

Adaptive Basis Function Selection for Computationally Efficient Predictions

Anton Kullberg, Frida Viset, Isaac Skog and Gustaf Hendeby

The self-archived postprint version of this journal article is available at Linköping University Institutional Repository (DiVA):

<https://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-206535>

N.B.: When citing this work, cite the original publication.

Kullberg, A., Viset, F., Skog, I., Hendeby, G., (2024), Adaptive Basis Function Selection for Computationally Efficient Predictions, *IEEE Signal Processing Letters*
<https://doi.org/10.1109/lsp.2024.3445272>

Original publication available at:

<https://doi.org/10.1109/lsp.2024.3445272>

©2024 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

<https://www.ieee.org/>

Adaptive Basis Function Selection for Computationally Efficient Predictions

Anton Kullberg, Frida Viset, Isaac Skog, *Senior Member, IEEE*, and Gustaf Hendeby, *Senior Member, IEEE*.

Abstract—*Basis Function* (BF) expansions are a cornerstone of any engineer’s toolbox for computational function approximation which shares connections with both neural networks and Gaussian processes. Even though BF expansions are an intuitive and straightforward model to use, they suffer from quadratic computational complexity in the number of BFs if the predictive variance is to be computed. We develop a method to automatically select the most important BFs for prediction in a sub-domain of the model domain. This significantly reduces the computational complexity of computing predictions while maintaining predictive accuracy. The proposed method is demonstrated using two numerical examples, where reductions up to 50–75% are possible without significantly reducing the predictive accuracy.

I. INTRODUCTION

THE focus of this paper is on *Basis Function* (BF) expansion models, as these are prevalent in a wide variety of disciplines [1]–[9]. The canonical BF model is given by

$$f(x) = \sum_{i=1}^L \phi_i(x)\theta_i, \quad (1)$$

where ϕ_i is BF i and θ_i is its corresponding weight. Note that ϕ_1, \dots, ϕ_L may be interpreted as the bases for $f(\cdot)$. For many applications, such as magnetic field mapping, weather forecasting, etc., the domain of the function $f(\cdot)$ can be vast, such that the number of BFs L needs to be very large to represent the underlying phenomena [10]–[12]. The downside of this is that the predictive model can become very computationally demanding to use. For instance, magnetic fields modeled with BF expansions can be used to improve indoor navigation [5], but unless the domain is split into patches to remedy the scaling issues, the online computational requirements scale quadratically with L . Equivalent models can also be used for multi-agent navigation and motion planning in large nonlinear fields, but the required communication between each pair of agents at each time instance also scales quadratically with L [13], [14]. In other applications, such as tire friction modeling in autonomous vehicles [4], L must usually be kept very small so that the computational budget is not exceeded at the expense of the predictive accuracy.

The aforementioned cases are limited for computational reasons, i.e., computing predictions becomes computationally

expensive as L grows larger. One solution to remedy this, typically done a priori, is to reduce L . This results in a reduction of the model’s representative power, effectively restricting the function space that it spans. Another solution is to tailor the selection of the BFs ϕ_i to the particular problem at hand, such that L can be kept small. This requires specialized domain knowledge, which is not always trivial nor feasible to obtain. A third alternative is to consider so-called compact *Radial Basis Functions* (RBFs) [10], [11], [15], [16] that exhibit favorable computational properties at the cost of representational power.

The aforementioned alternatives deal with this computational problem as an *a priori* problem. In contrast, we will approach the model reduction problem as an *a posteriori* problem. To that end, we seek to determine which BFs are necessary to represent a given BF expansion up to some tolerance, effectively reducing the model order purely for predictive purposes. Hence, the weights of the BF expansion are already learned and our task is to compress the model at prediction time so that it becomes computationally cheaper to use, without sacrificing too much accuracy. That is, the model order L used during the learning phase is still kept large, but the *predictive* model uses only a few ϕ_i corresponding to the most important BFs for the specific test input at hand. We remark that while hyperparameter optimization, in, e.g., RBF expansion models [17], [18], may be interpreted as a type of BF selection, this is done during training. It thus requires training data and can furthermore not adapt to the test input.

