
Bayesian Approaches to Distribution Regression

Ho Chung Leon Law*
University of Oxford
ho.law@spc.ox.ac.uk

Dougal J. Sutherland*
University College London
dougal@gmail.com

Dino Sejdinovic
University of Oxford
dino.sejdinovic@stats.ox.ac.uk

Seth Flaxman
Imperial College London
s.flaxman@imperial.ac.uk

Abstract

Distribution regression has recently attracted much interest as a generic solution to the problem of supervised learning where labels are available at the group level, rather than at the individual level. Current approaches, however, do not propagate the uncertainty in observations due to sampling variability in the groups. This effectively assumes that small and large groups are estimated equally well, and should have equal weight in the final regression. We account for this uncertainty with a Bayesian distribution regression formalism, improving the robustness and performance of the model when group sizes vary. We frame our models in a neural network style, allowing for simple MAP inference using backpropagation to learn the parameters, as well as MCMC-based inference which can fully propagate uncertainty. We demonstrate our approach on illustrative toy datasets, as well as on a challenging problem of predicting age from images.

1 INTRODUCTION

Distribution regression is the problem of learning a regression function from samples of a distribution to a single set-level label. For example, we might attempt to infer the sentiment of texts based on word-level features, to predict the label of an image based on small patches, or even perform traditional parametric statistical inference by learning a function from sets of samples to the parameter values.

Recent years have seen wide-ranging applications of this framework, including inferring summary statistics in Approximate Bayesian Computation (Mitrovic et al., 2016), estimating Expectation Propagation messages (Jitkrittum et al., 2015), predicting the voting behaviour of demo-

graphic groups (Flaxman et al., 2015, 2016), and learning the total mass of dark matter halos from observable galaxy velocities (Ntampaka et al., 2015, 2016). Closely related distribution classification problems also include identifying the direction of causal relationships from data (Lopez-Paz et al., 2015) and classifying text based on bags of word vectors (Yoshikawa et al., 2014; Kusner et al., 2015).

One particularly appealing approach to the distribution regression problem is to represent the input set of samples by their kernel mean embedding (described in Section 2.1), where distributions are represented as single points in a reproducing kernel Hilbert space. Standard kernel methods can then be applied for distribution regression, classification, anomaly detection, and so on. This approach was perhaps first popularized by Muandet et al. (2012); Szábo et al. (2016) provided a recent learning-theoretic analysis.

In this framework, however, each distribution is simply represented by the empirical mean embedding, ignoring the fact that large sample sets are much more precisely understood than small ones. Most studies also use point estimates for their regressions, such as kernel ridge regression or support vector machines, thus ignoring uncertainty both in the distribution embeddings and in the regression model.

Our Contributions We propose a set of Bayesian approaches to distribution regression. The simplest method, similar to that of Flaxman et al. (2015), is to use point estimates of the input embeddings but account for uncertainty in the regression model with simple Bayesian linear regression. Alternatively, we can treat uncertainty in the input embeddings but ignore model uncertainty with the proposed Bayesian mean shrinkage model, which builds on a recently proposed Bayesian nonparametric model of uncertainty in kernel mean embeddings (Flaxman et al., 2016), and then use a sparse representation of the desired function in the RKHS for prediction in the regression model. This model allows for a full account of uncertainty in the mean embedding, but requires a point estimate of the regression function for conjugacy; we thus use backpropagation to obtain a MAP estimate for it as well as various hyperparameters. We then combine the treatment of the two sources

of uncertainty into a fully Bayesian model and use Hamiltonian Monte Carlo for efficient inference. Depending on the inferential goals, each model can be useful. We demonstrate our approaches on an illustrative toy problem as well as a challenging real-world age estimation task.

2 BACKGROUND

2.1 Problem Overview

Distribution regression is the task of learning a classifier or a regression function that maps probability distributions to labels. The challenge of distribution regression goes beyond the standard supervised learning setting: we do not have access to exact input-output pairs since the true inputs, probability distributions, are observed only through samples from that distribution:

$$\left(\{x_j^1\}_{j=1}^{N_1}, y_1\right), \dots, \left(\{x_j^n\}_{j=1}^{N_n}, y_n\right), \quad (1)$$

so that each bag $\{x_j^i\}_{j=1}^{N_i}$ has a label y_i along with N_i individual observations $x_j^i \in \mathcal{X}$. We assume that the observations $\{x_j^i\}_{j=1}^{N_i}$ are i.i.d. samples from some unobserved distribution P_i , and that the true label y_i depends only on P_i . We wish to avoid making any strong parametric assumptions on the P_i . For the present work, we will assume the labels y_i are real-valued; Appendix B shows an extension to binary classification. We typically take the observation space \mathcal{X} to be a subset of \mathbb{R}^p , but it could easily be a structured domain such as text or images, since we access it only through a kernel (for examples, see e.g. Gärtner, 2008).

We consider the standard approach to distribution regression, which relies on kernel mean embeddings and kernel ridge regression. For any positive definite kernel function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, there exists a unique reproducing kernel Hilbert space (RKHS) \mathcal{H}_k , a possibly infinite-dimensional space of functions $f : \mathcal{X} \rightarrow \mathbb{R}$ where evaluation can be written as an inner product, and in particular $f(x) = \langle f, k(\cdot, x) \rangle_{\mathcal{H}_k}$ for all $f \in \mathcal{H}_k, x \in \mathcal{X}$. Here $k(\cdot, x) \in \mathcal{H}_k$ is a function of one argument, $y \mapsto k(y, x)$.

Given a probability measure P on \mathcal{X} , let us define the kernel mean embedding into \mathcal{H}_k as

$$\mu_P = \int k(\cdot, x) P(dx) \in \mathcal{H}_k. \quad (2)$$

Notice that μ_P serves as a high- or infinite-dimensional vector representation of P . For the kernel mean embedding of P into \mathcal{H}_k to be well-defined, it suffices that $\int \sqrt{k(x, x)} P(dx) < \infty$, which is trivially satisfied for all P if k is bounded. Analogously to the reproducing property of RKHS, μ_P represents the expectation function on \mathcal{H}_k : $\int h(x) P(dx) = \langle h, \mu_P \rangle_{\mathcal{H}_k}$. For so-called *characteristic* kernels (Sriperumbudur et al., 2010), every probability measure has a unique embedding, and thus μ_P completely determines the corresponding probability measure.

