2024), arXiv:2310.16116 [astro-ph.CO].

[29] S. Gauthier, B. Thérien, L. Alsène-Racicot, M. Chaudhary, I. Rish, E. Belilovsky, M. Eickenberg, and G. Wolf, "Parametric scattering networks," (2022), arXiv:2107.09539 [cs.LG].

[30] V. Khemani, M. H. Azarian, and M. G. Pecht, Electronics **11** (2022), 10.3390/electronics11030451.

[31] M. Craigie, P. L. Taylor, Y.-S. Ting, C. Cuesta-Lazaro, R. Ruggeri, and T. M. Davis, arXiv preprint arXiv:2405.13083 (2024).

[32] K. Kuijken, C. Heymans, A. Dvornik, H. Hildebrandt, J. T. A. de Jong, A. H. Wright, T. Erben, M. Bilicki, B. Giblin, H.-Y. Shan, F. Getman, A. Grado, H. Hoekstra, L. Miller, N. Napolitano, M. Paolilo, M. Radovich, P. Schneider, W. Sutherland, M. Tewes, C. Tortora, E. A. Valentijn, and G. A. Verdoes Kleijn, Astronomy & Astrophysics **625**, A2 (2019).

[33] Euclid Collaboration, Y. Mellier, Abdurro'uf, J. A. Acevedo Barroso, A. Achúcarro, J. Adamek, R. Adam, G. E. Addison, N. Aghanim, M. Aguena, V. Ajani, Y. Akrami, A. Al-Bahlawan, A. Alavi, I. S. Albuquerque, G. Alestas, G. Alguero, A. Allaoui, S. W.

Allen, V. Allevato, A. V. Alonso-Tetilla, B. Altieri, A. Alvarez-Candal, S. Alvi, A. Amara, L. Amendola, J. Amiaux, I. T. Andika, S. Andreon, A. Andrews, G. Angora, R. E. Angulo, F. Annibali, A. Anselmi, S. Anselmi, S. Arcari, M. Archidiacono, G. Aricò, M. Arnaud, S. Arnouts, M. Asgari, J. Asorey, L. Atayde, H. Atek, F. Atrio-Barandela, M. Aubert, E. Aubourg, T. Auphan, N. Auricchio, B. Aussel, H. Aussel, P. P. Avelino, A. Avgoustidis, S. Avila, S. Awan, R. Azzollini, C. Baccigalupi, E. Bachelet, D. Bacon, M. Baes, M. B. Bagley, B. Bahr-Kalus, A. Balaguera-Antolinez, E. Balbinot, M. Balcells, M. Baldi, I. Baldry, A. Balestra, M. Ballardini, O. Ballester, M. Balogh, E. Bañados, R. Barbier, S. Bardelli, M. Baron, T. Barreiro, R. Barrena, J. C. Barriere, B. J. Barros, A. Barthelemy, N. Bartolo, A. Basset, P. Battaglia, A. J. Battisti, C. M. Baugh, L. Baumont, L. Bazzanini, J. P. Beaulieu, V. Beckmann, A. N. Belikov, J. Bel, F. Bellagamba, M. Bella, E. Bellini, K. Benabed, R. Bender, G. Benevento, C. L. Bennett, K. Benson, P. Bergamini, J. R. Bermejo-Climent, F. Bernardeau, D. Bertacca, M. Berthe, J. Berthier, M. Bethermin, F. Beutler, C. Bevillon, S. Bhargava, R. Bhatawdekar, D. Bianchi, L. Bisigello, A. Biviano, R. P. Blake, A. Blanchard, J. Blazek, L. Blot, A. Bosco, C. Bodendorf, T. Boenke, H. Böhringer, P. Boldrini, M. Bolzonella, A. Bonchi, M. Bonici, D. Bonino, L. Bonino, C. Bonvin, W. Bon, J. T. Booth, S. Borgani, A. S. Borlaff, E. Borsato, A. Bosco, B. Bose, M. T. Botticella, A. Boucaud, F. Bouche, J. S. Boucher, D. Boutigny, T. Bouvard, R. Bouwens, H. Bouy, R. A. A. Bowler, V. Bozza, E. Bozzo, E. Branchini, G. Brando, S. Brau-Nogue, P. Brekke, M. N. Bremer, M. Brescia, M. A. Breton, J. Brinchmann, T. Brinckmann, C. Brockley-Blatt, M. Brodwin, L. Brouard, M. L. Brown, S. Bruton, J. Bucko, H. Buddelmeijer, G. Buenadicha, F. Buitrago, P. Burger, C. Burigana, V. Busillo, D. Busonero, R. Cabanac, L. Cabayol-Garcia, M. S. Cagliari, A. Caillat, L. Caillat, M. Calabrese, A. Calabro, G. Calderone, F. Calura, B. Camacho Quevedo, S. Camera, L. Campos, G. Canas-Herrera, G. P. Candini, M. Cantiello, V. Capobianco, E. Cappellaro, N. Cappelluti, A. Cappi, K. I. Caputi, C. Cara, C. Carbone, V. F. Cardone, E. Carella, R. G. Carlberg, M. Carle, L. Carminati, F. Caro, J. M. Carrasco, J. Carretero, P. Carrilho, and J. Carron Duque, arXiv e-prints , arXiv:2405.13491 (2024), arXiv:2405.13491 [astro-ph.CO].