In this paper, we present an approach for prediction time BF selection that

- 1) Avoids using specialized domain knowledge to select BFs a priori and defers that selection to a posteriori.
- 2) Only uses the BFs that are important for the test input, thereby adapting the predictive model to the test input while reducing the computational complexity as compared to using the full model.
- 3) Does not require additional data for selecting relevant BFs, in contrast to previous approaches.

We require a BF selection process that has low computational complexity, such that the overall complexity of computing the predictions is lower than that of the full model. Test input adaptivity is handled by optimally reducing the model in a sub-domain Ω , that covers the test input of interest at prediction time. The approach is developed for a general BF expansion model and extended to a commonly used sparse *Gaussian Process* (GP) approximation. Thus, our contribution is a computationally efficient, data-free, method of prediction time BF selection that adapts to the test input.

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Anton Kullberg and Gustaf Hendeby are with Linköping University, Sweden. (email: firstname.lastname@liu.se). Frida Viset is with Delft University of Technology, Netherlands. (email: F.M.Viset@tudelft.nl). Isaac Skog is with KTH Royal Institute of Technology, Sweden. (email: skog@kth.se).

Source code available: <https://github.com/AOKullberg/adaptive-bf-selection>

II. BASIS FUNCTION SELECTION

Without loss of generality, we assume that ϕ_1, \dots, ϕ_L are ordered in some way, such that reducing L corresponds to removing particular components of the BF expansion (1). Conceptually, this may be thought of as a Fourier series expansion where the removal of ϕ_i corresponds to removing some particular frequency in the representation. We first describe our selection procedure in the deterministic case and then move to the stochastic case where $\theta \sim p(\theta)$.

A. Adaptive basis function selection

The problem is to select a subset of the bases of the model $f(x)$ and use this subset as our predictive model, i.e.,

$$\hat{f}(x) = \sum_{j \in \mathcal{J}} \phi_j(x) \theta_j. \quad (2)$$

Here, \mathcal{J} is a set of indices, corresponding to the particular ϕ_j , that best represents the original model in some sense. Note that $|\mathcal{J}| = n_J < L$.

1) *Finding \mathcal{J}* : Formally, in the deterministic sense, finding $\hat{f}(x)$ can be formulated as an optimization problem given by

$$\min_{\mathcal{J}} \|f(x) - \hat{f}(x)\|_2^2 \triangleq \mathcal{L} \quad (3a)$$

$$\text{s.t.} \quad |\mathcal{J}| = n_J < L. \quad (3b)$$

Here, $f(\cdot)$ and $\hat{f}(\cdot)$ are given by (1) and (2), respectively. Further, $\|g(x)\|_2^2 = \int_{\Omega} |g(x)|^2 dx$. Note that (3) does *not* imply that the estimator (2) is going to be *Root Mean Square Error* (RMSE) optimal with respect to the true underlying function. It only means that $\hat{f}(\cdot)$ will be a good estimator of $f(\cdot)$. Hence, there may be another model in the model class of $\hat{f}(\cdot)$ that performs better than $\hat{f}(\cdot)$, but finding this would require access to the training data.

The problem (3) is reminiscent of a subset/feature selection problem [19]–[21]. Typical subset selection problems are solved by either greedy inclusion based on some heuristic or essentially by regularizing the optimization problem through, e.g., the *Least Absolute Shrinkage and Selection Operator* (LASSO) [22]. Further, typical solutions are focused on finding the best possible predictive model given a particular cardinality of the set \mathcal{J} . Thus, the general problem is i) training data dependent, ii) requires a fixed cardinality, and iii) not test-input dependent. There are subtle differences between the general subset selection problem and (3) that allow us to i) solve (3) efficiently without access to the training data, ii) perform automatic cardinality selection, and iii) adapt the predictive model to the test data.