2.2 Estimating Mean Embeddings

For a set of samples $\{x_j\}_{j=1}^n$ drawn iid from P , the empirical estimator of $\mu_P, \widehat{\mu_P} \in \mathcal{H}_k$, is given by

$$\widehat{\mu_P} = \mu_{\widehat{P}} = \int k(\cdot, x) \widehat{P}(dx) = \frac{1}{n} \sum_{j=1}^n k(\cdot, x_j). \quad (3)$$

This is the standard estimator used by previous distribution regression approaches, which the reproducing property of \mathcal{H}_k shows us corresponds to the kernel

$$\langle \widehat{\mu_{P_i}}, \widehat{\mu_{P_j}} \rangle_{\mathcal{H}_k} = \frac{1}{N_i N_j} \sum_{\ell=1}^{N_i} \sum_{r=1}^{N_j} k(x_\ell^i, x_r^j). \quad (4)$$

But (3) is an empirical mean estimator in a high- or infinite-dimensional space, and is thus subject to the well-known *Stein phenomenon*, so that its performance is dominated by the James-Stein shrinkage estimators. Indeed, Muandet et al. (2014) studied shrinkage estimators for mean embeddings, which can result in substantially improved performance for some tasks (Ramdas and Wehbe, 2015). Flaxman et al. (2016) proposed a Bayesian analogue of shrinkage estimators, which we now review.

This approach consists of (1) a Gaussian Process prior $\mu_P \sim \mathcal{GP}(m_0, r(\cdot, \cdot))$ on \mathcal{H}_k , where r is selected to ensure that $\mu_P \in \mathcal{H}_k$ almost surely and (2) a normal likelihood $\widehat{\mu_P}(\mathbf{x}) \mid \mu_P(\mathbf{x}) \sim \mathcal{N}(\mu_P(\mathbf{x}), \Sigma)$. Here, conjugacy of the prior and the likelihood leads to a Gaussian process posterior on the true embedding μ_P , given that we have observed $\widehat{\mu_P}$ at some set of locations \mathbf{x} . The posterior mean is then essentially identical to a particular shrinkage estimator of Muandet et al. (2014), but the method described here has the extra advantage of a closed form uncertainty estimate, which we utilise in our distributional approach. For the choice of r , we use a Gaussian RBF kernel k , and choose either $r = k$ or, following Flaxman et al. (2016), $r(x, x') = \int k(x, z) k(z, x') \nu(dz)$ where ν is proportional to a Gaussian measure. For details of our choices, and why they are sufficient for our purposes, see Appendix A.

This model accounts for the uncertainty based on the number of samples N_i , shrinking the embeddings for small sample sizes more. As we will see, this is essential in the context of distribution regression, particularly when bag sizes are imbalanced.

2.3 Standard Approaches to Distribution Regression

Following Szábo et al. (2016), assume that the probability distributions P_i are each drawn randomly from some unknown meta-distribution over probability distributions, and take a two-stage approach, illustrated as in Figure 1. Denoting the feature map $k(\cdot, x) \in \mathcal{H}_k$ by $\phi(x)$, one uses the empirical kernel mean estimator (3) to separately estimate

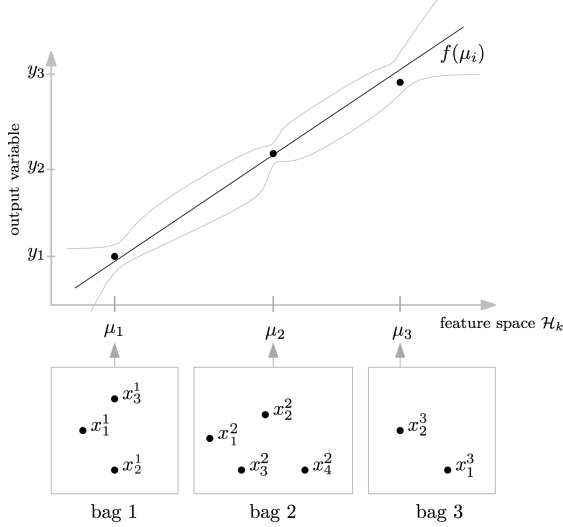


Figure 1: Each bag is summarised by a kernel mean embedding $\mu_i \in \mathcal{H}_k$; a regression function $f : \mathcal{H}_k \rightarrow \mathbb{R}$ predicts labels $y_i \in \mathbb{R}$. We propose a Bayesian approach to propagate uncertainty due to the number of samples in each bag, obtaining posterior credible intervals illustrated in grey.

the mean of each group:

$$\widehat{\mu}_1 = \frac{1}{N_1} \sum_{j=1}^{N_1} \phi(x_j^1), \quad \dots, \quad \widehat{\mu}_n = \frac{1}{N_n} \sum_{i=1}^{N_n} \phi(x_j^n). \quad (5)$$

Next, one uses kernel ridge regression (Saunders et al., 1998) to learn a function $f : \mathcal{H}_k \rightarrow \mathbb{R}$, by minimizing the squared loss with an RKHS complexity penalty:

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{H}_K} \sum_i (y_i - f(\hat{\mu}_i))^2 + \lambda \|f\|_{\mathcal{H}_K}^2.$$

Here $K : \mathcal{H}_k \times \mathcal{H}_k \rightarrow \mathbb{R}$ is a “second-level” kernel on mean embeddings. If K is a linear kernel on the RKHS \mathcal{H}_k , then the resulting method can be interpreted as a linear (ridge) regression on mean embeddings, which are themselves nonlinear transformations of the inputs. A nonlinear second-level kernel on \mathcal{H}_k sometimes improves performance (Muandet et al., 2012; Szábo et al., 2016).

Distribution regression as described is not scalable for even modestly-sized datasets, as computing each of the $\mathcal{O}(n^2)$ entries of the relevant kernel matrix requires time $\mathcal{O}(N_i N_j)$. Many applications have thus used variants of random Fourier features (Rahimi and Recht, 2007). In this paper we instead expand in terms of landmark points drawn randomly from the observations, yielding radial basis networks (Broomhead and Lowe, 1988) with mean pooling.

3 MODELS

We consider here three different Bayesian models, with each model encoding different types of uncertainty. We

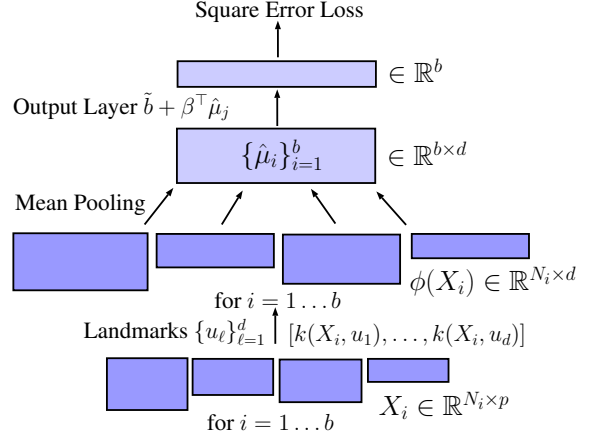


Figure 2: Our baseline model, a RBF network for distribution regression. X_i represents the matrix of samples for bag i , while $k(X_i, u_\ell)$ represents the element wise operation on each row of X_i , with b representing the batch size for stochastic gradient descent.

begin with a non-Bayesian RBF network formulation of the standard approach to distribution regression as a baseline, before refining this approach to better propagate uncertainty in bag size, as well as model parameters.

