

Scalable Statistical Inference of Photometric Redshift via Data Subsampling

Arindam Fadikar¹[0000–0001–7396–0350], Stefan M. Wild¹[0000–0002–6099–2772],
and Jonas Chaves-Montero²[0000–0002–9553–4261]

¹Mathematics and Computer Science Division, Argonne National Laboratory, Lemont
IL 60439, USA

{afadikar,wild}@anl.gov

²Donostia International Physics Centre, Paseo Manuel de Lardizabal 4, 20018
Donostia-San Sebastian, Spain.

jonas.chaves@dipc.org

Abstract. Handling big data has largely been a major bottleneck in traditional statistical models. Consequently, when accurate point prediction is the primary target, machine learning models are often preferred over their statistical counterparts for bigger problems. But full probabilistic statistical models often outperform other models in quantifying uncertainties associated with model predictions. We develop a data-driven statistical modeling framework that combines the uncertainties from an ensemble of statistical models learned on smaller subsets of data carefully chosen to account for imbalances in the input space. We demonstrate this method on a photometric redshift estimation problem in cosmology, which seeks to infer a distribution of the redshift—the stretching effect in observing the light of far-away galaxies—given multivariate color information observed for an object in the sky. Our proposed method performs balanced partitioning, graph-based data subsampling across the partitions, and training of an ensemble of Gaussian process models.

Keywords: Gaussian process · data subsampling · photometric redshift.

1 Introduction

Data analysis techniques have become an essential part of a scientist’s toolbox for making inferences about an underlying system or phenomenon. With the advancement of modern computing, processing power has gone up manyfold, and data volumes have grown at a similar pace. Thus, there is a growing demand for scalable machine learning (ML) and statistical techniques that can handle large data in an efficient and reasonable way. Recent trends include the use of ML and statistical regression techniques such as deep neural networks [21, 18], tree-based models [16], and Gaussian process (GP) models [28, 14]. Each technique has advantages and disadvantages that promote or limit its use for a specific application. For example, deep neural networks are known for achieving superior prediction accuracy but at the cost of significant data and compute

cost for training. Statistical models such as GPs approximate the global input-output relationship and provide full uncertainty estimates at a fraction of the data required by deep neural networks. However, statistical models do not generally scale well with data size. Some GP models have been proposed to find workarounds such as introducing sparse approximation of large correlation matrices [20], using locally learned models [14] or considering only a subset of the data [1].

In this paper we propose a statistical modeling framework that can handle large training data by leveraging data subsampling combined with advanced statistical regression techniques to provide a full density estimate. The basic idea entails using smaller subsets of data to train regression models and building an ensemble of these models. Our modeling approach differs from other approximations in that we attempt to learn the global input-output relationship in each individual subsample and model. This is counterintuitive and opposite approaches that seek to build accurate local response surfaces [14, 20]. However, uncertainty estimates can be undesirably constricted in locally learned models. Furthermore, data partitioning and subsampling in our proposed approach are driven entirely by data distribution and are free from model influences present in other locally learned models. We focus on the analysis paradigm in cosmology that deals with estimation of the redshifts of objects (e.g., galaxies) as a motivating application, which is discussed in the following section.

The rest of the paper is structured as follows. In §2 we describe the estimation problem and data used. We present our proposed methodology in §3 and numerical results in §4. In §5 we summarize our approach and its benefits.

2 Photometric redshift estimation

The cosmological analysis of galaxy surveys, from gathering information on dark energy to unveiling the nature of dark matter, relies on the precise projection of galaxies from two-dimensional sky maps into the three-dimensional space [32, 20]. However, measuring the distance of distant objects in the Universe from the Earth is challenging. Furthermore, the accelerated expansion of the Universe causes the wavelength of light from a distant object to be stretched or *redshifted*. Interestingly, the redshift of an object is proportional to its receding speed and can be used to estimate the radial distance to this source. For an accurate redshift estimation, one would have to obtain the full spectrum of each galaxy at a very high resolution, which is a demanding and time-consuming task [19, 3]. An alternative to this approach is to infer the redshift based on the intensity of galaxy light observed through a reduced number of wavebands [2]. Such an approximation is known as photometric redshift, whereas the redshift obtained from the full spectrum is called spectroscopic redshift.

Photometric redshift estimation methods can be divided into two categories: spectral energy distribution template fitting [25, 10] and statistical regression or ML methods [11, 7, 20]. Supervised ML methods such as deep neural networks have recently seen success in approximating the mapping between broadband

fluxes and redshift [3]. Many of the successes of ML methods, however, rely on the availability of large training datasets coupled with large computing power. Because of their increased computational complexity, statistical regression methods have not been a preferred choice, even when the data size is moderately large.

The case study we consider is the publicly available data release 7 of the Sloan Digital Sky Survey (SDSS) [33], which contains approximately 1 million galaxies with spectroscopic redshifts estimates. From this initial sample, we select $\simeq 100\,000$ galaxies with signal-to-noise ratio larger than 10 in all photometric bands, clean photometry, and high-confident spectroscopic redshift below $z = 0.3$. The photometry of each SDSS galaxy consists of flux measurement in five broadband filters u, g, r, i , and z , which serve as input to the predictive model for redshift.

