UnAPNet: UNCERTAINTY AWARE PHOTOMETRIC REDSHIFT ESTIMATION USING RESNET AND SPECTRAL-NORMALIZED NEURAL GAUSSIAN PROCESS

Swagata Biswas, Shubhrangshu Ghosh, Shailesh Deshpande, Arpan Pal

TCS Research, India
Email:{swagata.biswas|shubhrangshu.ghosh|shailesh.deshpande|arpan.pal}@tcs.com

ABSTRACT

Modern wide-field surveys such as SDSS, DES, HSC-SSP, Euclid and LSST are generating petabyte-scale datasets that require accurate and scalable photometric redshift (Pz) estimation. Although existing ML and DL methods have improved Pz prediction, most existing approaches provide limited uncertainty quantification and do not account for the varying scientific impact of prediction errors. Reliable modeling of both prediction accuracy and uncertainty is essential for cosmological applications including weak lensing, galaxy clustering, dark energy studies and galaxy evolution. We propose **UnAPNet** (**Un**certainty **A**ware **P**hotometric Redshift Estimation using Res**Net** and Spectral Normalized Neural Gaussian Process (SNGP)), a unified framework that jointly models *aleatoric uncertainty* from observational errors and *epistemic uncertainty* arising from model and data limitations. UnAPNet combines a ResNet-based *Regressor* with an SNGP-based *Uncertainty Estimator* and directly incorporates catalog-level photometric measurement errors as input features. In addition, we introduce a LSST-guided domain-aware loss function that assigns progressively larger penalties to outliers and catastrophic outliers, improving robustness against severe prediction failures. We evaluate UnAPNet on the HSC-SSP dataset, a representative precursor to LSST. Using metrics including RMS, bias, catastrophic outlier fraction, and normalized epistemic variance, we demonstrate that UnAPNet delivers competitive Pz accuracy together with well-calibrated uncertainty estimates even in higher redshift ranges $0 < Pz < 3.99$. The proposed approach substantially improves prediction reliability and uncertainty calibration, achieving performance close to LSST Science Requirements and highlighting the promise of uncertainty-aware deep learning for next-generation cosmological surveys.

***Index Terms*—** Astronomy, Photometric Redshift, Uncertainty, Large Scale Surveys

1. INTRODUCTION

Modern wide-field astronomical surveys such as SDSS, DES, HSC-SSP, Euclid and LSST are generating petabyte-scale datasets containing billions of celestial objects [1]. These surveys provide unprecedented opportunities to study galaxy evolution, dark energy, and the large-scale structure of the universe. A fundamental requirement for such studies is accurate redshift estimation [2]. Redshift characterizes the shift of spectral lines due to the Doppler effect and the expansion of the universe. Receding sources exhibit a displacement toward longer wavelengths, whereas approaching sources show a shift toward shorter wavelengths. As these spectral shifts serve as key indicators of an object's velocity and distance, redshift estimation plays a pivotal role in studying the large-scale structure and evolution of the universe. Although spectroscopic redshifts are highly accurate, obtaining them for billions of objects is time consuming and challenging [3]. Consequently, photometric redshift (Pz) estimation, derived from multi-band imaging data, has become an essential and scalable alternative.

Existing Pz estimation methods can be broadly classified into template-fitting and empirical approaches. Template-fitting techniques [4] estimate redshifts by matching the observed photometry band fluxes with spectral energy distribution (SED) templates, offering physical interpretability but suffering from template incompleteness and systematic biases. Empirical methods, including machine learning (ML) and deep learning (DL) models [5], learn the mapping between photometric features and redshift from spectroscopic training samples and often achieve superior predictive performance. However, their effectiveness depends strongly on the quality, completeness, and representativeness of the training data, particularly at high redshifts where labeled samples are scarce. Although hybrid methods [6] attempt to combine the strengths of both approaches, challenges related to biased training samples, photometric uncertainties, and distribution shifts remain. Beyond accurate redshift prediction, reliable uncertainty quantification [7] has become increasingly important for downstream applications such as weak gravitational lensing, galaxy clustering, dark energy studies, and galaxy evolution analyses. Most conventional ML-based Pz estimators [8] [9] provide only point estimates and do not adequately characterize their associated uncertainties, potentially leading to biased scientific inferences.