3.1 Baseline Model

The baseline RBF network formulation we employ here is a variation of the approaches of Broomhead and Lowe (1988), Que and Belkin (2016), Law et al. (2017), and Zacheer et al. (2017). As shown in Figure 2, the initial input is a minibatch consisting of several bags X_i , each containing N_i points. Each point is then converted to an explicit featureisation, taking the role of ϕ in (5), by a radial basis layer: $x_j^i \in \mathbb{R}^p$ is mapped to

$$\phi(x_j^i) = [k(x_j^i, u_1), \dots, k(x_j^i, u_d)]^\top \in \mathbb{R}^d$$

where $\mathbf{u} = \{u_\ell\}_{\ell=1}^d$ are landmark points. A mean pooling layer yields the estimated mean embedding $\hat{\mu}_i$ corresponding to each of the bags j represented in the minibatch, where $\hat{\mu}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} \phi(x_j^i)$.¹ Finally, a fully connected output layer gives real-valued labels $\hat{y}_i = \beta^\top \hat{\mu}_i + b$. As a loss function we use the mean square error $\frac{1}{n} \sum_i (\hat{y}_i - y_i)^2$. For learning, we use backpropagation with the Adam optimizer (Kingma and Ba, 2015). To regularise the network, we use early stopping on a validation set, as well as an L_2 penalty corresponding to a normal prior on β .

¹In the implementation, we stack all of the bags X_i into a single matrix of size $\sum_j N_j \times d$ for the first layer, then perform pooling via sparse matrix multiplication.

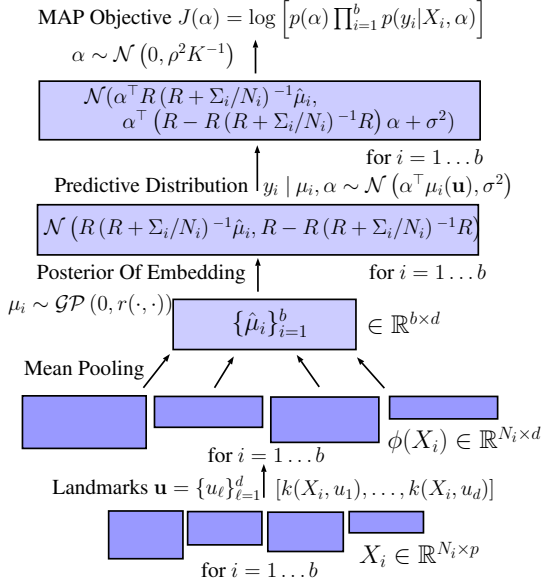


Figure 3: Our Bayesian mean shrinkage pooling model. This diagram takes $m_0 = \mathbf{0}$, $\eta = 1$ and $\mathbf{u} = \mathbf{z}$, so that $R = R_{\mathbf{z}} = R_{\mathbf{z}\mathbf{z}}$, and $K_{\mathbf{z}} = K$.

3.2 Bayesian Linear Regression Model

The most obvious approach to adding uncertainty to the model of Section 3.1 is to encode uncertainty over regression parameters β only, as follows:

$$\beta \sim \mathcal{N}(0, \rho^2) \quad y_i | \mathbf{x}_i, \beta \sim \mathcal{N}(\beta^\top \hat{\mu}_i, \sigma^2).$$

This is essentially Bayesian linear regression on the empirical mean embeddings, and is closely related to the model of Flaxman et al. (2015). Here, we are working directly with the finite-dimensional $\hat{\mu}_i$, unlike the infinite-dimensional μ_i before. Due to the conjugacy of the model, we can easily obtain the predictive distribution $y_i | \mathbf{x}_i$, integrating out the uncertainty over β . This provides us with uncertainty intervals for the predictions y_i .

For model tuning, we can maximise the model evidence, i.e. the marginal log-likelihood (see Bishop (2006) for details), and use backpropagation through the network to learn σ and ρ and any kernel parameters of interest.²

3.3 Bayesian Mean Shrinkage Model

A shortcoming of the prior models, and of the standard approach in Szábo et al. (2016), is that they ignore uncertainty in the first level of estimation due to varying number of samples in each bag. Ideally we would estimate not just the mean embedding per bag, but also a measure of the sample

²Note that unlike the other models considered in this paper, we cannot easily do minibatch stochastic gradient descent, as the marginal log-likelihood does not decompose for each individual data point.

variance, in order to propagate this information regarding uncertainty from the bag size through the model. Bayesian tools provide a natural framework for this problem.

We can use the Bayesian nonparametric prior over kernel mean embeddings (Flaxman et al., 2016) described in Section 2.2, and observe the empirical embeddings at the landmark points \mathbf{u}_i . For \mathbf{u}_i , we take a fixed set of landmarks, which we can choose via k -means clustering or sample without replacement (Que and Belkin, 2016). Using the conjugacy of the model to the Gaussian process prior $\mu_i \sim \mathcal{GP}(m_0, \eta r(\cdot, \cdot))$, we obtain a closed-form posterior Gaussian process whose evaluation at points $\mathbf{h} = \{h_s\}_{s=1}^{n_h}$ is:

$$\mu_i(\mathbf{h}) | \mathbf{x}_i \sim \mathcal{N}\left(R_{\mathbf{h}\mathbf{h}}(R + \Sigma_i/N_i)^{-1}(\hat{\mu}_i - m_0) + m_0, R_{\mathbf{h}\mathbf{h}} - R_{\mathbf{h}}(R + \Sigma_i/N_i)^{-1}R_{\mathbf{h}}^\top\right)$$

where $R_{st} = \eta r(u_s, u_t)$, $(R_{\mathbf{h}\mathbf{h}})_{st} = \eta r(h_s, h_t)$, $(R_{\mathbf{h}})_{st} = \eta r(h_s, u_t)$, and \mathbf{x}_i denotes the set $\{x_j^i\}_{j=1}^{N_i}$. We take the prior mean m_0 to be the average of the $\hat{\mu}_i$; under a linear kernel K , this means we shrink predictions towards the mean prediction. Note η essentially controls the strength of the shrinkage: a smaller η means we shrink more strongly towards m_0 . We take Σ_i to be the average of the empirical covariance of $\{\varphi(x_j^i)\}_{j=1}^{N_i}$ across all bags, to avoid poor estimation of Σ_i for smaller bags. More intuition about the behaviour of this estimator can be found in Appendix C.