To find the solution to (3), notice that \mathcal{L} can be written as

$$\begin{aligned} \mathcal{L} &= \left\| \sum_{i=1}^L \phi_i(x) \theta_i - \sum_{j \in \mathcal{J}} \phi_j(x) \theta_j \right\|_2^2 \\ &= \left\| \sum_{j \notin \mathcal{J}} \phi_j(x) \theta_j \right\|_2^2 = \int_{\Omega} \left| \sum_{j \notin \mathcal{J}} \phi_j(x) \theta_j \right|^2 dx. \quad (4) \end{aligned}$$

The second step follows by definition and corresponds to the norm of the omitted components. The third step is simply the definition of the norm over some domain Ω . Note that

the domain Ω is a user-defined choice, essentially telling us where we want the model to be well-approximated. Now, we can bound the integrand by

$$\left| \sum_{j \notin \mathcal{J}} \phi_j(x) \theta_j \right|^2 \leq \sum_{j \notin \mathcal{J}} |\phi_j(x) \theta_j|^2 = \sum_{j \notin \mathcal{J}} |\phi_j(x)|^2 |\theta_j|^2, \quad (5)$$

by using the triangle inequality. Thus, we can bound \mathcal{L} by

$$\mathcal{L} \leq \sum_{j \notin \mathcal{J}} \int_{\Omega} |\phi_j(x)|^2 |\theta_j|^2 dx = \sum_{j \notin \mathcal{J}} \int_{\Omega} |\phi_j(x)|^2 dx |\theta_j|^2. \quad (6)$$

We can use this bound to select important BFs on a particular domain Ω , which may be a subset of the domain of the BF expansion. This allows us to adaptively select the BFs that contribute the most to the prediction *at* the particular test input we are interested in.

For models where $\int_{\Omega} |\phi_j(x)|^2 dx \approx \int_{\Omega} |\phi_i(x)|^2 dx, \forall i, j$ the bound can be further simplified to

$$\mathcal{L} \leq C \sum_{j \notin \mathcal{J}} |\theta_j|^2, \quad (7)$$

where $C = \int_{\Omega} |\phi_j(x)|^2 dx$. Thus, in this case, to minimize \mathcal{L} , we can simply select the BFs with the largest absolute weights.

The bounds (6) and (7) can be used in two ways. Firstly, given a cardinality n_J , they allow us to quickly choose the n_J BFs contributing most to the accuracy of the predictive model at the test input at hand. Secondly, they allow us to adaptively select the number of BFs used when predicting in a domain Ω such that the approximation error is smaller than (6). To be more precise, assume w.l.o.g. that θ is sorted such that $|\theta_1| < |\theta_2| < \dots < |\theta_L|$. Starting from $i = 1$, we can then sum $|\theta_i|$ until (6) reaches some predefined threshold. At that point, we simply choose the remaining θ_i as the reduced basis. We remark that (6) and (7) can be used for many different BF expansion models. The requirement for (6) is simply that we can evaluate $\int_{\Omega} |\phi_j(x)|^2 dx$, whereas (7) applies generally as long as $\int_{\Omega} |\phi_j(x)|^2 dx \approx \int_{\Omega} |\phi_i(x)|^2 dx, \forall i, j$.

2) *Probabilistic solution*: In the probabilistic case, where $\theta \sim \mathcal{N}(m, S)$, we substitute \mathcal{L} for $\mathbb{E}_p[\mathcal{L}]$. The delta method can then be used to approximate the expectation as $\mathbb{E}_p[\mathcal{L}(\theta)] \approx \mathcal{L}(\mathbb{E}_p[\theta]) = \mathcal{L}(m)$. Thus, the problem turns into selecting the BFs with the largest absolute mean. Obviously, this does not consider the quality of the approximate predictive variance. We nevertheless pursue this strategy as an approximate solution to the corresponding probabilistic BF selection problem. Empirically, as shown in Section III, this still works well. In this probabilistic setting, the BF selection process is even more beneficial than in the deterministic case, as computing the predictive variance is usually an $\mathcal{O}(L^2)$ operation.

3) *Discussion*: The considered problem is closely related to offline knowledge distillation, a type of model compression, that typically attempts to find a more compact representation of a large neural network model [23], [24]. However, knowledge distillation typically requires access to a dataset such that the reduced model can be trained to mimic the original model on particular data. In contrast, we avoid using data at the cost of a suboptimally reduced model, as we essentially throw away information from BFs we do not select. As such,

the method can be seen as a type of pruning, also common in neural network model compression [25], [26]. Pruning was pioneered in [27], where weights of a neural network were removed using Hessian information, again requiring access to a training dataset. Here, the information loss can theoretically be remedied by projecting the whole function $f(\cdot)$ onto the reduced bases. However, as this is computationally intensive ($\mathcal{O}(L^3)$), it is not pursued further here.