Uncertainty in Pz estimation arises from two primary sources [10]: aleatoric uncertainty, which originates from

observational noise, calibration errors, and missing band information, and epistemic uncertainty, which results from model limitations, incomplete training data coverage, and distribution mismatch. While aleatoric uncertainty is inherent to the data acquisition process, epistemic uncertainty can be reduced through refining model architecture and augmenting more representative training datasets. Since modern astronomical surveys generally record observational uncertainties associated with both photometric and spectroscopic measurements, these uncertainties can be explicitly exploited to develop uncertainty-aware Pz estimation models.

Motivated by the significance of reliable uncertainty quantification, several probabilistic approaches for Pz estimation have been proposed in recent years including Bayesian Neural Networks (BNNs) [7], GPz [11], ANNz2 [12], and more recently Starkindler [13]. These approaches provide uncertainty-aware Pz estimates, with BNNs offering predictive uncertainties, GPz modeling heteroscedastic noise, ANNz2 generating redshift probability density functions, and Starkindler incorporating spectroscopic redshift uncertainties directly into training. Together, they highlight the growing role of uncertainty-aware learning in Pz estimation.

Despite these advances, key limitations remain. Most existing methods focus primarily on either aleatoric or epistemic uncertainty rather than modeling both jointly. BNNs and Gaussian Process-based approaches provide principled uncertainty estimates but often face computational and scalability challenges for large surveys. In addition, many methods do not fully utilize observational uncertainties available in survey catalogs. Furthermore, the state-of-the-art approaches typically do not account for the severity of prediction errors during training, treating all errors uniformly regardless of their scientific impact.

To address these challenges, we propose **Un**certainty **A**ware **P**hotometric Redshift estimation using Res**Net** and Spectral Normalized Neural Gaussian Process (**UnAPNet**) framework, that explicitly accounts for both aleatoric and epistemic uncertainties. **UnAPNet** combines a Residual Neural Network (ResNet) model [14] and a Spectral Normalized Neural Gaussian Process (SNGP) model [15]. The ResNet model estimates Pz, while the SNGP model provides uncertainty estimates by measuring the confidence of predictions with respect to the training data distribution. The ResNet architecture has demonstrated remarkable success in extracting hierarchical representations from complex data and mitigating optimization difficulties in deep networks. The SNGP framework augments deep neural networks with Gaussian process-inspired uncertainty estimation capabilities while maintaining computational efficiency. Also, we introduce a domain-aware loss function in the **UnAPNet** that penalizes prediction errors according to their scientific impact. Based on the LSST Science Requirements [16], prediction errors are categorized as inliers, outliers, and catastrophic outliers. The loss increases with error severity, generating stronger optimization gradients for highly erroneous predictions, thus improving the model robustness.

Compared with existing methods, UnAPNet directly exploits catalog-level photometric uncertainties, provides scalable epistemic uncertainty estimation without costly Bayesian inference, and generates reliable confidence measures alongside accurate redshift predictions. These capabilities make it particularly well suited for next-generation cosmological surveys, where uncertainty quantification is as critical as predictive accuracy. We demonstrate that this unified ResNet-SNGP framework improves both Pz estimation and uncertainty calibration on large-scale astronomical datasets. We evaluate **UnAPNet** using the Hyper Suprime-Cam Subaru Strategic Program (HSC-SSP)[1] dataset, which is widely regarded as the representative of the forthcoming LSST survey. We report the catastrophic outlier, RMS, bias and normalized epistemic variance, and benchmark these metrics against the LSST Science Requirements Document (SRD) [16]. The proposed framework achieves performance approaching the LSST SRD thresholds.

2. DATA AND METHODS

2.1. Dataset used

Our objective is to estimate the Pz and the predictive uncertainty associated with it for large scale surveys. Here, we have used the Hyper Suprime-Cam Strategic Survey Program (HSC-SSP) catalog. This survey is considered to be a representative of the Legacy Survey of Space and Time (LSST)[2] whose first light was received on June 2026. The dataset comprises of 286401 samples including galaxies and quasars.