Now, supposing we have normal observation error σ^2 , and use a linear kernel as our second level kernel K , we have:

$$y_i | \mu_i, f \sim \mathcal{N}(f, \mu_i | \mathcal{H}_k, \sigma^2) \quad (6)$$

where $f \in \mathcal{H}_k$. Clearly, this is difficult to work with; hence we parameterise f as $f = \sum_{\ell=1}^s \alpha_\ell k(\cdot, z_\ell)$, where $\mathbf{z} = \{z_\ell\}_{\ell=1}^s$ is a set of landmark points for f , which we can learn or fix. (Appendix D gives a motivation for this approximation using the representer theorem.) Using the reproducing property, our likelihood model becomes:

$$y_i | \mu_i, \alpha \sim \mathcal{N}(\alpha^\top \mu_i(\mathbf{z}), \sigma^2) \quad (7)$$

where $\mu_i(\mathbf{z}) = [\mu_i(z_1), \dots, \mu_i(z_s)]^\top$. For fixed α and \mathbf{z} we can analytically integrate out the dependence on μ_i , and the predictive distribution of a bag label becomes

$$y_i | \mathbf{x}_i, \alpha \sim \mathcal{N}(\xi_i^\alpha, \nu_i^\alpha)$$

$$\xi_i^\alpha = \alpha^\top R_{\mathbf{z}} \left(R + \frac{\Sigma_i}{N_i}\right)^{-1}(\hat{\mu}_i - m_0) + \alpha^\top m_0$$

$$\nu_i^\alpha = \alpha^\top \left(R_{\mathbf{z}\mathbf{z}} - R_{\mathbf{z}} \left(R + \frac{\Sigma_i}{N_i}\right)^{-1} R_{\mathbf{z}}^\top\right) \alpha + \sigma^2.$$

The prior $\alpha \sim \mathcal{N}(0, \rho^2 K_{\mathbf{z}}^{-1})$, where $K_{\mathbf{z}}$ is the kernel matrix on \mathbf{z} , gives the standard regularisation on f of $\|f\|_{\mathcal{H}_k}^2$. The log-likelihood objective becomes

$$\frac{1}{2} \sum_{i=1}^n \left\{ \log \nu_i^\alpha + \frac{(y_i - \xi_i^\alpha)^2}{\nu_i^\alpha} \right\} + \frac{\alpha^\top K_{\mathbf{z}} \alpha}{2\rho^2}.$$

We can use backpropagation to learn the parameters α , σ , and if we wish η , \mathbf{z} , and any kernel parameters. The full model is illustrated in Figure 3. This approach allows us to directly encode uncertainty based on bag size in the objective function, and gives probabilistic predictions.

3.4 Bayesian Distribution Regression

It is natural to combine the two Bayesian models above, fully propagating uncertainty in estimation of the mean embedding and of the regression coefficients α . Unfortunately, conjugate Bayesian inference is no longer available. Thus, we consider a Markov Chain Monte Carlo (MCMC) sampling based approach, and here use Hamiltonian Monte Carlo (HMC) for efficient inference, though any MCMC-type scheme would work. Whereas inference above used gradient descent to maximise the marginal likelihood, with the gradient calculated using automatic differentiation, here we use automatic differentiation to calculate the gradient of the joint log-likelihood and follow this gradient as we perform sampling over the parameters we wish to infer.

We can still exploit the conjugacy of the mean shrinkage layer, obtaining an analytic posterior over the mean embeddings. Conditional on the mean embeddings, we have a Bayesian linear regression model with parameters α . We sample this model with the NUTS HMC sampler (Hoffman and Gelman, 2014; Stan Development Team, 2014).

4 RELATED WORK

As previously mentioned, Szábo et al. (2016) provides a thorough learning-theoretic analysis of the regression model discussed in Section 2.3. This formalism considering a kernel method on distributions using their embedding representations, or various scalable approximations to it, has been widely applied (e.g. Muandet et al., 2012; Yoshikawa et al., 2014; Flaxman et al., 2015; Jitkrittum et al., 2015; Lopez-Paz et al., 2015; Mitrovic et al., 2016). There are also several other notions of similarities on distributions in use (not necessarily falling within the framework of kernel methods and RKHSs), as well as local smoothing approaches, mostly based on estimates of various probability metrics (Moreno et al., 2003; Jebara et al., 2004; Póczos et al., 2011; Oliva et al., 2013; Poczos et al., 2013; Kusner et al., 2015). For a partial overview, see Sutherland (2016).

Other related problems of learning on instances with group-level labels include learning with label proportions (Quadrianto et al., 2009; Patrini et al., 2014), ecological inference (King, 1997; Gelman et al., 2001), pointillistic pattern search (Ma et al., 2015), multiple instance learning (Dietterich et al., 1997; Kück and de Freitas, 2005; Zhou et al., 2009; Krummenacher et al., 2013) and learning with sets (Zaheer et al., 2017).³

There have also been some Bayesian approaches in related contexts, though most do not follow our setting where the label is a function of the underlying distribution rather than the observed sample set. Kück and de Freitas (2005) consider an MCMC method with group-level labels but focus on individual-level classifiers, while Jackson et al. (2006) use hierarchical Bayesian models on both individual-level and aggregate data for ecological inference.

Jitkrittum et al. (2015) and Flaxman et al. (2015) quantify the uncertainty of distribution regression models by interpreting the kernel ridge regression on embeddings as Gaussian process regression. However, the former’s setting has no uncertainty in the mean embeddings, while the latter’s treats empirical embeddings as fixed inputs to the learning problem (as in Section 3.2).

There has also been generic work on input uncertainty in Gaussian process regression (Girard, 2004; Damianou et al., 2016). These methods could provide a framework towards allowing for second-level kernels in our models. One could also, though, consider regression with uncertain inputs as a special case of distribution regression, where the label is a function of the distribution’s mean and $N_i = 1$.

5 EXPERIMENTS

We will now demonstrate our various Bayesian approaches: the mean-shrinkage pooling method with $r = k$ (*shrinkage*) and with $r(x, x') = \int k(x, z)k(z, x')\nu(dz)$ for ν proportional to a Gaussian measure (*shrinkageC*), Bayesian linear regression (*BLR*), and the full Bayesian distribution regression model with $r = k$ (*BDR*). We also compare the non-Bayesian baselines *RBF network* (Section 3.1) and *freq-shrinkage*, which uses the shrinkage estimator of Muandet et al. (2014) to estimate mean embeddings. Code for our methods and to reproduce the experiments is available at <https://github.com/hcllaw/bdr>.

