



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Standing-wave-decomposition Gaussian process

Chi-Ken Lu, Scott Cheng-Hsin Yang, and Patrick Shafto

Phys. Rev. E **98**, 032303 — Published 11 September 2018

DOI: [10.1103/PhysRevE.98.032303](https://doi.org/10.1103/PhysRevE.98.032303)

Standing Wave Decomposition Gaussian Process

Chi-Ken Lu, Scott Cheng-Hsin Yang, and Patrick Shafto

Department of Mathematics and Computer Science, Rutgers University, Newark, New Jersey 07302

(Dated: August 21, 2018)

We propose a Standing Wave Decomposition (SWD) approximation to Gaussian Process regression (GP). GP involves a costly matrix inversion operation, which limits applicability to large data analysis. For an input space that can be approximated by a grid and when correlations among data are short-ranged, the kernel matrix inversion can be replaced by analytic diagonalization using the SWD. We show that this approach applies to uni- and multi-dimensional input data, extends to include longer-range correlations, and the grid can be in a latent space and used as inducing points. Through simulations, we show that our approximate method applied to the squared exponential kernel outperforms existing methods in predictive accuracy per unit time in the regime where data are plentiful. Our SWD-GP is recommended for regression analyses where there is a relatively large amount of data and/or there are constraints on computation time.

I. INTRODUCTION

Gaussian Process (GP) regression [1] is powerful as well as elegant because of its ability to transform correlation into prediction. This nonparametric and Bayesian approach needs not specify a predetermined function that the data points fit to, and only assumes that any subset of data follows a joint Gaussian distribution characterized by the mean and covariance functions. GP is equivalent to a class of neural networks which have infinite number of hidden units [2, 3]. As a flexible model for highly complicated functions and physical processes, GP is a familiar tool for physicists in many context, such as predicting melting temperature [4], inter-atomic potential [5], fractional Brownian motion [6], DNA replication kinetics [7, 8], gravitational waveform [9], etc. Furthermore, in the machine learning community, GP can be used for classification [1, 10], data dimension reduction, and signal reconstruction [11].

However, the scalability of GP regression is limited by its $O(N^3)$ training time for N training data points. Inverting the kernel matrix is the critical step which essentially transforms correlation into prediction. The task of speeding up the GP training algorithm has been taken up by computer scientists who use inducing points and Nyström method to approximate the kernel matrix [12], and, recently, by physicists [13, 14] who employ quantum algorithms capable of solving system of linear equations [15]. In this paper, we propose an approximating GP based on sparsifying the kernel matrix of a latent grid in the input space, which allows us to analytically diagonalize the kernel matrix by employing standing wave decomposition. Observing that correlation among nearest neighbors is most relevant to prediction, we approximate the full matrix with its tridiagonal matrix.¹ This tridiagonal matrix has analytic eigenvalues and eigenvectors resembling standing waves, a fact that has been extensively exploited in the studies of electronic structure of

solids [16]. If the training data is on a grid, like time-series data, then the approximated kernel matrix can be immediately used in the usual GP prediction algorithm. When the given data is not on a grid in the input space, a latent grid is treated as inducing points and the function values are obtained by Bayesian projection from off-grid data. Our approach works best for the squared exponential (SE) kernel matrix, a very popular kernel for many machine learning tasks, because the off-diagonal matrix elements decay rapidly. For high dimensional grid data, the SE kernel matrix has a Kronecker product form, which makes the present approach applicable as well.

One constraint of the approximated GP is that the ratio of length scale for SE kernel to the grid spacing, ℓ/Δ , has an upper bound because the kernel matrix must be semi-definite positive, which is guaranteed if the smallest eigenvalue is positive. The tridiagonal approximation works well when the neglected off-diagonal matrix elements have little effect on those small eigenvalues and the corresponding eigenvectors. An improvement to the tridiagonal approximation and the upper bound for ℓ/Δ is the pentadiagonal approximation, which has the next nearest correlation included in the kernel matrix. The pentadiagonal matrix, however, does not have simple form of eigenvalues and eigenvectors. Thus, using the eigenvector of tridiagonal kernel matrix, we reconstruct the kernel matrix so that the next-nearest neighbor correlation terms are included and the eigenvectors are still exact with eigenvalues modified. In the end, a simulation with synthetic data is carried out with comparison with other GP approximations to highlight the condition under which the proposed method is at its best.

The rest of the paper is structured as follows. In Sec. II, an overview of GP regression along with an analytic example in terms of eigen basis decomposition is given. In Sec. III, we introduce the approximate tridiagonal kernel matrix for a one-dimensional grid input and derive the standing wave eigen basis. The extension to two-dimensional grid input using the Kronecker product is provided in Sec. IV. To account for longer-ranged correlations, we demonstrate in Sec. V the reconstruction of a nearly pentadiagonal kernel matrix which allows the

¹ In a tridiagonal matrix, the only non-zero values are on the main diagonal, superdiagonal, and subdiagonal.

standing wave decomposition Gaussian process (SWD-GP) to be applied with longer length scale. Based on the above techniques, in Sec. VI our SWD-GP is then combined with a latent grid (LG) to model data that do not lie on a grid. Our LGSWD-GP method is compared with popular existing methods, and the simulated accuracy and run time are reported in Sec. VII. Sec. VIII is devoted to the investigation of the time complexity of our approximated methods and a comparison with some popular algorithms. In the end, the related works and a brief discussion are given in Sec. IX and Sec. X, respectively.