B. Hilbert-space GP dual selection

Here, we consider a special case of Section II-A applied to a commonly used sparse GP approximation.

1) *GP approximation*: A standard GP is a collection of random variables, any finite number of which have a joint Gaussian distribution [28]. Formally, we denote a GP by $f \sim \mathcal{GP}(0, \kappa(x, x'))$, where $\kappa(x, x')$ is the kernel, representing the covariance between the pair of inputs x and x' . The standard conjugate GP assumes independent observations of the latent process values, i.e., $y \sim \mathcal{N}(f, \sigma^2 I)$. The predictive distribution of the GP is then given by $\mathcal{N}(f^*; \mu^*, V^*)$, where

$$\mu^* \triangleq \mathbb{E}[f^* | \mathbf{y}] = K_{*f}(K_{ff} + \sigma^2 I)^{-1} \mathbf{y} \quad (8a)$$

$$V^* \triangleq \text{var}[f^* | \mathbf{y}] = K_{*f}(K_{ff} + \sigma^2 I)^{-1} K_{f*}, \quad (8b)$$

where $[K_{ab}]_{ij} = \kappa([X^a]_i, [X^b]_j)$ and X^* and X^f are the collection of test and training inputs, respectively.

Due to the inversion of $(K_{ff} + \sigma^2 I)$, a standard GP has a computational complexity of $\mathcal{O}(N^3)$, where N is the number of data points. Approaches to remedy this have been studied extensively, see, e.g., [29]–[35]. Here, we consider one of these approaches, namely the *Hilbert-space Gaussian Process* (HGP) [35], as it has seen widespread use [5], [36], [37]. In the HGP, the kernel matrix K_{ff} is approximated by an approximate eigenvalue decomposition, i.e., $K_{ff} \approx \Phi \Lambda \Phi^\top$, where Φ correspond to the eigenvectors and Λ the corresponding eigenvalues; see [35] for details.

The predictive distribution of the HGP is then given by [35]

$$\mu^* = \Phi_*^\top \Sigma \Phi^\top \mathbf{y}, \quad V^* = \sigma^2 \Phi_*^\top \Sigma \Phi_*,$$

where $\Sigma = (\Phi^\top \Phi + \sigma^2 \Lambda^{-1})^{-1}$. This can be recognized as a BF expansion model with posterior over the weights given by

$$p(\theta | \mathbf{y}) = \mathcal{N}(\Sigma \Phi^\top \mathbf{y}, \sigma^2 \Sigma) \triangleq \mathcal{N}(m, S).$$

Hence, again, we can find a reduced approximate model by choosing BFs according to their absolute mean through (7).

2) *Dual parametrization*: It is sometimes beneficial to parametrize the HGP using the ‘‘dual’’ parametrization

$$\alpha = \Phi^\top \mathbf{y}, \quad B = \Phi^\top \Phi,$$

as these are the only factors in the posterior predictive distribution that depend on the data. Further, both α and B have an additive structure, e.g.,

$$[B]_{ij} = \sum_{n=1}^N \phi_i(x_n) \phi_j(x_n), \quad (9)$$

meaning that the inclusion of a new data point is extremely simple. However, the dual parametrization means that the BF

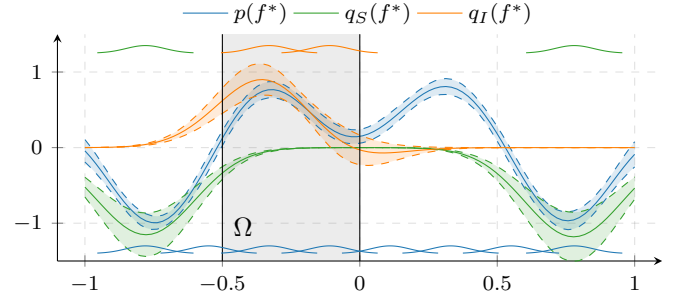


Fig. 1: Posterior predictive distributions in an RBF model. The full model is given by $p(f^*)$. The posteriors $q_I(f^*)$ and $q_S(f^*)$ use reduced models with bases selected by (6) and (7), respectively. The BFs for the original model are depicted at the bottom of the plot and the chosen BFs for the reduced models are depicted at the top of the plot. Note that the proportions of the BFs are exaggerated. The subdomain Ω of interest is highlighted in gray.