We first demonstrate the characteristics of our models on a synthetic dataset, and then evaluate them on a real life age prediction problem. Throughout, for simplicity, we take $\mathbf{u} = \mathbf{z}$, i.e. $R = R_{\mathbf{z}} = R_{\mathbf{z}\mathbf{z}}$, and $K_{\mathbf{z}} = K$ – although \mathbf{u} and \mathbf{z} could be different, with \mathbf{z} learnt. Here k is the standard RBF kernel. We tune the learning rate, number of landmarks, bandwidth of the kernel and regularisation parameters on a validation set. For BDR, we use weakly informative normal priors (possibly truncated at zero); for other models, we learn the remaining parameters.

5.1 Gamma Synthetic Data

We create a synthetic dataset by repeatedly sampling from the following hierarchical model, where y_i is the label for the i th bag, each $x_j^i \in \mathbb{R}^5$ has entries i.i.d. according to the given distribution, and ε is an added noise term which

³For more, also see giorgiopatrini.org/nips15workshop.

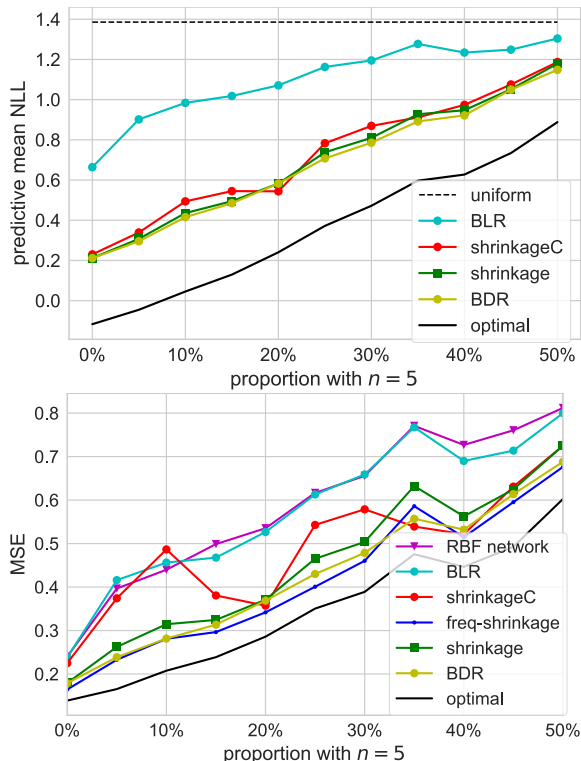


Figure 4: Top: negative log-likelihood. Bottom: mean-squared error. For context, performance of the Bayes-optimal predictor is also shown, and for NLL ‘uniform’ shows the performance of a uniform prediction on the possible labels. For MSE, the constant overall mean label predictor achieves about 1.3.

differs for the two experiments below:

$$y_i \sim \text{Uniform}(4, 8)$$

$$[x_j^i]_\ell \mid y_i \stackrel{\text{iid}}{\sim} \frac{1}{y_i} \left[\Gamma \left(\frac{y_i}{2}, \frac{1}{2} \right) \right] + \varepsilon \text{ for } j \in [N_i], \ell \in [5].$$

In these experiments, we generate 1000 bags for training, 500 bags for a validation set for parameter tuning, 500 bags to use for early-stopping of the models, and 1000 bags for testing. Tuning is performed to maximize log-likelihoods for Bayesian models, MSE for non-Bayesian models. Landmark points \mathbf{u} are chosen via k -means (fixed across all models). We also show results of the Bayes-optimal model, which gives true posteriors according to the data-generating process; this is the best performance any model could hope to achieve. Our learning models, which treat the inputs as five-dimensional, fully nonparametric distributions, are at a substantial disadvantage even in how they view the data compared to this true model.

Varying bag size: Uncertainty in the inputs. In order to study the behaviour of our models with varying bag size, we fix four sizes $N_i \in \{5, 20, 100, 1000\}$. For each gener-

ated dataset, 25% of the bags have $N_i = 20$, and 25% have $N_i = 100$. Among the other half of the data, we vary the ratio of $N_i = 5$ and $N_i = 1000$ bags to demonstrate the methods’ efficacy at dealing with varied bag sizes: we let s_5 be the overall percentage of bags with $N_i = 5$, ranging from $s_5 = 0$ (in which case no bags have size $N_i = 5$) to $s_5 = 50$ (in which case 50% of the overall bags have size $N_i = 5$). Here we do not add additional noise: $\varepsilon = 0$.

Results are shown in Figure 4. BDR and shrinkage methods, which take into account bag size uncertainty, perform well here compared to the other methods. The full BDR model very slightly outperforms the Bayesian shrinkage models in both likelihood and in mean-squared error; frequentist shrinkage slightly outperforms the Bayesian shrinkage models in MSE probably because it is tuned for that metric. We also see that the choice of r affects the results; $r = k$ does somewhat better.

Figure 5 demonstrates in more detail the difference between these models. It shows test set predictions of each model on the bags of different sizes. Here, we can see explicitly that the shrinkage and BDR models are able to take into account the bag size, with decreasing variance for larger bag sizes, while the BLR model gives the same variance for all outputs. Furthermore, the shrinkage and BDR models can shrink their predictions towards the mean more for smaller bags than larger ones: this improves performance on the small bags while still allowing for good predictions on large bags, contrary to the BLR model.

Fixed bag size: Uncertainty in the regression model.

The previous experiment showed the efficacy of the shrinkage estimator in our models, but demonstrated little gain from posterior inference for regression weights β over their MAP estimates, i.e. there is no discernible improvement of BLR over RBF network. To isolate the effect of quantifying uncertainty in the regression model, we now consider the case where there is no variation in bag size at all and normal noise is added onto the observations. In particular we take $N_i = 1000$ and $\varepsilon \sim \mathcal{N}(0, 1)$, and sample landmarks randomly from the training set.

Results are shown in Table 1. Here, BLR or BDR outperform all other methods on all runs, highlighting that uncertainty in the regression model is also important for predictive performance. Importantly, the BDR method performs well in this regime as well as in the previous one.