II. OVERVIEW OF GP

Here we give an overview of GP by considering two unobserved values y_1 and y_2 at two distinct locations x_1 and x_2 , respectively. To make it a tractable problem, we may assume the values follow a joint normal distribution, namely $p(y_1, y_2) = \mathcal{N}(0, \Sigma)$ with zero mean and covariance matrix

$$\Sigma = \sigma^2 \begin{pmatrix} 1 & c \\ c & 1 \end{pmatrix}. \quad (1)$$

The parameter c is a function of distance $|x_1 - x_2|$, which quantifies the correlation between y_1 and y_2 . Here, $|c| \leq 1$ so that Σ is positive semidefinite. It is then interesting to see how the joint distribution adjusts itself after one of the variable, say, y_2 has been observed. The way the correlation is transformed into prediction is seen by first decomposing the inverse of kernel matrix using the eigenvalues and eigenvectors, which yields $\Sigma^{-1} = \frac{1}{\lambda_+} \mathbf{v}_+ \mathbf{v}_+^T + \frac{1}{\lambda_-} \mathbf{v}_- \mathbf{v}_-^T$. The eigenvalues are $\lambda_{\pm} = \sigma^2(1 \pm c)$ and the corresponding eigenvectors read

$$\mathbf{v}_{\pm} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ \pm 1 \end{pmatrix}. \quad (2)$$

After some algebra, we find that the joint distribution is factorized into,

$$p(y_1, y_2) \propto \exp \left[-\frac{(y_1 - cy_2)^2}{2(1-c^2)\sigma^2} \right] \exp \left(-\frac{y_2^2}{2\sigma^2} \right). \quad (3)$$

Thus, the adjusted distribution for y_1 becomes conditional on the observed y_2 . It is easy to verify that $p(y_1|y_2) = p(y_1, y_2)/p(y_2)$ is still Gaussian, and the updated mean and variance $\mu_{1|2} = cy_2$ and $\sigma_{1|2}^2 = (1-c^2)\sigma^2$, respectively.

Generalizing the above procedure to a multivariate joint distribution over the set of observed and unobserved variables $\{y_1, y_2, \dots, y_N, y_*\}$ with a general kernel matrix K constitutes the GP regression, which is to seek the underlying latent function $y = f(\mathbf{x})$ mapped from the input points $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N, \mathbf{x}_*\}$. Because of the Gaussianity, the correlation among the data leads to the conditional distribution over the unobserved variable,

$p(y_*|y_{1:N})$, which is another Gaussian $\mathcal{N}(\mu_*, \sigma_*^2)$ with updated mean,

$$\mu_* = \mathbf{k}_*^T K^{-1} \mathbf{y}, \quad (4)$$

and variance,

$$\sigma_*^2 = \sigma^2 - \mathbf{k}_*^T K^{-1} \mathbf{k}_*, \quad (5)$$

where K denotes the kernel matrix associated with training data $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$ and \mathbf{k}_* denotes the kernel matrix between training and test points (\mathbf{x}_*, y_*) . Corresponding to the above example, $K = \sigma^2$, $k_* = c\sigma^2$, and $y = y_2$.

In the following, we shall focus on the kernel matrix taken from the squared exponential kernel function $[K]_{lm} = k(\mathbf{x}_l, \mathbf{x}_m)$,

$$k(\mathbf{x}, \mathbf{x}') = \sigma^2 \prod_{i=1}^d \exp \left[-\frac{(x_i - x'_i)^2}{2\ell^2} \right], \quad (6)$$

where the product form appears due to the separable sum of squared distance along each dimension. The hyperparameters include the variance σ^2 and length scale ℓ .

III. STANDING WAVE EIGEN BASIS

It is observed that the kernel matrix K given by (6) has a simplified form if all the input points $\{\mathbf{x}\}_{i=1}^M$ are on an one-dimensional grid since the matrix elements $k(\mathbf{x}_i, \mathbf{x}_j)$ only depend on $|i - j|$, as in a Toeplitz matrix [17]. Here we make further simplification by retaining only the most relevant terms, namely the matrix elements $[K]_{ij}$ with $|i - j| \leq 1$, which is quite legitimate when the grid unit length Δ is larger than length scale. Notice that the rest of matrix elements decrease exponentially as $\alpha^{(i-j)^2}$ with $\alpha = \exp(-\frac{\Delta^2}{2\ell^2})$. The eigenvalue equations $K\mathbf{v} = \lambda\mathbf{v}$ in terms of the coefficients of eigenvector $\mathbf{v} = [c_1, c_2, \dots, c_M]$ read