2.2. Network Architecture

Our proposed network architecture, **UnAPNet**, consists of two modules - *Regressor* and *Uncertainty Estimator* as shown in Fig 1. The Regressor module is responsible for Pz estimation (z_{phot}), while the Uncertainty Estimator module quantifies the uncertainty (z_{phot_un}) associated with the redshift predictions generated by the Regressor module. In this work, for Pz estimation, the Regressor module employs a ResNet architecture [14] without a classifier head. To assess the prediction uncertainty, the Uncertainty Estimator module utilizes the Spectral-normalized Neural Gaussian Process (SNGP) framework [15]. The primary objective of our approach is to account for both aleatoric and epistemic uncertainties in Pz estimation. Aleatoric uncertainty arises from inherent noise and measurement errors in the observational data. To incorporate this source of uncertainty, we provide the observational errors associated with the *grizy* bands as additional inputs to the ResNet model. Epistemic uncertainty, on the

[1] https: //zenodo.org/records/5528827

[2] https://rubinobservatory.org/explore/how-rubin-works/lsst

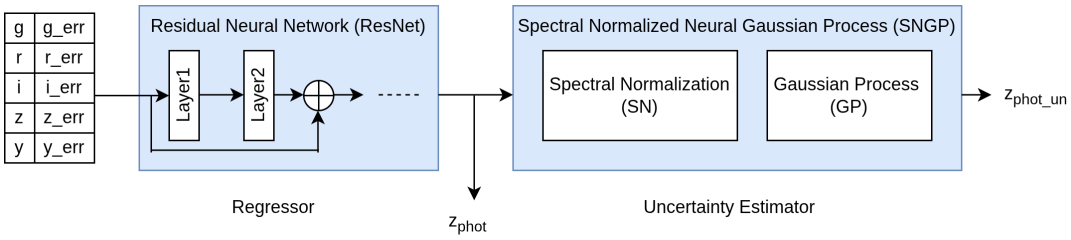


Fig. 1. **UnAPNet** framework: Consists of a *Regressor* module and an *Uncertainty Estimator* module. *grizy* band photometry data alongwith its associated observational error (*g_err,r_err,i_err,z_err,y_err*) are fed as input to the UnAPNet. The *Regressor* predicts the photometric redshift (z_{phot}) and the *Uncertainty Estimator* provides the normalized variance uncertainty (z_{phot_un}) in estimating z_{phot}. The Regressor can be any ML/ DL model. Here, we have used ResNet. The Uncertainty Estimator used is Spectral Normalized Neural Gaussian Process.

other hand, stems from the limited knowledge of the underlying data distribution and potential model misspecification. To capture this uncertainty, we employ the SNGP model, which estimates the confidence of the ResNet predictions and provides an uncertainty measure corresponding to the Pz estimates.

Regressor Module: In this module, we have employed Residual Neural Network (ResNet), a deep neural network architecture designed to facilitate the training of very deep models through the use of residual learning. Each residual block learns a residual function $G(x)$, and the output is computed as $z_{phot} = x + G(x)$, where x denotes the input feature map including $grizy$ band photometry data and the error in observation associated with each band (*g_err,r_err,i_err,z_err,y_err*), and $G(x)$ comprises a sequence of convolutional, normalization, and activation layers. The identity skip connections improve gradient flow during backpropagation, mitigating vanishing-gradient issues and enabling the learning of rich hierarchical feature representations. For regression task, the final layer is replaced with a fully connected output layer that predicts continuous-valued targets. The network is trained by minimizing the loss function, allowing it to learn the mapping between input features and continuous output variables. In this study, ResNet without the classifier head is employed as a feature extraction backbone to learn robust representations for accurate regression prediction.