5.2 IMDb-WIKI: Age Estimation

We now demonstrate our methods on a celebrity age estimation problem, using the IMDb-WIKI database (Rothe et al., 2016) which consists of 397 949 images of 19 545 celebrities⁴, with corresponding age labels. This database

⁴We used only the IMDb images, and removed some implausible images, including one of a cat and several of people with

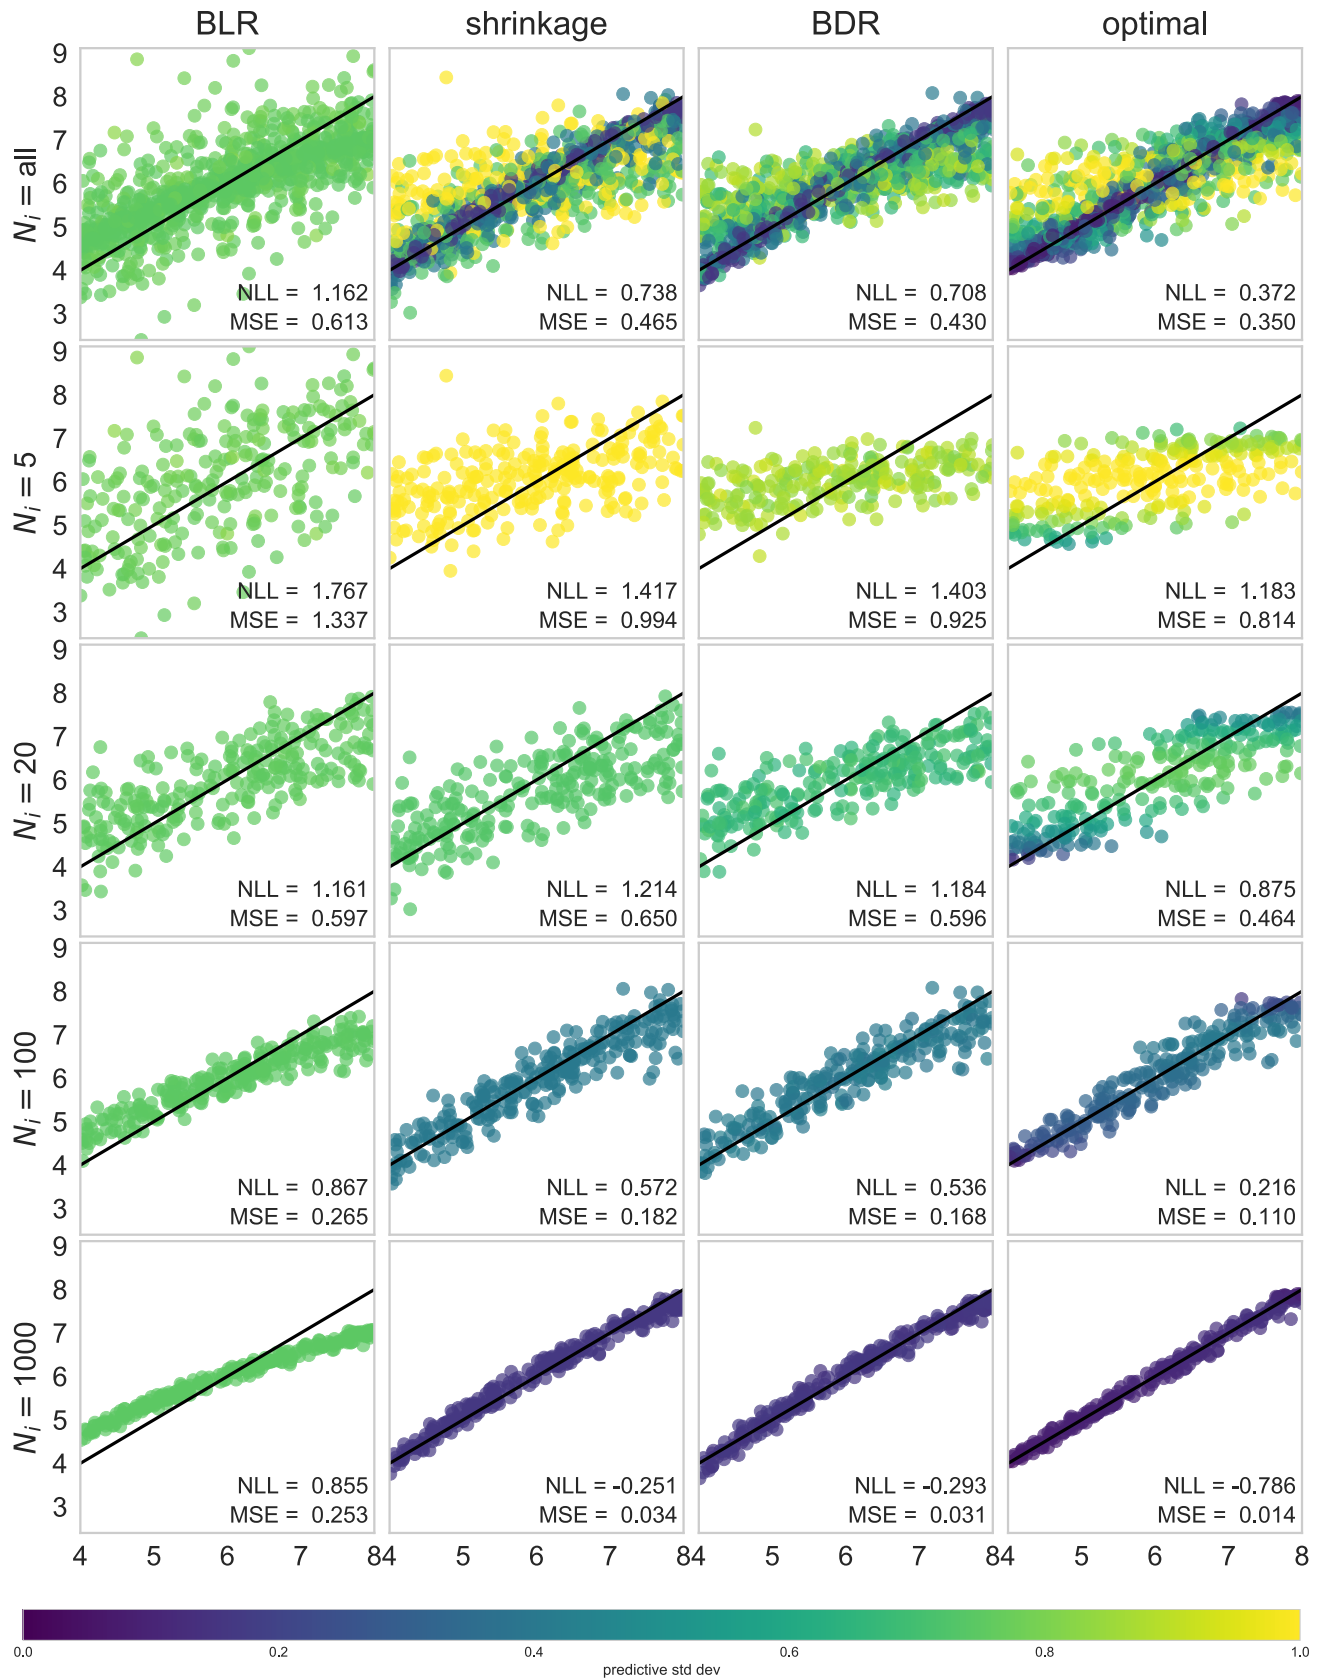


Figure 5: Predictions for the varying bag size experiment of Section 5.1. Each column corresponds to a single prediction method. Each point in an image represents a single bag, with its horizontal position the true label y_i , and its vertical position the predicted label. The black lines show theoretical perfect predictions. The rows represent different subsets of the data: the first row shows all bags, the second only bags with $N_i = 5$, and so on. Colours represent the predictive

Table 1: Results on the fixed bag size dataset, over 10 dataset draws (standard deviations in parentheses). BLR/BDR perform best on all runs in both metrics.