In general we consider a dataset of $N (= 10^5)$ scattered data points $\mathbf{x}_1, \dots, \mathbf{x}_N$ contained in a compact input space $\mathcal{X} \subset \mathbb{R}^d$, with corresponding logged redshifts y_1, \dots, y_N . Without loss of generality, we assume that the input space has been normalized so that $\mathcal{X} = [0, 1]^d$ is the unit cube and the logged redshift values are transformed to mean zero and a standard deviation of 1. In our case study, d is 5, with the dimensions corresponding to four colors computed by taking the ratio of the flux consecutive filters and the magnitude in the i -band. In what follows, we refer to these variables as color space.

Approximating the relationship between broadband filter values and the photometric redshift is challenging because of several factors. For example, lack of coverage of the input space in the training dataset results in poor predictions at the extremes, as well as degeneracies arising in the prediction of the redshift in some cases. Hence, simple Gaussian predictive uncertainty may not accurately represent the distribution of the redshift given a set of colors. For example, a galaxy at a high redshift with high luminosity and a galaxy at a low redshift with low luminosity may register the same flux values to a telescope. Considering such uncertainty in the predictive distribution would require discovering and modeling these latent processes. Mixed-density networks [9] and full Bayesian photometric redshift estimation [4] are examples that model the predictive distribution by a finite number of Gaussians. While our objective is similar to [9], we give greater importance to obtaining a full predictive distribution that can accurately reflect multiple predictive modes.

3 Statistical methodology

Our proposed approach consists of three intermediate stages: (1) partitioning the input color space \mathcal{X} into a user-defined number of partitions in order to ensure balance among the input data, (2) subsampling datapoints from these partitions, and (3) training a regression model on the sampled data. The final predictive model is then based on an ensemble of models obtained by repeated execution of the latter two stages. Although we present the basic implementations here for simplicity, scalability is emphasized in each stage: the partitioning can exploit domain decomposition parallelism, the subsampling allows for model construc-

tion using datasets of user-specified greatly reduced size, and the ensemble is naturally parallelizable.

3.1 Partitioning the input space

We begin by partitioning the color space \mathcal{X} into a set of mutually exclusive subsets $\mathcal{C} = \{\mathcal{C}_k : k \in \mathcal{I}\}$ such that $\cup_{k \in \mathcal{I}} \mathcal{C}_k = \mathcal{X}$ and $\mathcal{C}_{k_i} \cap \mathcal{C}_{k_j} = \emptyset$ for $k_i \neq k_j$.

We especially target datasets that are highly nonuniform, including those arising in redshift estimation, and thus we seek balanced partitions. By balanced, here we intend for the number of training points in all partitions to satisfy $|\mathcal{C}_k| \approx \frac{N}{|\mathcal{I}|}$, where $|\mathcal{C}_k|$ denotes the cardinality of \mathcal{C}_k .

The number of partitions (m) is not predetermined for our balanced hyperrectangle partitioning algorithm. However, the values of the inputs to the algorithm, N_{\min} and N_{\max} —minimum and maximum number of datapoints in each partition, respectively—are influenced by a desired m . To achieve a (n approximately) balanced hyperrectangle partition, we propose a three-step procedure that consists of an initialization of coarse partitions, pruning to satisfy the minimality condition $\min_{k \in \mathcal{I}} |\mathcal{C}_k| \geq N_{\min}$, and then splitting to satisfy the maximality condition $\max_{k \in \mathcal{I}} |\mathcal{C}_k| \leq N_{\max}$. We also require $N_{\max} \geq 2N_{\min}$ to guarantee termination of our procedure.

Initialize Interval boundaries $0 = x_0^p < x_1^p < \dots < x_{m_p}^p = 1$ of size m_p are defined along each input dimension $p = 1, \dots, d$. Then an initial set of hyperrectangle partitions \mathcal{C} in the d -dimensional input space is constructed as the Cartesian product of all 1D intervals:

$$\mathcal{C} = \{A_{j_1}^1 \times \dots \times A_{j_d}^d : j_p = 1, \dots, m_p, p = 1, \dots, d\}. \quad (1)$$

This leaves open a choice of interval boundary values along each dimension. In this study, we have opted for quantile-based splits. The result of the initialization step is $\prod_{p=1}^d m_p$ hyperrectangles, some of which may be empty.

Merge In the next step, partitions with cardinality less than N_{\min} are merged successively with their neighbors until the minimality condition is satisfied. Note that at the end of the merge step, some partitions may have cardinality greater than N_{\max} . The merging algorithm is briefly described below.

We define $\mathcal{S}_{\min} = \{\mathcal{C}_k \in \mathcal{C} : |\mathcal{C}_k| < N_{\min}\}$ to be the set of partitions with cardinality less than N_{\min} . We begin by identifying a target partition $\mathcal{C}_{(0)} \in \arg \min_{\mathcal{C}_k \in \mathcal{S}_{\min}} |\mathcal{C}_k|$, namely, a partition with the smallest cardinality. We also define the directional neighborhood function $\mathcal{N}_{\omega}^p(\cdot)$, which represents neighbors of \cdot along dimension $p \in \{1, \dots, d\}$ and where $\omega \in \{\text{lower} \equiv l, \text{upper} \equiv u\}$ denotes the relative position of the neighbors with respect to dimension p .