selection process is not as simple, since we would first need to convert the posterior from the (α, B) parametrization to the (m, S) , which is an $\mathcal{O}(L^3)$ operation, defeating the whole purpose to begin with. To see this, consider (6) and write

$$\mathbb{E} \left[\sum_{j \notin \mathcal{J}} |\theta_j|^2 \right] \approx \sum_{j \notin \mathcal{J}} |m_j|^2 = \sum_{j \notin \mathcal{J}} \left| \sum_{i=1}^n [\Sigma]_{ji} \alpha_i \right|^2 \leq \sum_{j \notin \mathcal{J}} \sum_{i=1}^n |[\Sigma]_{ji} \alpha_i|^2. \quad (10)$$

This obviously requires Σ , which costs $\mathcal{O}(L^3)$ to compute. One may also consider using, e.g., a truncated singular value decomposition of B to select important components, but this also bears with it a computational complexity of $\mathcal{O}(L^3)$. Instead, we propose to resolve this by only considering the diagonal elements of B such that

$$[\Sigma]_{ii} \approx \frac{1}{[B]_{ii} + \sigma^2 / \lambda_i}, \quad [\Sigma]_{ij} \approx 0, \quad i \neq j$$

such that computing Σ becomes an $\mathcal{O}(L)$ operation. Under the interpretation of B as the precision matrix of a BF expansion, this corresponds to neglecting any mutual information between BFs. With this approximation, the bound becomes

$$\sum_{j \notin \mathcal{J}} \sum_{i=1}^n |[\Sigma]_{ji} \alpha_i|^2 = \sum_{j \notin \mathcal{J}} |[\Sigma]_{jj} \alpha_j|^2. \quad (11)$$

Note that we only use the diagonal of B for BF selection, *not* for the final predictive distribution, where we use the full B . This results in a suboptimal selection strategy but we still expect it to perform well. This is motivated by the fact that, in the HGP, the BFs are orthonormal and the inner product between ϕ_i and ϕ_j is given by

$$\langle \phi_i, \phi_j \rangle = \int_{\Omega} \phi_i(x) \phi_j(x) dx = \delta_{ij}.$$

Hence, $[B]_{ij}$, see (9), can be viewed as an unweighted stochastic estimate of $\langle \phi_i, \phi_j \rangle$ based on the data. As long as the data distribution sufficiently covers the domain of the BFs, then $[B]_{ij} \rightarrow 0$, as $N \rightarrow \infty$, if $i \neq j$.

III. NUMERICAL EXAMPLES

Our first numerical example considers selection in the case of RBFs, as we, in this case, can intuitively determine which BFs are relevant, confirming the validity of our approach. Our second example considers a random function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ drawn from a GP prior. It is mainly used to illustrate the computational benefits of the BF selection approach in models with thousands of parameters. Lastly, a third example, where the method is applied to a real dataset for magnetic field mapping, is provided in the additional material.

A. Radial basis functions

We consider an RBF expansion where the RBFs are given by $\phi_i(x) = \exp(-\|x - c_i\|_2^2/l^2)$, where l and c_i are parameters of BF i . We generate synthetic data from a function on the domain $[-1, 1]$ and place $L = 10$ BFs equidistantly spaced in the domain. The posterior over the BF weights (m, S) is found through linear regression. We then select $|\mathcal{J}| = 2$ BFs for prediction in the subdomain $\Omega = [-0.5, 0]$ using both (6) and (7). The intuitive BF choice is BFs for which c_i are close, in some sense, to Ω , since the BFs are exponentially decaying away from c_i ; see [10], [11].

The results are presented in Fig. 1. Clearly, the BF selection using (6) finds that the intuitively correct closest BFs are the most important for prediction in Ω . The selection using (7), on the other hand, identifies the BFs on the edge of the domain as the most important. This is natural, considering that, in this case, $\int_{\Omega} |\phi_j(x)|^2 dx \not\approx \int_{\Omega} |\phi_i(x)|^2 dx, \forall i, j$, violating the assumption of (7). Nevertheless, this confirms the validity of the BF selection approach.