METHOD	MSE	NLL
Optimal	0.170 (0.009)	0.401 (0.018)
RBF network	0.235 (0.014)	–
freq-shrinkage	0.232 (0.012)	–
shrinkage	0.237 (0.014)	0.703 (0.027)
shrinkageC	0.236 (0.013)	0.700 (0.029)
BLR	0.228 (0.012)	0.681 (0.025)
BDR	0.227 (0.012)	0.683 (0.025)

Table 2: Results on the grouped IMDB-WIKI dataset over ten runs (standard deviations in parentheses). Here shrinkage methods perform the best across all 10 runs.

METHOD	RMSE	NLL
CNN	10.25 (0.22)	3.80 (0.034)
RBF network	9.51 (0.20)	–
freq-shrinkage	9.22 (0.19)	–
shrinkage	9.28 (0.20)	3.54 (0.021)
BLR	9.55 (0.19)	3.68 (0.021)

was constructed by crawling IMDB for images of its most popular actors and directors, with potentially many images for each celebrity over time. Rothe et al. (2016) use a convolutional neural network (CNN) with a VGG-16 architecture to perform 101-way classification, with one class corresponding to each age in $\{0, \dots, 100\}$.

We take a different approach, and assume that we are given several images of a single individual (i.e. samples from the distribution of celebrity images), and are asked to predict their mean age based on several pictures. For example, we have 757 images of Brad Pitt from age 27 up to 51, while we have only 13 images of Chelsea Peretti at ages 35 and 37. Note that 22.5% of bags have only a single image. We obtain 19 545 bags, with each bag containing between 1 and 796 images of a particular celebrity, and the corresponding bag label calculated from the average of the age labels of the images inside each bag.

In particular, we use the representation $\varphi(x)$ learnt by the CNN in Rothe et al. (2016), where $\varphi(x) : \mathbb{R}^{256 \times 256} \rightarrow \mathbb{R}^{4096}$ maps from the pixel space of images to the CNN’s last hidden layer. With these new representations, we can now treat them as inputs to our radial basis network, shrinkage (taking $r = k$ here) and BLR models. Although we could also use the full BDR model here, due to the computational time and memory required to perform proper pa-

rameter tuning, we relegate this to a later study.

We use 9820 bags for training, 2948 bags for early stopping, 2946 for validation and 3928 for testing. Landmarks are sampled without replacement from the training set.

We repeat the experiment on 10 different splits of the data, and report the results in Table 2. The baseline CNN results give performance by averaging the predictive distribution from the model of Rothe et al. (2016) for each image of a bag; note that this model was trained on all of the images used here. From Table 2, we can see that the shrinkage methods have the best performance; they outperforms all other methods in all 10 splits of the dataset, in both metrics. Non-Bayesian shrinkage again yields slightly better RMSEs, probably because it is tuned for that metric. This demonstrates that modelling bag size uncertainty is vital.

6 CONCLUSION

Supervised learning on groups of observations using kernel mean embeddings typically disregards sampling variability within groups. To handle this problem, we construct Bayesian approaches to modelling kernel mean embeddings within a regression model, and investigate advantages of uncertainty propagation within different components of the resulting distribution regression. The ability to take into account the uncertainty in mean embedding estimates is demonstrated to be key for constructing models with good predictive performance when group sizes are highly imbalanced. We also demonstrate that the results of a complex neural network model for age estimation can be improved by shrinkage.

Our models employ a neural network formulation to provide more expressive feature representations and learn discriminative embeddings. Doing so makes our model easy to extend to more complicated featurisations than the simple RBF network used here. By training with backpropagation, or via approximate Bayesian methods such as variational inference, we can easily ‘learn the kernel’ within our framework, for example fine-tuning the deep network of Section 5.2 rather than using a pre-trained model. We can also apply our networks to structured settings, learning regression functions on sets of images, audio, or text. Such models naturally fit into the empirical Bayes framework.

On the other hand, we might extend our model to more Bayesian feature learning by placing priors over the kernel hyperparameters, building on classic work on variational approaches (Barber and Schottky, 1998) and fully Bayesian inference (Andrieu et al., 2001) in RBF networks. Such approaches are also possible using other featurisations, e.g. random Fourier features (as in Oliva et al., 2015).

Future distribution regression approaches will need to account for uncertainty in observation of the distribution. Our methods provide a strong, generic building block to do so.

supposedly negative age, or ages of several hundred years.