Given $\mathcal{C}_{(0)}$, the selection of partition(s) to merge with is equivalent to finding an appropriate dimension and direction to merge along. At most $2d$ such merging choices exist. We require that the newly formed partition be a hyperrectangle.

For a given partition $\mathcal{C}_{(0)}$, the merging dimension and direction are selected according to following condition:

$$(p^*, \omega^*) \in \arg \min_{\substack{p=1, \dots, d \\ \omega \in \{l, u\}}} \sum_{\mathcal{C}_k \in \mathcal{N}_\omega^p(\mathcal{C}_{(0)})} |\mathcal{C}_k|, \quad (2)$$

in other words, the dimension and direction for which the updated partition contains the least number of datapoints among all possible combinations. In the case of multiple (p^*, ω^*) possibilities in (2), the one that results in the most uniform sides after merging is selected. In particular, we choose the merger that results in the smallest ratio between the longest and shortest sides. Once the merged hyperrectangle achieves at least cardinality N_{\min} , the set \mathcal{S}_{\min} is updated. The merging step continues until \mathcal{S}_{\min} is empty.

Split In the split step, partitions with cardinality greater than N_{\max} are successively broken into smaller partitions until the maximality condition is satisfied. We define $\mathcal{S}_{\max} = \{\mathcal{C}_k : |\mathcal{C}_k| > N_{\max}\}$ to be the set of hyperrectangles with cardinality greater than N_{\max} . A partition with highest cardinality in \mathcal{S}_{\max} is selected for a split, $\mathcal{C}_{(m)} \in \arg \max_{\mathcal{C}_k \in \mathcal{S}_{\max}} |\mathcal{C}_k|$. A split of a hyperrectangle is defined as breaking it into two hyperrectangles along one dimension. Hence, for any split operation the appropriate dimension needs to be identified. The break point in that dimension can be any location for which the two new partitions satisfy the minimality condition (such a location exists because we require $N_{\max} \geq 2N_{\min}$). Our case study uses the median as the break point. To promote uniformity in the shape of hyperrectangle partitions, we always select the dimension corresponding to the longest side to perform the split, unless any of the resulting hyperrectangles breaks the minimality conditions, in which case we move to the next best dimension. Successive splits are carried out until \mathcal{S}_{\max} is empty.

3.2 Conditional sampling from partitions

After assigning N datapoints to m partitions $\mathcal{C}_{k_1}, \dots, \mathcal{C}_{k_m}$, we move to the next stage where one or more samples are drawn from each partition according to our proposed sampling rule. The primary goal of such a sampling scheme is to explore and discover latent processes by sequentially sampling from the partitions obtained in the previous step.

Induced graph on the partitions An essential ingredient of our proposed sampling technique is a graph structure based on the partitions on \mathcal{C} . We define an undirected graph $\mathcal{G} = (V, E)$ induced by the partitions, where nodes are defined to be m partitions $V = \{\mathcal{C}_{k_1}, \dots, \mathcal{C}_{k_m}\}$ and there is an undirected edge between \mathcal{C}_{k_i} and \mathcal{C}_{k_j} , $k_i \neq k_j$ if their closures share more than one point. In other words $|cl(\mathcal{C}_{k_i}) \cap cl(\mathcal{C}_{k_j})| > 1$. This criterion means that partitions that share only a corner are not considered neighbors. Note that the edges can entirely be determined during the partition phase for hyperrectangle partitions.

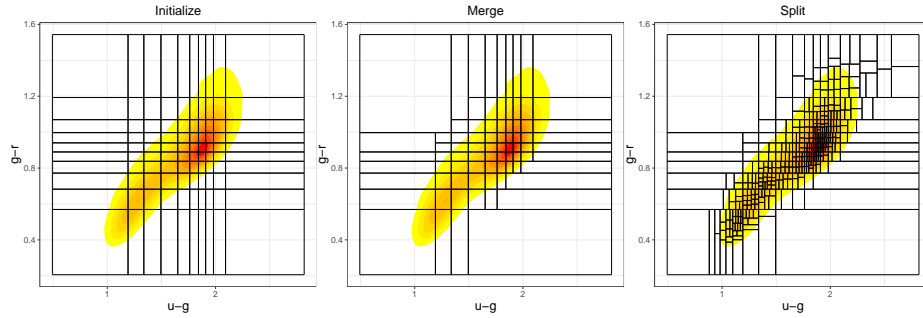


Fig. 1: Illustration of hyperrectangle partitioning of a 2D space containing approximately 80,000 datapoints, with input $(N_{\min}, N_{\max}) = (100, 300)$. At the beginning, 100 partitions are initialized based on marginal percentiles 10%, \dots , 90% (left); merging then occurs for partitions with cardinality less than N_{\min} (middle); and lastly partitions with cardinality more than N_{\max} are successively split into smaller partitions, resulting in 393 partitions (right).

Other forms of graphs are also possible and sometimes necessary to reflect the known underlying manifold structure of the input space. Our proposed sampling procedure is independent of the graph-generating process and hence allows for more flexibility and adaptability to diverse arrays of applications.

The algorithm We define importance sampling [24] alike strategy that leverages the spatial dependence among the input values via the partition-induced graph \mathcal{G} . The basic idea is a sampler that traverses through the partitions via edges E in \mathcal{G} and successively draws a datapoint from a partition by conditioning on the previously sampled datapoints from neighbor partitions. Such a strategy is motivated primarily by the intent of untangling convoluted latent processes that the data is arising from, without adding an extra computational burden. We note that since a single datapoint is drawn from each partition, the overall complexity of the sampling depends on N/m and m rather than N .