B. Random function

We consider a function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ drawn from a GP prior with zero mean and a squared-exponential kernel given by $\kappa(x, x') = \sigma_f^2 \exp(-\|x - x'\|_2^2/2l^2)$, with lengthscale $l = 0.1$ and variance $\sigma_f^2 = 0.05$. The $N = 1000$ training inputs are uniformly drawn from $[-1, 1]^3$ and i.i.d. Gaussian noise with variance $\sigma^2 = 0.01$ is added to the outputs. The baseline model is an HGP defined on $[-2, 2]^3$. The BFs are thus given by [35]

$$\phi_i(\mathbf{x}) = \prod_{d=1}^3 \frac{1}{\sqrt{2}} \sin\left(\frac{\pi i_d([\mathbf{x}]_d + 2)}{4}\right),$$

where i is a multi-index and i_d indexes each individual dimension. The number of BFs along each dimension is varied in $L_d \in [2, 20]$, corresponding to a total number of BFs between $L = 8$ and $L = 8000$.

To evaluate the quality of the approximate predictions, we consider the *Negative Log Predictive Density* (NLPD), RMSE, and *Kullback-Leibler* (KL) divergence of the reduced model's and full model's predictive distributions on a dense grid of $N_t = 3375$ test points. The NLPD and KL are computed toward a standard GP model with the same kernel. Lastly, we benchmark the time necessary to compute the predictions. For ease of presentation, the four metrics for the reduced and full model are then converted into "relative" metrics, by, e.g.,

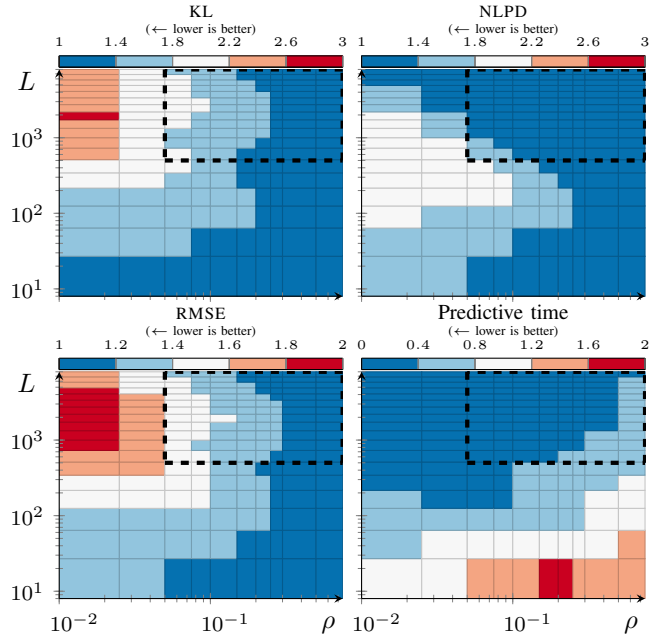


Fig. 2: Performance of the reduced model relative to the full (original) model. ρ is the fraction of BFs chosen for prediction, i.e., $\rho = 0.1$ chooses the 10% most important BFs. L is the total number of basis functions available for the full model. The metrics are all relative, where blue is better and red is worse.

$\frac{\text{NLPD}_{\text{app}}}{\text{NLPD}_{\text{full}}}$. All the metrics are plotted against the number of BFs L and fraction ρ of retained BFs in Fig. 2.

The results are intuitively sound, as the approximation error grows as more BFs are removed, i.e., for low values of ρ . For low values of L , the method does not yield any speed-ups. This is natural as the cost of the BF selection process itself is then greater than the speed-up achieved by the reduction. This is particularly apparent in the bottom right part of the relative predictive time plot, where the BF selection process actually makes the predictions *slower*. However, already at $L \approx 200$, the BF selection process starts speeding up the prediction process. Considering all of the metrics, the sweet spot is around $L > 500$ and $\rho > 0.05$; see the black dashed area in Fig. 2, where there is a low approximation error and low relative predictive time. For $\rho \approx 0.3$, the KL, NLPD, and RMSE are all close to the full model, while the predictive time is greatly reduced as long as $L > 200$. Note that the optimal value of ρ is problem-dependent. However, as mentioned in Section II, by selecting a proper threshold for the bounds (6) and (7), they may be used to automatically select the number of BFs to keep. We leave this for future exploration.