Uncertainty Estimator: Conventional deep neural networks can assign high-confidence predictions to out-of-distribution samples. Spectral-normalized Neural Gaussian Process (SNGP) addresses this limitation by making the neural networks more distance-aware, increasing prediction uncertainty as samples deviate from the training data distribution. SNGP enhances a deterministic neural network with reliable uncertainty estimation by combining spectral normalization and a Gaussian Process (GP) output layer. First, spectral normalization constrains the spectral norm (largest singular value) of each layer's weight matrix, encouraging an approximately bi-Lipschitz, distance-preserving representation where distances between inputs are reasonably maintained in the latent space. This preserves the geometric structure of the training data, ensuring that latent-space distances remain meaningful indicators of similarity and enabling more reliable identification of in-distribution versus out-of-distribution inputs through improved uncertainty estimation. Second, SNGP replaces the conventional dense output layer with a Gaussian Process (GP) layer, which provides predictions alongwith well-calibrated uncertainty estimates. Since Gaussian Processes are inherently distance-aware, their predictive uncertainty naturally increases as a test sample moves farther from the training data. To make this approach scalable for deep learning applications, SNGP employs an efficient GP approximation based on random Fourier features and a Laplace-approximated posterior, enabling computationally efficient training and inference while preserving the GP's ability to quantify uncertainty for out-of-distribution and sparsely observed regions of the input space. By combining distance-preserving latent representations with a distance-aware Gaussian Process output layer, SNGP produces more reliable and well-calibrated uncertainty estimates (z_{phot_un}) that reflect a model's familiarity with a given input. This improves robustness, calibration, and out-of-distribution awareness, achieving uncertainty estimation performance comparable to ensemble methods while requiring only a single model and a single forward pass, making it efficient for large-scale real-time applications.

2.3. Experimental Setup

Dataset Preprocessing: We used the HSC-SSP catalog. For each sample, we considered the $grizy$ band photometry flux data and the associated error in observation with each photometry band. The selected dataset comprises of samples spanning within the redshift (z) range of $0 < z < 3.99$. We randomly split the dataset in $80:20$ for training and testing.

Performance metrics: We have evaluated our proposed model using four performance metrics as defined in the LSST SRD - *outlier* (O), *catastrophic outlier* (O_c), *bias* (b), and *RMS* (rms). These are computed using Eqs. 1–4, where z_{spec}

and z_{phot} denote the ground truth spectroscopic and predicted photometric redshifts, respectively, and n is the sample size:

$$O = \frac{|z_{phot} - z_{spec}|}{1 + z_{spec}} \tag{1}$$

$$O_c = |z_{phot} - z_{spec}| \tag{2}$$

$$b = \frac{z_{phot} - z_{spec}}{1 + z_{spec}} \tag{3}$$

$$rms = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(\frac{z^i_{phot} - z^i_{spec}}{1 + z^i_{spec}}\right)^2} \tag{4}$$

A sample is identified as an outlier if $O > 0.15$ and as a catastrophic outlier if $O_c > 1.0$. As per the LSST standards, the thresholds should be: $rms < 0.2$, $b < 0.003$, and a catastrophic outlier fraction $O_c < 10\%$ of the total sample.

Loss Function: We classify the prediction error, E, into three categories: *inlier*(E_{in}), *outlier*(E_{out}) and *catastrophic outlier*(E_{cout}). An error is considered an inlier when $E \leq 0.15$, an outlier when $0.15 < E \leq 1$, and a catastrophic outlier when $E > 1$. These thresholds follow the LSST SRD guidelines, with severity increasing as $E_{in} < E_{out} < E_{cout}$.

To reflect this severity hierarchy, the proposed loss (L) function (Eq 5) combines three components: L_{in}, L_{out} and L_{cout} corresponding to loss caused by inliers, outliers, and catastrophic outliers, respectively. Higher-error categories are assigned larger weights ($w_1 \leq w_2 \ll w_3$), with L_{cout} receiving the highest weight(w_3) and L_{in} the lowest(w_1). The outlier and catastrophic outlier terms are conditionally activated, using indicators I_1 and I_2, only when the prediction error falls within their respective ranges, enabling stronger penalization of severe errors. I_1 is activated when $O \geq 0.15$ and I_2 is activated when $O_c \geq 1$. As the error increases, the loss grows proportionally, generating larger gradients and driving more substantial parameter updates during optimization.