Without loss of generality, we assume that \mathcal{G} is connected. If it is not, then the sampling stage can be carried out independently for each connected component. The sampling stage is an iterative process. We denote the datapoints in partition \mathcal{C}_{k_i} as $\{(\mathbf{x}_1^{(i)}, y_1^{(i)}), \dots, (\mathbf{x}_{n_i}^{(i)}, y_{n_i}^{(i)})\}$. The sampler is initialized by randomly sampling a datapoint $(\mathbf{x}_*^{(1)}, y_*^{(1)})$ from a partition \mathcal{C}_{k_1} . (The choice of k_1 is discussed later.) Next, the sampler moves to a partition that is an unsampled neighbor of \mathcal{C}_{k_1} . If there is more than one neighbor to choose from, the sampler moves to an available neighbor according to a criterion similar to the initialization step. The chosen partition is denoted by \mathcal{C}_{k_2} , and the datapoints $\{(\mathbf{x}_1^{(2)}, y_1^{(2)}), \dots, (\mathbf{x}_{n_2}^{(2)}, y_{n_2}^{(2)})\}$ in \mathcal{C}_{k_2} are weighted by a symmetric Gaussian kernel denoted by $K(\cdot|y_*^{(1)})$ centered at $y_*^{(1)}$. Let $f_y^{(i)}$ denote the distribution of y in

\mathcal{C}_{k_i} . Then the weighted sampling distribution $g_y^{(i)}$ is defined as

$$g_y^{(i)}(y|\mathcal{C}_{k_i}) \propto f_y^{(i)}(y|\mathcal{C}_{k_i}) \times \prod_{j \in \mathcal{A}_i} K(y|y_*^{(j)}), \quad i = 2, \dots, m, \quad (3)$$

where $\mathcal{A}_i = I(\mathcal{N}(\mathcal{C}_{k_i}) \cap \mathcal{V}_{(i-1)})$, $\mathcal{N}(\cdot)$ are the neighbors of \cdot in \mathcal{G} , $\mathcal{V}_{(i)}$ is the set of the first i partitions visited by the sampler, $I(\cdot)$ is the index set of \cdot , and $K(y|y_*^{(j)}) \propto \exp\{-\frac{1}{2\eta^2}(y - y_*^{(j)})^2\}$. η controls the width of the kernel K . Then a sample is drawn from \mathcal{C}_{k_2} according to $g_y^{(2)}$, and the sampler walks through the partitions until all partitions are visited. We note that in practice we set η to a value less than the standard deviation of $y_*^{(j)}$ in (3).

As previously noted, one of the objectives of such a sampling scheme is to discover and model latent processes that the data might be arising from. Our initialization criterion is geared toward facilitating this goal. The initial partition is chosen to be one for which the variance of y within the partition is maximal. In other words,

$$\mathcal{C}_{k_1} \in \arg \max_{\mathcal{C}_k \in \mathcal{C}_{k_1}, \dots, \mathcal{C}_{k_m}} \text{Var}(y|\mathcal{C}_k).$$

Subsequent selection of partitions is carried out in a similar way. At any iteration, the next sampling partition is chosen from the unvisited neighbors of sampled partitions that have maximum variance,

$$\mathcal{C}_{k_i} \in \arg \max_{\mathcal{C}_k \in \mathcal{N}(\mathcal{V}_{(i-1)}) \cap \mathcal{V}_{(i-1)}^c} \text{Var}(y|\mathcal{C}_k), \quad i = 2, \dots, m.$$

This process continues until samples are drawn from all partitions. The weighting strategy of datapoints based on its neighbors encourages the sampling scheme to discover latent global structures in the input-output relationship.

In a complete setup, this sampling is performed multiple times (independently and in parallel), generating multiple datasets to train our statistical model.

3.3 Modeling via ensembles

After a set of training data is generated, the last stage is to train a regression model. Our model of choice is Gaussian process [28], which is a semi-parametric regression model, fully characterized by a mean function and a covariance function. Historically GP models have been popular in both ML and statistics because of their ability to fit a large class of response surfaces [31, 23, 6, 29, 13]. The covariance function in a GP model often does the heavy lifting of describing the response variability by means of distance-based correlations among datapoints. Certain classes of GP covariance structures ensure smoothness and continuity in the response surface.

We define y to be a noisy realization from GP z . Then the data model can be written as

$$y(\mathbf{x}) = z(\mathbf{x}) + \epsilon, \quad \epsilon \stackrel{iid}{\sim} N(0, \sigma^2) \quad z(\mathbf{x}) \sim GP(0, \mathcal{C}_\Phi(\mathbf{x}, \mathbf{x}')),$$

where $\mathcal{C}_\Phi(\cdot, \cdot)$ is a covariance function with length-scale parameters Φ . The likelihood of the data is then given by the probability density function of a multivariate normal distribution:

$$\mathbf{y}^T = (y_1, \dots, y_n)^T \sim \text{MVN}(\mathbf{0}, \mathcal{C}_n), \quad (4)$$