IV. CONCLUSIONS

A method for a posteriori selecting relevant BF in an already learned BF expansion model has been proposed. The method enables model predictions to be calculated efficiently over subsets of the model domain with minor accuracy reduction. The method can be applied to a variety of BF expansion models, such as RBF expansions, sparse GP approximations, etc., and enable these to tackle larger problems. An important application area for the proposed method is in multi-agent magnetic-field localization, where local and reduced-size GP-based magnetic-field maps must be distributed to the agents.

REFERENCES

- [1] G. Pillonetto, L. Schenato, and D. Varagnolo, "Distributed Multi-Agent Gaussian Regression via Finite-Dimensional Approximations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 9, pp. 2098–2111, Sep. 2019.
- [2] W. Chen and L. Liu, "Long-Term Autonomous Ocean Monitoring with Streaming Samples," in *OCEANS 2019 MTS/IEEE SEATTLE*, Seattle, WA, USA, Oct. 2019, pp. 1–8.
- [3] I. Torroba, M. Cella, A. Terán, N. Rolleberg, and J. Folkesson, "Online Stochastic Variational Gaussian Process Mapping for Large-Scale Bathymetric SLAM in Real Time," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3150–3157, Jun. 2023.
- [4] K. Berntorp, "Bayesian Tire-Friction Learning by Gaussian-Process State-Space Models," in *18th European Control Conference (ECC)*, Naples, Italy, Jun. 2019, pp. 231–236.
- [5] M. Kok and A. Solin, "Scalable Magnetic Field SLAM in 3D Using Gaussian Process Maps," in *Proceedings of the 21st International Conference on Information Fusion (FUSION)*, Cambridge, United Kingdom, Jul. 2018, pp. 1353–1360.
- [6] M. Vetterli, "Wavelets, approximation, and compression," *IEEE Signal Processing Magazine*, vol. 18, no. 5, pp. 59–73, 2001.
- [7] G.-J. Jang, T.-W. Lee, and Y.-H. Oh, "Single-channel signal separation using time-domain basis functions," *IEEE Signal Processing Letters*, vol. 10, no. 6, pp. 168–171, 2003.
- [8] Y. Wang, E. Athanassoula, and S. Mao, "Basis function expansions for galactic dynamics: Spherical versus cylindrical coordinates," *A&A*, vol. 639, 2020.
- [9] M. S. Petersen, M. D. Weinberg, and N. Katz, "exp: N-body integration using basis function expansions," *Monthly Notices of the Royal Astronomical Society*, vol. 510, no. 4, pp. 6201–6217, 01 2022.
- [10] F. M. Viset, R. Helmons, and M. Kok, "Spatially scalable recursive estimation of Gaussian process terrain maps using local basis functions," 2023, arXiv:2210.09168.
- [11] A. Kullberg, I. Skog, and G. Hendeby, "Online Joint State Inference and Learning of Partially Unknown State-Space Models," *IEEE Transactions on Signal Processing*, vol. 69, pp. 4149–4161, 2021.
- [12] M. Heaton, A. Datta, A. Finley, R. Furrer, J. Guinness, R. Guhaniyogi, F. Gerber, R. Gramacy, D. Hammerling, M. Katzfuss, F. Lindgren, D. Nychka, F. Sun, and A. Zammit-Mangion, "A case study competition among methods for analyzing large spatial data," *Journal of Agricultural, Biological and Environmental Statistics*, vol. 24, pp. 398–425, 12 2018.
- [13] D. Jang, J. Yoo, C. Y. Son, D. Kim, and H. J. Kim, "Multi-Robot Active Sensing and Environmental Model Learning With Distributed Gaussian Process," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5905–5912, Oct. 2020.
- [14] F. Viset, R. Helmons, and M. Kok, "Distributed multi-agent magnetic field norm SLAM with Gaussian processes," in *Proceedings of the 26th International Conference on Information Fusion (FUSION)*, Charleston, SC, USA, Jun. 2023, pp. 1–8.
- [15] M. D. Buhmann, *Radial Basis Functions: Theory and Implementations*, 1st ed. Cambridge University Press, Jul. 2003.
- [16] —, "Radial functions on compact support," *Proceedings of the Edinburgh Mathematical Society*, vol. 41, no. 1, pp. 33–46, Feb. 1998.