$$\begin{aligned} L &= L_{in} + L_{out} + L_{cout} \\ &= \sum_{i=1}^{n}\{w_1 log_2(\frac{exp(z^i_{phot} - z^i_{spec})}{\sum_{j=1}^{n} exp(1 + z^j_{phot} - z^j_{spec})}) \\ &+ w_2 I_1 log_2(\frac{exp(z^i_{phot} - z^i_{spec})}{\sum_{j=1}^{n} exp(1 + z^j_{spec})}) \\ &+ w_3 I_2 log_2(exp(z^i_{phot} - z^i_{spec}))\} \end{aligned} \tag{5}$$

$$\begin{aligned} I_1 &= 1 \quad if \quad exp(\frac{z^i_{phot} - z^i_{spec}}{1 + z^i_{spec}}) \geq exp(0.15) \\ &= 0 \quad otherwise \end{aligned}$$

$$\begin{aligned} I_2 &= 1 \quad if \quad exp(z^i_{phot} - z^i_{spec}) \geq exp(1) \\ &= 0 \quad otherwise \end{aligned}$$

$$w_1 + w_2 + w_3 = 1$$

$$w_1 \leq w_2 \ll w_3$$

The loss formulation uses an exponential function to make it insensitive to the sign of the error, while a logarithmic term prevents excessively small loss values. Each component expression is derived intuitively by keeping the actual definition of inliers, outlier (Eq 1) and catastrophic outlier (Eq 2) in mind, ensuring that the network focuses on reducing high-severity prediction errors during training.

Uncertainty estimation: Aleatoric uncertainty is incorporated as an input to the model, while predictive uncertainty is estimated using SNGP, which combines a deep neural network with a Gaussian Process output layer. For a test input x, a random Fourier feature representation $F(x)$ is computed (Eq. 6), from which the predictive mean $\mu(x)$ (Eq. 7) and variance $var(x)$ (Eq. 8) are derived. Here, W_L, $h(x)$, b_L, D, L, and β denote the random Fourier projection matrix, learned feature representation, bias vector, Fourier feature dimension, number of random Fourier features, and posterior mean weight vector, respectively. The variance increases as inputs move farther from the training data in the learned feature space. The predictive probability $p(y|x)$ (Eq. 9) obtained by integrating the softmax output over a Gaussian distribution parameterized by the mean and variance, resulting in confident predictions for in-distribution inputs and higher uncertainty for out-of-distribution samples.

$$F(x) = \sqrt{\frac{D}{L}}\cos W_L h(x) + b_L \tag{6}$$

$$\mu(x) = F(x)^T \beta \tag{7}$$

$$var(x) = F(x)^T \sum F(x) \tag{8}$$

$$p(y|x) = N(y; \mu(x); var(x)) \tag{9}$$

Runtime environment: We used an NVIDIA A5500 GPU for model training and evaluation. The training time required was about 6 hours.

3. RESULTS AND DISCUSSION

We evaluated our proposed model using the HSC-SSP dataset. The redshift ranges between $0 < z < 3.99$. The ground truth spectroscopic redshift vs the predicted photometric redshift is visualized using a hexbinplot in Fig 2 where the minimum bin density count is 25. The samples contained within the green tolerance bounds are inliers with prediction error, E, lying within ± 0.15. The samples lying outside the green tolerance bounds but within the red tolerance bounds are outliers, E ranging between $\pm(0.15–1.0)$. The samples beyond the red tolerance bounds, i.e. with $E > \pm 1.0$, are catastrophic outliers. The embedded scatter plot shows the data points associated with bins containing fewer than 25 samples.

The samples are predominantly clustered around the perfect match line. However, each prediction is associated with an inherent uncertainty. While most deviations are minor and fall within the inlier region, a subset of predictions exhibit significant errors resulting into outliers and catastrophic outliers. Identifying these and quantifying predictive uncertainty are essential to prevent error propagation and accumulation, which can adversely impact downstream analyses and decision-making processes.