where \mathcal{C}_n is an $n \times n$ matrix, obtained by $\mathcal{C}_n = [\mathcal{C}_\Phi(\mathbf{x}_i, \mathbf{x}_j) + \sigma^2 \delta_{i=j}]_{1 \leq i, j \leq n}$. Distribution of z at an untried input setting \mathbf{x}^* conditioned on n observations is also Gaussian, with mean and covariance given by

$$\begin{aligned} \mathbb{E}(y(\mathbf{x}^*)|\mathbf{x}, \mathbf{z}, \cdot) &= c_n(\mathbf{x}^*)^T \mathcal{C}_n^{-1} \mathbf{z}, \\ \text{Var}(y(\mathbf{x}^*)|\mathbf{x}, \mathbf{z}, \cdot) &= \mathcal{C}_n - c_n(\mathbf{x}^*)^T \mathcal{C}_n^{-1} c_n(\mathbf{x}^*), \end{aligned} \quad (5)$$

where $c_n(\mathbf{x}^*) = (\mathcal{C}_\Phi(\mathbf{x}^*, \mathbf{x}_1), \dots, \mathcal{C}_\Phi(\mathbf{x}^*, \mathbf{x}_n))$. Our implementation uses a scaled separable Gaussian covariance kernel $\mathcal{C}_\Phi(\mathbf{x}, \mathbf{x}') = \exp\left(-\sum_{p=1}^d \frac{(x_p - x'_p)^2}{\phi_p}\right)$. The length-scale parameter $\Phi = (\phi_1, \dots, \phi_d)$ controls the correlation strength along each dimension. Despite a GP's attractive properties, a major drawback of the standard GP model is the associated computational cost for estimating its parameters (i.e., Φ and σ). Each evaluation of the likelihood (4) involves inverting \mathcal{C}_n —an operation of $O(n^3)$ complexity. Hence, model training (i.e., estimation of Φ) becomes computationally infeasible as n grows. Alternatives have been proposed to deal with large n , including local GP approximations [14], knot-based low-rank representation [1, 30], process convolution [17], and compactly supported sparse covariance matrices [20]. While all these methods achieve a certain computational efficiency at large n , none is intended to discover and model all the latent stochastic processes that the data is possibly arising from. Dirichlet process-based mixture models attempt to solve this problem [27], but scalability remains a challenge. In contrast, our proposed method uses $n = m \ll N$ training samples for each GP model, and the sampling scheme actively searches for all latent smooth processes that can be inferred from the data.

Full predictive model Our final predictive model is constructed by taking an ensemble of N_m trained GPs models (each of these models trained with m samples as above). Denoting the predictive distribution at a new input \mathbf{x}^* as $f_i(y^*|\mathbf{x}^*)$, we define the ensemble predictor as

$$\hat{f}(y^*|\mathbf{x}^*) = \frac{1}{N_m} \sum_{i=1}^{N_m} f_i(y^*|\mathbf{x}^*). \quad (6)$$

Each f_i has a Gaussian distribution with mean and variance given by (5). Prediction intervals are computed based on this mixed Gaussian distribution. The individual GPs can be trained independently and hence can be done in parallel. The choice of N_m can be guided by the computing budget and the choices of N_{\min}, N_{\max} (which induce a value m). Figure 2 shows an illustration of 10 GP surfaces on a 2D color space, each trained on 50 datapoints sampled from

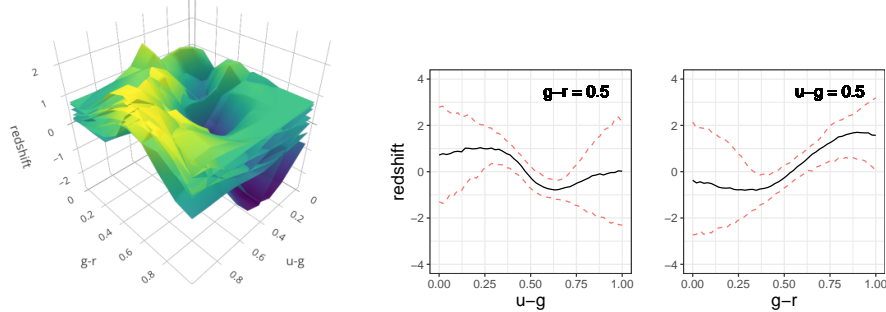


Fig. 2: Example ensemble of size 10 of mean GP surfaces; inputs are scaled to $[0,1]^2$, and the response is transformed to mean 0 and standard deviation 1. A 1D view of ensemble predictions of redshift at different values of $u-g$ and $g-r$ while keeping the other fixed at 0.5 is shown at the bottom. Black solid lines and red dashed lines represent median prediction and 90% confidence interval, respectively.

data of size 2,000. Predicted distribution of redshift at any tuple $(u-g, g-r)$ is then a combination of Gaussian distributions from the 10 GP models. We note that the estimated prediction uncertainty from ensemble models would be much larger compared with that of a single model on the full data. Moreover, in a single GP case, the family of predictive distributions is restricted to the family of unimodal symmetric Gaussian distributions, which is often inappropriate for noisy data. In contrast, the ensemble estimate can be interpreted as an approximation of the unknown distribution of the redshift conditional on the colors in a semi-parametric way.