References

- Christophe Andrieu, Nando De Freitas, and Arnaud Doucet. Robust full bayesian learning for radial basis networks. *Neural Computation*, 13(10):2359–2407, 2001.
- David Barber and Bernhard Schottky. Radial basis functions: a bayesian treatment. *NIPS*, pages 402–408, 1998.
- C.M. Bishop. *Pattern recognition and machine learning*. Springer New York, 2006.
- David S Broomhead and David Lowe. Radial basis functions, multi-variable functional interpolation and adaptive networks. Technical report, DTIC Document, 1988.
- Andreas C. Damianou, Michalis K. Titsias, and Neil D. Lawrence. Variational inference for latent variables and uncertain inputs in Gaussian processes. *JMLR*, 17(42):1–62, 2016.
- Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 99(1):31–71, 1997.
- Seth Flaxman, Yu-Xiang Wang, and Alexander J Smola. Who supported Obama in 2012?: Ecological inference through distribution regression. In *KDD*, pages 289–298. ACM, 2015.
- Seth Flaxman, Dino Sejdinovic, John P. Cunningham, and Sarah Filippi. Bayesian learning of kernel embeddings. In *UAI*, 2016.
- Seth Flaxman, Dougal J. Sutherland, Yu-Xiang Wang, and Yee-Whye Teh. Understanding the 2016 US presidential election using ecological inference and distribution regression with census microdata. 2016. arXiv:1611.03787.
- Thomas Gärtner. *Kernels for Structured Data*, volume 72. World Scientific, Series in Machine Perception and Artificial Intelligence, 2008.
- Andrew Gelman, David K Park, Stephen Ansolabehere, Phillip N Price, and Lorraine C Minnite. Models, assumptions and model checking in ecological regressions. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 164(1):101–118, 2001.
- Agathe Girard. *Approximate methods for propagation of uncertainty with Gaussian process models*. PhD thesis, University of Glasgow, 2004.
- Matthew D. Hoffman and Andrew Gelman. The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *JMLR*, pages 1593–1623, 2014.
- Christopher Jackson, Nicky Best, and Sylvia Richardson. Improving ecological inference using individual-level data. *Statistics in medicine*, 25(12):2136–2159, 2006.
- Tony Jebara, Risi Imre Kondor, and Andrew Howard. Probability product kernels. *JMLR*, 5:819–844, 2004.
- Wittawat Jitkrittum, Arthur Gretton, Nicolas Heess, S. M. Ali Eslami, Balaji Lakshminarayanan, Dino Sejdinovic, and Zoltán Szabó. Kernel-Based Just-In-Time Learning for Passing Expectation Propagation Messages. In *UAI*, 2015.
- Gary King. *A Solution to the Ecological Inference Problem*. Princeton University Press, 1997. ISBN 0691012407.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. arXiv:1412.6980.
- Gabriel Krummenacher, Cheng Soon Ong, and Joachim M Buhmann. Ellipsoidal multiple instance learning. In *ICML (2)*, pages 73–81, 2013.
- Hendrik Kück and Nando de Freitas. Learning about individuals from group statistics. In *UAI*, pages 332–339, 2005.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. From word embeddings to document distances. In *ICML*, pages 957–966, 2015.
- H. C. L. Law, C. Yau, and D. Sejdinovic. Testing and learning on distributions with symmetric noise invariance. In *NIPS*, 2017. arXiv:1703.07596.
- David Lopez-Paz, Krikamol Muandet, Bernhard Schölkopf, and Ilya Tolstikhin. Towards a learning theory of cause-effect inference. In *ICML*, 2015.
- Milan Lukić and Jay Beder. Stochastic processes with sample paths in reproducing kernel hilbert spaces. *Transactions of the American Mathematical Society*, 353(10):3945–3969, 2001.
- Yifei Ma, Dougal J. Sutherland, Roman Garnett, and Jeff Schneider. Active pointillistic pattern search. In *AIS-TATS*, 2015.
- J. Mitrovic, D. Sejdinovic, and Y.W. Teh. DR-ABC: Approximate Bayesian Computation with Kernel-Based Distribution Regression. In *ICML*, pages 1482–1491, 2016.
- Pedro J Moreno, Purdy P Ho, and Nuno Vasconcelos. A Kullback-Leibler divergence based kernel for SVM classification in multimedia applications. In *NIPS*, 2003.
- Krikamol Muandet, Kenji Fukumizu, Francesco Dinuzzo, and Bernhard Schölkopf. Learning from distributions via support measure machines. In *NIPS*, 2012. arXiv:1202.6504.
- Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, Arthur Gretton, and Bernhard Schoelkopf. Kernel mean estimation and stein effect. In *ICML*, 2014.
- Michelle Ntampaka, Hy Trac, Dougal J. Sutherland, Nicholas Battaglia, Barnabás Póczos, and Jeff Schneider. A machine learning approach for dynamical mass measurements of galaxy clusters. *The Astrophysical Journal*, 803(2):50, 2015. ISSN 1538-4357. arXiv:1410.0686.

- Michelle Ntampaka, Hy Trac, Dougal J. Sutherland, S. Fromenteau, B. Poczos, and Jeff Schneider. Dynamical mass measurements of contaminated galaxy clusters using machine learning. *The Astrophysical Journal*, 831(2):135, 2016. arXiv:1509.05409.
- Junier B Oliva, Barnabás Póczos, and Jeff Schneider. Distribution to distribution regression. In *ICML*, 2013.
- Junier B Oliva, Avinava Dubey, Barnabás Póczos, Jeff Schneider, and Eric P Xing. Bayesian nonparametric kernel-learning. In *AISTATS*, 2015. arXiv:1506.08776.
- Giorgio Patrini, Richard Nock, Tiberio Caetano, and Paul Rivera. (Almost) no label no cry. In *NIPS*. 2014.
- Natesh S Pillai, Qiang Wu, Feng Liang, Sayan Mukherjee, and Robert L Wolpert. Characterizing the function space for bayesian kernel models. *JMLR*, 8(Aug):1769–1797, 2007.
- Barnabás Póczos, Liang Xiong, and Jeff Schneider. Non-parametric divergence estimation with applications to machine learning on distributions. In *UAI*, 2011.
- Barnabas Poczos, Aarti Singh, Alessandro Rinaldo, and Larry Wasserman. Distribution-free distribution regression. In *AISTATS*, pages 507–515, 2013. arXiv:1302.0082.
- Novi Quadrianto, Alex J Smola, Tiberio S Caetano, and Quoc V Le. Estimating labels from label proportions. *JMLR*, 10:2349–2374, 2009.
- Qichao Que and Mikhail Belkin. Back to the future: Radial basis function networks revisited. In *AISTATS*, 2016.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *NIPS*, pages 1177–1184, 2007.
- Aaditya Ramdas and Leila Wehbe. Nonparametric independence testing for small sample sizes. In *IJCAI*, 2015. arXiv:1406.1922.
- Rasmus Rothe, Radu Timofte, and Luc Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision (IJCV)*, July 2016.
- Craig Saunders, Alexander Gammerman, and Volodya Vovk. Ridge regression learning algorithm in dual variables. In *ICML*, 1998.
- Bernhard Schölkopf, Ralf Herbrich, and Alex J. Smola. A generalized representer theorem. In *COLT*, 2001.
- Bharath K Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert RG Lanckriet. Hilbert space embeddings and metrics on probability measures. *JMLR*, 99:1517–1561, 2010.
- Stan Development Team. Stan: A c++ library for probability and sampling, version 2.5.0, 2014. URL <http://mc-stan.org/>.
- Ingo Steinwart. Convergence types and rates in generic Karhunen-Loève expansions with applications to sample path properties. *arXiv preprint arXiv:1403.1040v3*, March 2017.
- Dougal J. Sutherland. *Scalable, Flexible, and Active Learning on Distributions*. PhD thesis, Carnegie Mellon University, 2016.
- Zoltán Szábo, Bharath K. Sriperumbudur, Barnabás Póczos, and Arthur Gretton. Learning theory for distribution regression. *JMLR*, 17(152):1–40, 2016. arXiv:1411.2066.
- Grace Wahba. *Spline models for observational data*, volume 59. Siam, 1990.
- Yuya Yoshikawa, Tomoharu Iwata, and Hiroshi Sawada. Latent support measure machines for bag-of-words data classification. In *NIPS*, pages 1961–1969, 2014.
- Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, and Alexander Smola. Deep sets. In *NIPS*, 2017.
- Zhi-Hua Zhou, Yu-Yin Sun, and Yu-Feng Li. Multi-instance learning by treating instances as non-iid samples. In *ICML*, 2009.