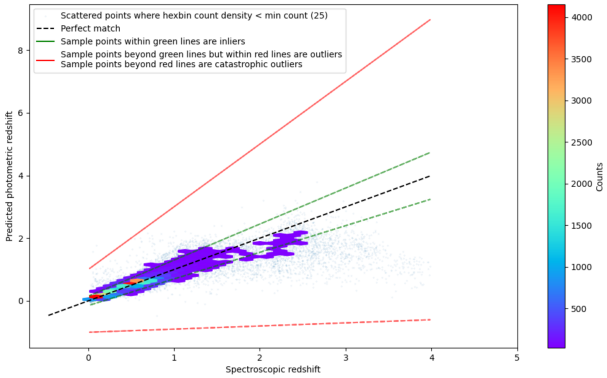


Fig. 2. Hexbin plot for the interval $0 < z < 3.99$, with a min bin density of 25 samples. Embedded scatter plot shows data points associated with bins containing fewer than 25 samples. Prediction error, E, are set at ± 0.15 for inliers, $\pm(0.15\text{–}1.0)$ for outliers, and $> \pm 1.0$ for catastrophic outliers.

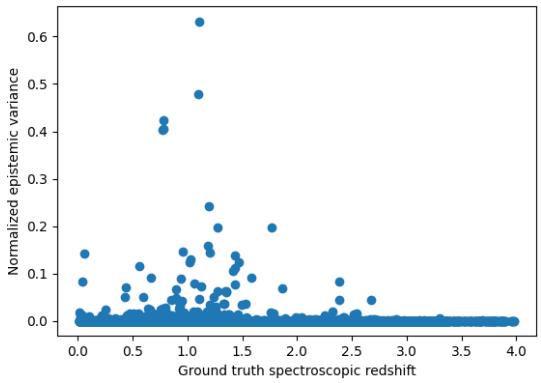


Fig. 3. Uncertainty estimation. Ground truth spectroscopic redshift vs normalized epistemic variance for samples within the redshift range $0 < z < 3.99$.

In Fig. 3, we show the predictive uncertainty of the Pz estimates using normalized epistemic variance. Most test samples exhibit near-zero variance, indicating high confidence and a well-learned feature-redshift mapping. A small fraction of samples, mainly at higher redshifts, show significantly high uncertainty due to their lower signal-to-noise ratio (SNR) and sparse representation in the training set, which limits the model's ability to learn reliable feature-redshift relationships in these regions.

In Table 1, we report the performance metrics defined in Sec. 2.3, including *rms*, *bias*, and the *catastrophic outlier rate* (O_c). Unlike most state-of-the-art studies, which are typically limited to $z \leq 1.5$ [7], our evaluation extends to $z \leq 3.99$. This broader range is more challenging due to lower SNR and fewer training samples at high redshifts. Despite these challenges, our method achieves performance comparable to state-of-the-art approaches while satisfying two LSST Science Requirements.

Table 1. Comparison with SOTA methods.

Architecture(z)	**RMS**	**Bias**	**O_c**
Mizuki (≤ 1.5)	0.102	0.307	0.011
DEmP (≤ 1.5)	0.092	0.277	0.003
XGBoost (≤ 1.5)	0.022	0.149	0.002
SPIDERz (≤ 1.5)	0.051	0.199	0.002
BNN (≤ 1.5)	0.174	0.013	0.023
RF (≤ 1.5)	0.088	0.001	0.006
PhoRQNet (≤ 1.5)	0.137	0.0007	0
Scaled Ensemble (≤ 3.99)	0.13	0.0127	0.022
UnAPNet (≤ 3.99)	0.108	0.002	0.028
LSST Req	< 0.02	0.003	< 0.1

4. CONCLUSION

We presented **UnAPNet**, an uncertainty-aware Pz estimation framework that jointly models aleatoric and epistemic uncertainties by combining a ResNet-based regressor with an SNGP-based uncertainty estimator. By incorporating catalog-level photometric uncertainties and introducing a LSST-guided domain-aware loss function, the proposed framework improves robustness to outliers and catastrophic outliers while providing reliable confidence estimates alongside redshift predictions. Experiments on the HSC-SSP dataset demonstrate competitive photometric redshift accuracy, well-calibrated uncertainty estimates across the redshift range ($0 < Pz < 3.99$), and performance approaching LSST Science Requirements. These results highlight the potential of UnAPNet as a scalable and trustworthy solution for Pz estimation along with uncertainty quantification in next-generation cosmological surveys.

5. REFERENCES

[1] R Kanthavel, R Adline Freeda, and R Dhaya, "Astroinformatics and big data challenges," *The Intelligent Universe: AI's Role in Astronomy*, pp. 279–312, 2025.