4 Results

In this section we discuss numerical results from our proposed approach to model the redshift as a function of colors. After removing outliers, the size of the final dataset was reduced to 99,826, of which 20% were held out for out-of-sample prediction. Predictive distributions for the redshift conditional on colors were constructed by using our implementation of balanced hyperrectangle partitioning and sampling in R [26]; GP models were fitted by using the *mleHomGP* function in the *hetGP* package [5].

Partition and sample Inputs to the hyperrectangle partitioning scheme N_{\min} and N_{\max} were set at 50 and 150, respectively. These yielded a total of 749 nonempty rectangular partitions with an average cardinality of 107 in the 5D color space. Considering the tradeoff of execution time in the partition step and accuracy of the final prediction, we arrived at this particular choice of N_{\min}

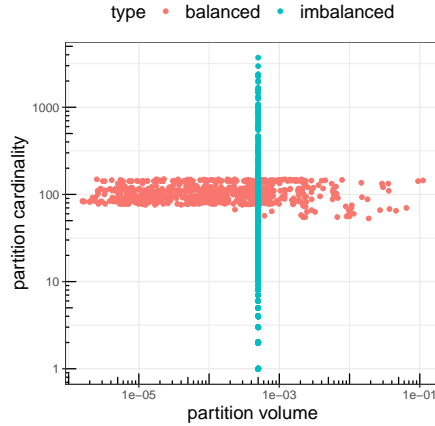


Fig. 3: Properties obtained from balanced hyperrectangle (in red) and equal-volume partitioning (in cyan), with 749 and 711 nonempty partitions, respectively.

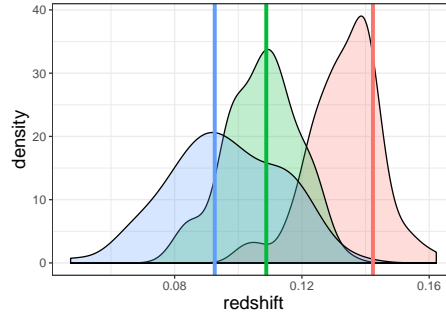


Fig. 4: Predicted density of redshift at three different color inputs (each indicated by a different shading in the figure). The vertical lines represent the truth for each color input.

and N_{\max} after a few trials. At each iteration of sampling and modeling, we drew one sample per partition. The samples were used to train a zero-mean GP model with a nugget term [5]. Effectively, each individual GP model was trained on 749 examples. In contrast to partitions being balanced in terms of cardinality, partitions can be constructed to be balanced with respect to volumes, which for hyperrectangles equates to dividing the input space into equal-volume hyperrectangles. The scatter plot in Fig. 3 shows the cardinality and volume of partitions as points on \mathbb{R}^2 from balanced hyperrectangle and equal-volume partitions. By fixing the number of partitions and their volume, equal-volume partitioning returned 5,769 empty and 711 nonempty partitions with cardinality ranging from 1 to 10,000. Datapoints were not uniformly distributed along any dimension, thus making the equal-volume partitions have highly imbalanced data compositions. A sample from equal-volume partitions would always result in a biased training set for a model. In contrast, balanced hyperrectangle partitioning optimizes over cardinality by adaptively forming the hyperrectangles that enforce roughly homogeneous data size across partitions.

Model and prediction As described in § 3.3, zero-mean GP models were fitted to each data subsample for $N_m = 50$ times. Maximum likelihood estimates of length-scale parameters Φ and the noise variance σ^2 were obtained and used to construct the ensemble estimate in (6). Figure 4 shows examples of predicted distributions at three different 5D color inputs and the true redshifts. To obtain smooth empirical density estimates, we sampled from the distribution given by (6); the samples were then used in a kernel density estimator to construct

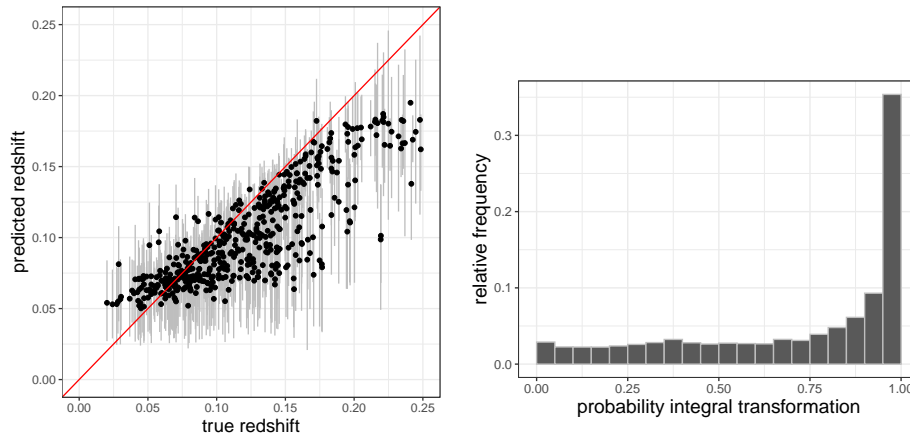


Fig. 5: True versus predicted redshifts for 500 random test points. Median predicted redshifts are denoted by black dots, and vertical grey bars show 90% prediction intervals on the left. A probability integral transformation (PIT) histogram plot is on the right.