- [17] Z. Zhong and J. Hua, "Kernel-based adaptive sampling for image reconstruction and meshing," *Computer Aided Geometric Design*, vol. 43, pp. 68–81, 2016.
- [18] S. Cammarasana and G. Patanè, "Kernel-based sampling of arbitrary signals," *Computer-Aided Design*, vol. 141, p. 103103, 2021.
- [19] G. H. John, R. Kohavi, and K. Pfleger, "Irrelevant features and the subset selection problem," in *Machine Learning Proceedings 1994*. San Francisco (CA): Morgan Kaufmann, 1994, pp. 121–129.
- [20] J. Zhu, C. Wen, J. Zhu, H. Zhang, and X. Wang, "A polynomial algorithm for best-subset selection problem," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 117, no. 52, pp. 33 117–33 123, Dec. 2020.
- [21] B. Venkatesh and J. Anuradha, "A review of feature selection and its methods," *Cybernetics and Information Technologies*, vol. 19, no. 1, p. 3–26, Mar. 2019.
- [22] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.
- [23] J. Gou, B. Yu, S. J. Maybank, and D. Tao, "Knowledge distillation: A survey," *International Journal of Computer Vision*, vol. 129, no. 6, p. 1789–1819, Mar. 2021.
- [24] C. Bucilua, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '06. New York, NY, USA: Association for Computing Machinery, 2006, p. 535–541.
- [25] Z. Li, H. Li, and L. Meng, "Model compression for deep neural networks: A survey," *Computers*, vol. 12, no. 3, 2023.
- [26] M. B. Hossain, N. Gong, and M. Shaban, "Computational complexity reduction techniques for deep neural networks: A survey," in *IEEE International Conference on Artificial Intelligence, Blockchain, and Internet of Things (AIBThings)*, 2023, pp. 1–6.
- [27] Y. LeCun, J. Denker, and S. Solla, "Optimal brain damage," in *Advances in Neural Information Processing Systems*, vol. 2. Denver, CO, USA: Morgan-Kaufmann, 1989.
- [28] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2005.
- [29] L. Csató, "Gaussian Processes – Iterative Sparse Approximations," Ph.D. dissertation, Aston University, 2002.
- [30] J. Quiñero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate Gaussian process regression," *Journal of Machine Learning Research*, vol. 6, no. 65, pp. 1939–1959, 2005.
- [31] M. Titsias, "Variational learning of inducing variables in sparse Gaussian processes," in *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, ser. Proceedings of Machine Learning Research, vol. 5. PMLR, 16–18 Apr 2009, pp. 567–574.
- [32] A. Wilson and H. Nickisch, "Kernel interpolation for scalable structured Gaussian processes (KISS-GP)," in *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, ser. Proceedings of Machine Learning Research, vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 1775–1784.
- [33] J. Hensman, N. Fusi, and N. D. Lawrence, "Gaussian processes for big data," in *Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence (UAI)*, ser. UAI'13. Arlington, Virginia, USA: AUAI Press, 2013, p. 282–290.
- [34] P. Izmailov, A. Novikov, and D. Kropotov, "Scalable Gaussian processes with billions of inducing inputs via tensor train decomposition," in *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, ser. Proceedings of Machine Learning Research, vol. 84. PMLR, 09–11 Apr 2018, pp. 726–735.
- [35] A. Solin and S. Särkkä, "Hilbert space methods for reduced-rank Gaussian process regression," *Statistics and Computing*, vol. 30, no. 2, pp. 419–446, Mar. 2020.
- [36] K. Berntorp, "Online Bayesian inference and learning of Gaussian-process state-space models," *Automatica*, vol. 129, p. 109613, Jul. 2021.
- [37] A. Svensson, A. Solin, S. Särkkä, and T. Schön, "Computationally Efficient Bayesian Learning of Gaussian Process State Space Models," in *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, ser. Proceedings of Machine Learning Research. Cadiz, Spain: PMLR, May 2016, pp. 213–221.