[2] Jeffrey A Newman and Daniel Gruen, "Photometric redshifts for next-generation surveys," *Annual Review of Astronomy and Astrophysics*, vol. 60, no. 1, pp. 363–414, 2022.

[3] Alberto Fernandez-Soto, Kenneth M Lanzetta, Hsiao-Wen Chen, Sebastian M Pascarelle, and Noriaki Yahata, "On the compared accuracy and reliability of spectroscopic andphotometric redshift measurements," *The Astrophysical Journal Supplement Series*, vol. 135, no. 1, pp. 41, 2001.

[4] Gabriel B Brammer, Pieter G Van Dokkum, and Paolo Coppi, "Eazy: a fast, public photometric redshift code," *The Astrophysical Journal*, vol. 686, no. 2, pp. 1503, 2008.

[5] Ben Henghes, Jeyan Thiyagalingam, Connor Pettitt, Tony Hey, and Ofer Lahav, "Deep learning methods for obtaining photometric redshift estimations from images," *Monthly Notices of the Royal Astronomical Society*, vol. 512, no. 2, pp. 1696–1709, 2022.

[6] Jonas Chris Ferrao, Dickson Dias, Pranav Naik, Glory D'cruz, Anish Naik, Siya Khandeparkar, and Manisha Gokuldas Fal Dessai, "Template-fitting meets deep learning: redshift estimation using physics-guided neural networks," *Journal of Astrophysics and Astronomy*, vol. 47, no. 1, pp. 28, 2026.

[7] Evan Jones, Tuan Do, Bernie Boscoe, Jack Singal, Yujie Wan, and Zooey Nguyen, "Improving photometric redshift estimation for cosmology with lsst using bayesian neural networks," *The Astrophysical Journal*, vol. 964, no. 2, pp. 130, 2024.

[8] Swagata Biswas, Shubhrangshu Ghosh, Avyarthana Ghosh, Yogesh Wadadekar, Abhishek Roy Choudhury, Arijit Mukherjee, Shailesh Deshpande, and Arpan Pal, "Photometric redshift estimation using scaled ensemble learning," *The Astrophysical Journal*, vol. 996, no. 1, pp. 107, 2026.

[9] Swagata Biswas, Shubhrangshu Ghosh, Avyarthana Ghosh, Shailesh Deshpande, and Arpan Pal, "Phorqnet: Photometric redshift estimator using quantum neural network," in *2025 15th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*. IEEE, 2025, pp. 1–6.

[10] Matthew A Chan, Maria J Molina, and Christopher A Metzler, "Estimating epistemic and aleatoric uncertainty with a single model," *Advances in neural information processing systems*, vol. 37, pp. 109845–109870, 2024.

[11] Ibrahim A Almosallam, Matt J Jarvis, and Stephen J Roberts, "Gpz: non-stationary sparse gaussian processes for heteroscedastic uncertainty estimation in photometric redshifts," *Monthly Notices of the Royal Astronomical Society*, vol. 462, no. 1, pp. 726–739, 2016.

[12] Iftach Sadeh, Filipe B Abdalla, and Ofer Lahav, "Annz2: photometric redshift and probability distribution function estimation using machine learning," *Publications of the Astronomical Society of the Pacific*, vol. 128, no. 968, pp. 104502, 2016.

[13] Raahul Singh and Ashutosh Pandey, "Starkindler: An uncertainty aware objective for photometric redshift estimation," *arXiv preprint arXiv:2512.22566*, 2025.

[14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[15] Jeremiah Liu, Zi Lin, Shreyas Padhy, Dustin Tran, Tania Bedrax Weiss, and Balaji Lakshminarayanan, "Simple and principled uncertainty estimation with deterministic deep learning via distance awareness," *Advances in neural information processing systems*, vol. 33, pp. 7498–7512, 2020.

[16] Rachel Mandelbaum, Tim Eifler, Renée Hložek, Thomas Collett, Eric Gawiser, Daniel Scolnic, David Alonso, Humna Awan, Rahul Biswas, Jonathan Blazek, et al., "The lsst dark energy science collaboration (desc) science requirements document," *arXiv preprint arXiv:1809.01669*, 2018.