the predicted density function. As expected, the distributions show asymmetric, multimodal behaviors that would be impossible to capture by a single GP model. On the other hand, unlike in a traditional setting, comparing mean prediction with the true response would not be appropriate here and would result in low prediction accuracy. In other cases, degeneracies in the prediction are resolved by considering the highest mode among all mixture components or the mean from the mixture component with the highest probability [9]. Figure 5 shows the prediction performance of 500 random test (out-of-sample) examples. For each example, we obtained the median prediction (denoted by black dots) and 90% high-density probability region (denoted by vertical grey bars). Because of the selection effect of observing cosmological objects in a restricted region, imbalance arises with respect to cosmological scale [9]. Objects with higher redshifts are usually less represented in the full dataset in consideration. This is also evident from the probability integral transformation (PIT) plot in Fig. 5. The PIT plot is used to assess the quality of the probabilistic predictions—a perfect uniform distribution suggests perfect accuracy. In our case it is far from uniform, because of the shift in prediction at higher redshifts.

Scalability Our framework allows for fairly straightforward parallel computations; each stage is amenable to parallelism. Once the data is partitioned according to any partitioning scheme, the ensemble members can be obtained in an embarrassingly parallel way. Each ensemble member consists of sampling data from the partitions and training the GP model, independently from other ensemble members. Almost linear speedup can be achieved when each ensemble member is run in parallel by using multiple cores in a single computing unit.

5 Discussion and conclusion

In the SDSS dataset considered here, redshift ranges from 0 to 0.3, with a high concentration around 0.1. To make sense of the Gaussianity assumption for each GP model, we modeled logged redshifts instead of raw ones. Input values were transformed to $[0, 1]$ to have consistent length-scale parameter estimates in the covariance. These transformations were handled at the beginning of the data-preprocessing step, before splitting the data into training and testing sets. Such an approach ensured that no extrapolation was being done in predicting the redshifts for the testing examples. Selection of N_{\min} and N_{\max} for balanced hyperrectangle partitioning is conceived to be a sequential updating process. In most applications, the choices are highly influenced by the intended size of the training set for each individual model in the ensemble. Prior knowledge of the expected smoothness of the global response surface can assist in choosing the size of the data subsample. For example, estimating a smooth surface using a GP with Gaussian covariance would require a substantially small number of training examples. We used $N_{\min} = 50$ and $N_{\max} = 150$ based on prior experience. Construction of the graph induced by partitions is simple and intuitive. One could argue for considering two hyperrectangles \mathcal{C}_{k_i} and \mathcal{C}_{k_j} to be neighbors when their boundaries intersect only at a corner, namely, $|cl(\mathcal{C}_{k_i}) \cap cl(\mathcal{C}_{k_j})| = 1$. In fact, neighbor relationships on the partitions can be as arbitrary as an application desires. Considering corner-sharing partitions as neighbors will produce a graph with many more edges, which in turn requires more compute time in the subsequent sampling and modeling stages.

Here we have presented a novel approach to regression modeling for large data that leverages efficient data sampling and an advanced statistical model such as GP modeling. The problem of estimating the full predictive distribution of the photometric redshift provides a base case for demonstrating novel aspects of our proposed methodology. Redshift estimation based on photometric surveys is an important problem. Almost all recent estimation protocols emphasize being able to use large photometry data to train their models [9, 3, 20]: since the new generation of telescopes promises a much larger survey area of the sky, models need to ingest this huge amount of data and produce forecasts for new observations in a reasonable amount of time. One way to meet this need is by using a combination of new algorithms and more powerful supercomputers. New modeling paradigms also need to be developed that find a sweet spot between careful choices of efficient algorithms, good models, and opportunity to scale. Indeed, the present work is motivated by this pursuit.

Another aspect of our work involves prediction targets, which are different from simply estimating the mean or median under parametric uncertainty assumptions. Best-case scenarios would involve nonparametric approximation of the unknown predictive distribution, an example of which can be found in Dirichlet process-based models [12]. Full Bayesian estimation here is rarely feasible with large data. Combining several simple models is a step toward the same goal but at a fraction of the cost. Successful ensemble techniques [15, 22, 14] and ML models [21, 8] achieve superior accuracy by combining locally learned struc-

tures in the response surface. While we draw motivation from such endeavors, our approach differs from them in a significant way. Each individual model in our ensemble does not target local structures but, rather, learns a global response surface every time a new model is trained on a sample of the data, allowing for a broad range of distribution to be covered. To the best of our knowledge, this work is the first of its kind that draws local inference by combining global information. Evaluating the accuracy of different combinations of partitioning schemes, statistical models, and ensemble sizes is a promising direction for future work.

Acknowledgments This material was based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics and SciDAC programs under Contract No. DE-AC02-06CH11357.