[34] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew, J. R. P. Angel, G. Z. Angeli, R. Ansari, P. Antilogus, C. Araujo, R. Armstrong, K. T. Arndt, P. Astier, É. Aubourg, N. Auza, T. S. Axelrod, D. J. Bard, J. D. Barr, A. Barrau, J. G. Bartlett, A. E. Bauer, B. J. Bauman, S. Baumont, E. Bechtol, K. Bechtol, A. C. Becker, J. Becla, C. Beldica, S. Bellavia, F. B. Bianco, R. Biswas, G. Blanc, J. Blazek, R. D. Blandford, J. S. Bloom, J. Bogart, T. W. Bond, M. T. Booth, A. W. Borgland, K. Borne, J. F. Bosch, D. Boutigny, C. A. Brackett, A. Bradshaw, W. N. Brandt, M. E. Brown, J. S. Bullock, P. Burchat, D. L. Burke, G. Cagnoli, D. Calabrese, S. Callahan, A. L. Callen, J. L. Carlin, E. L. Carlson, S. Chandrasekharan, G. Charles-Emerson, S. Chesley, E. C. Cheu, H.-F. Chiang, J. Chiang, C. Chirino, D. Chow, D. R. Ciardi, C. F. Claver, J. Cohen-Tanugi, J. J. Cockrum, R. Coles, A. J. Connolly, K. H. Cook, A. Cooray, K. R. Covey, C. Cribbs, W. Cui, R. Cutri, P. N. Daly, S. F. Daniel, F. Daruich, G. Daubard, G. Daues, W. Dawson, F. Delgado, A. Dellapenna, R. de Peyster, M. de Val-Borro, S. W. Digel, P. Doherty, R. Dubois, G. P. Dubois-Felsmann, J. Durech, F. Economou, T. Eifler, M. Eracleous, B. L. Emmons, A. Fausti Neto, H. Ferguson, E. Figueroa, M. Fisher-Levine, W. Focke, M. D. Foss, J. Frank, M. D. Freemon, E. Gangler, E. Gawiser, J. C. Geary, P. Gee, M. Geha, C. J. B. Gessner, R. R. Gibson, D. K. Gilmore, T. Glanzman, W. Glick, T. Goldina, D. A. Goldstein, I. Goodenow, M. L. Graham, W. J. Gressler, P. Gris, L. P. Guy, A. Guyonnet, G. Haller, R. Harris, P. A. Hascall, J. Haupt, F. Hernandez, S. Herrmann, E. Hileman, J. Hoblitt, J. A. Hodgson, C. Hogan, J. D. Howard, D. Huang, M. E. Huffer, P. Ingraham, W. R. Innes, S. H. Jacoby, B. Jain, F. Jammes, M. J. Jee, T. Jenness, G. Jernigan, D. Jevremović, K. Johns, A. S. Johnson, M. W. G. Johnson, R. L. Jones, C. Juramy-Gilles, M. Jurić, J. S. Kalirai, N. J. Kallivayalil, B. Kalmbach, J. P. Kantor, P. Karst, M. M. Kasliwal, H. Kelly, R. Kessler, V. Kinnison, D. Kirkby, L. Knox, I. V. Kotov, V. L. Krabbendam, K. S. Krughoff, P. Kubánek, J. Kuczewski, S. Kulkarni, J. Ku, N. R. Kurita, C. S. Lage, R. Lambert, T. Lange, J. B. Langton, L. Le Guillou, D. Levine, M. Liang, K.-T. Lim, C. J. Lintott, K. E. Long, M. Lopez, P. J. Lotz, R. H. Lupton, N. B. Lust, L. A. MacArthur, A. Mahabal, R. Mandelbaum, T. W. Markiewicz, D. S. Marsh, P. J. Marshall, S. Marshall, M. May, R. McKercher, M. McQueen, J. Meyers, M. Migliore, M. Miller, and D. J. Mills, Astrophys. J. **873**, 111 (2019), arXiv:0805.2366 [astro-ph].

[35] S. Cheng and B. Ménard, Monthly Notices of the Royal Astronomical Society **507**, 1012 (2021).

[36] K. Cranmer, J. Brehmer, and G. Louppe, Proceedings of the National Academy of Sciences **117**, 30055 (2020), https://www.pnas.org/doi/pdf/10.1073/pnas.1912789117.

[37] A. Tejero-Cantero, J. Boelts, M. Deistler, J.-M. Lueckmann, C. Durkan, P. J. Gonçalves, D. S. Greenberg, and J. H. Macke, Journal of Open Source Software **5**, 2505 (2020).

[38] G. Papamakarios, T. Pavlakou, and I. Murray, "Masked autoregressive flow for density estimation," (2018), arXiv:1705.07057 [stat.ML].

[39] J.-M. Lueckmann, J. Boelts, D. S. Greenberg, P. J. Gonçalves, and J. H. Macke, "Benchmarking simulation-based inference," (2021), arXiv:2101.04653 [stat.ML].

[40] S. Cheng and B. Ménard, (2021), 10.48550/ARXIV.2112.01288.

[41] D. Erhan, Y. Bengio, A. Courville, and P. Vincent, *Visualizing higher-layer features of a deep network*, Tech. Rep. 4323 (Université de Montréal, Dept. IRO, 2009).

[42] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," (2013), arXiv:1311.2901 [cs.CV].

[43] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson, "Understanding neural networks through deep visualization," (2015), arXiv:1506.06579 [cs.CV].

[44] S. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," (2017), arXiv:1705.07874

[cs.AI].
[45] N. S. de Santi, F. Villaescusa-Navarro, L. Raul Abramo, H. Shao, L. A. Perez, T. Castro, Y. Ni, C. C. Lovell, E. Hernández-Martínez, F. Marinacci, D. N. Spergel, K. Dolag, L. Hernquist, and M. Vogelsberger, Journal of Cosmology and Astroparticle Physics **2025**, 082 (2025).