References

1. Banerjee, S., Fuentes, M.: Bayesian modeling for large spatial datasets. *Wiley Interdisciplinary Reviews: Computational Statistics* **4**(1), 59–66 (2012)
2. Baum, W.A.: Photoelectric determinations of redshifts beyond 0.2 c. *The Astrophysical Journal* **62**, 6–7 (1957)
3. Beck, R., Dobos, L., Budavári, T., Szalay, A.S., Csabai, I.: Photometric redshifts for the SDSS Data Release 12. *Monthly Notices of the Royal Astronomical Society* **460**(2), 1371–1381 (2016)
4. Benitez, N.: Bayesian photometric redshift estimation. *The Astrophysical Journal* **536**(2), 571 (2000)
5. Binois, M., Gramacy, R.B.: *hetGP: Heteroskedastic Gaussian Process Modeling and Design under Replication* (2019), <https://CRAN.R-project.org/package=hetGP>, r package version 1.1.2
6. Brahim-Belhouari, S., Bermak, A.: Gaussian process for nonstationary time series prediction. *Computational Statistics & Data Analysis* **47**(4), 705–712 (2004)
7. Cavuoti, S., Brescia, M., Longo, G., Mercurio, A.: Photometric redshifts with the quasi Newton algorithm (MLPQNA) Results in the PHAT1 contest. *Astronomy and Astrophysics* **546**, A13 (2012)
8. Dietterich, T.G., et al.: Ensemble learning. *The Handbook of Brain Theory and Neural Networks* **2**, 110–125 (2002)
9. D’Isanto, A., Polsterer, K.L.: Photometric redshift estimation via deep learning-generalized and pre-classification-less, image based, fully probabilistic redshifts. *Astronomy & Astrophysics* **609**, A111 (2018)
10. Fernández-Soto, A., Lanzetta, K.M., Yahil, A.: A new catalog of photometric redshifts in the Hubble deep field. *The Astrophysical Journal* **513**, 34–50 (1999)
11. Firth, A.E., Lahav, O., Somerville, R.S.: Estimating photometric redshifts with artificial neural networks. *Monthly Notices of the Royal Astronomical Society* **339**(4), 1195–1202 (2003)
12. Gelfand, A.E., Kottas, A., MacEachern, S.N.: Bayesian nonparametric spatial modeling with Dirichlet process mixing. *Journal of the American Statistical Association* **100**(471), 1021–1035 (2005)

13. Gramacy, R.B.: Surrogates: Gaussian Process Modeling, Design and Optimization for the Applied Sciences. Chapman Hall/CRC, Boca Raton, Florida (2020)
14. Gramacy, R.B., Apley, D.W.: Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics* **24**(2), 561–578 (2015)
15. Hastie, T., Rosset, S., Zhu, J., Zou, H.: Multi-class AdaBoost. *Statistics and its Interface* **2**(3), 349–360 (2009)
16. Hastie, T., Tibshirani, R., Friedman, J.: Random forests. In: *The Elements of Statistical Learning*, pp. 587–604. Springer, New York (2009)
17. Higdon, D.: Space and space-time modeling using process convolutions. In: *Quantitative Methods for Current Environmental Issues*, pp. 37–56. Springer, London (2002)
18. Hu, Y.H., Hwang, J.N.: *Handbook of Neural Network Signal Processing*. Acoustical Society of America (2002)
19. Ilbert, O., et al.: COSMOS Photometric Redshifts with 30-bands for 2-deg2. *The Astrophysical Journal* **690**(2), 1236–1249 (Jan 2009)
20. Kaufman, C.G., Bingham, D., Habib, S., Heitmann, K., Frieman, J.A., et al.: Efficient emulators of computer experiments using compactly supported correlation functions, with an application to cosmology. *The Annals of Applied Statistics* **5**(4), 2470–2492 (2011)
21. Lawrence, S., Giles, C.L., Tsoi, A.C., Back, A.D.: Face recognition: A convolutional neural-network approach. *IEEE Transactions on Neural Networks* **8**(1), 98–113 (1997)
22. Liaw, A., Wiener, M., et al.: Classification and regression by randomForest. *R News* **2**(3), 18–22 (2002)
23. Neal, R.M.: Regression and classification using Gaussian process priors. In: Bernardo, J.M., Berger, J.O., Dawid, A., Smith, A.F.M., et al. (eds.) *Bayesian Statistics*. vol. 6, pp. 476–501. Oxford University Press, Oxford (1998)
24. Neal, R.M.: Annealed importance sampling. *Statistics and Computing* **11**(2), 125–139 (2001)
25. Puschell, J.J., Owen, F.N., Laing, R.A.: Near-infrared photometry of distant radio galaxies - Spectral flux distributions and redshift estimates. *The Astrophysical Journal Letters* **257**, L57–L61 (1982)
26. R Core Team: *R: A Language and Environment for Statistical Computing* (2020), <https://www.R-project.org>
27. Rasmussen, C.E., Ghahramani, Z.: Infinite mixtures of Gaussian process experts. *Advances in Neural Information Processing Systems* **14**, 881–888 (2001)
28. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA (2005)
29. Sacks, J., Welch, W.J., Mitchell, T.J., Wynn, H.P.: Design and analysis of computer experiments. *Statistical Science* **4**, 409–423 (1989)
30. Snelson, E., Ghahramani, Z.: Sparse Gaussian processes using pseudo-inputs. *Advances in Neural Information Processing Systems* **18**, 1257–1264 (2005)
31. Wang, J., Hertzmann, A., Fleet, D.J.: Gaussian process dynamical models. *Advances in Neural Information Processing Systems* **18**, 1441–1448 (2005)
32. Weinberg, D.H., Mortonson, M.J., Eisenstein, D.J., Hirata, C., Riess, A.G., Rozo, E.: Observational probes of cosmic acceleration. *Physics Reports* **530**, 87–255 (2013)
33. York, D.G., et al.: The Sloan Digital Sky Survey: Technical summary. *The Astronomical Journal* **120**(3), 1579 (2000)