$$\alpha c_2 + (1 - \tilde{\lambda})c_1 = 0, \quad (7a)$$

$$\alpha c_{i+1} + (1 - \tilde{\lambda})c_i + \alpha c_{i-1} = 0, \quad (7b)$$

$$(1 - \tilde{\lambda})c_N + \alpha c_{N-1} = 0, \quad (7c)$$

where the index $i = 2 \dots M-1$ applies in (7b). The eigenvalues are rescaled as $\tilde{\lambda} = \lambda/\sigma^2$. We observe that (7) is similar with the energy Hamiltonian for an electron hopping on semi-infinite one-dimensional tight-binding lattice [16]. Defining the z-transform associated with an auxiliary $\{c_i\}_{i=1}^{\infty}$ as $F(z) = \sum_{i=1}^{\infty} c_i z^{1-i}$, we find from (7a) and (7b) that $F(z) = c_1 [1 - (\alpha + (1 - \tilde{\lambda})z)/(\alpha z^2 + (1 - \tilde{\lambda})z + \alpha)]$. Moreover, the sequence $\{c_i\}$ can be reproduced from the contour integral [16],

$$c_i = c_1 \left[\delta_{i1} - \frac{1}{2\pi i} \oint F(z) z^{i-2} dz \right], \quad (8)$$

followed by employing the residue theorem. We note that the pair of poles of $F(z)$ are $z_{1,2} = e^{\pm i\theta}$ with which the eigenvalues $\tilde{\lambda} = 1 + 2\alpha \cos \theta$ can be consistent. Thus, the coefficients are shown to be,

$$c_i = \frac{\sin i\theta}{\sin \theta} c_1, \quad (9)$$

for all nonnegative i . The final stage is to take care of (7c), which is not necessarily consistent with the solution in (9) with arbitrary θ and $\tilde{\lambda}$. The only way to make (9) a valid solution is to set the auxiliary term $c_{M+1} = 0$, which results in

$$\theta_k = k \frac{\pi}{M+1}, \quad k = 1, 2, 3, \dots, M, \quad (10)$$

and the corresponding eigenvalues,

$$\lambda_k = \sigma^2(1 + 2\alpha \cos \theta_k). \quad (11)$$

Consequently, we may conclude the orthonormal standing wave basis,

$$[v_k]_j = \frac{\sin j\theta_k}{\sqrt{\frac{M+1}{2}}}. \quad (12)$$

Fig. 1 demonstrates four pairs of symmetric eigenvectors $\{v_k, v_{M-k+1}\}$ of an 20×20 kernel matrix K . As seen from the identity in (10), it is interesting to remark that the arguments θ_k and θ_{M-k+1} always sum up to π , which suggests the symmetric coefficients $[v_k]_i = (-1)^{i+1}[v_{M-k+1}]_i$. Such symmetry is also manifested in Fig. 1 where, for instance, in (a) the star symbols representing v_{20} alternate its sign whereas the dot symbols associating with v_1 do not. Finally, the inverted kernel matrix can be simply decomposed as,

$$K^{-1} = \sum_{k=1}^M \frac{1}{\lambda_k} v_k v_k^T. \quad (13)$$

Here, we note that the same sparse approximation which results in the tridiagonal K must also be applied to the kernel matrix \mathbf{k}_* so that the product $\mathbf{k}_* K^{-1}$ shall result in a row vector which is one at the j th element and zero otherwise when the test point coincides with the j th training point. The expression for K^{-1} in (13) suggests that the prediction mean of SWD-GP shall pass through the training points if all the eigenvalues are positive. Next, we employ the inverted kernel matrix (13) in the GP learning algorithm (4) and (5), and apply this SWD-GP to the on-grid data, $\{x_i\}_{i=1}^{10}$, evenly spaced in $[0, 1]$, and $\{y_i = f(x_i)\}$ with $f(x) = x \cos(2\pi x) \sin[24\pi(x + 0.03)]$. Figure 2(a) and (b) represent the results using $\ell/\Delta = 0.27$ and 0.54 , respectively. The prediction mean associated with smaller ℓ/Δ passes through all the training points, but is not able to capture the longer-distance features. Away from the training points, the variance σ_* is close to the value of σ used

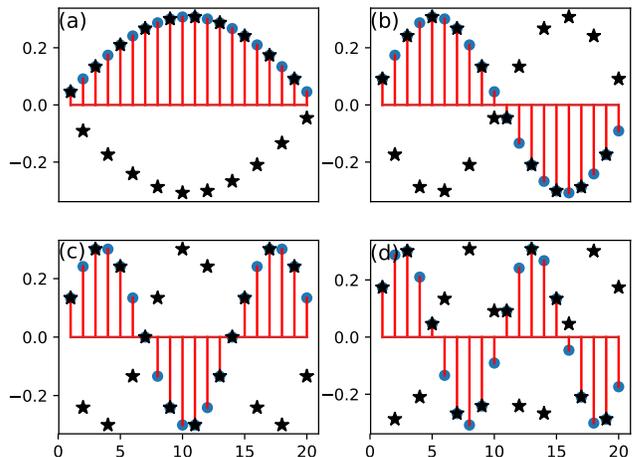


FIG. 1. Four pairs of eigenvectors v_k and v_{M-k} are shown for the tridiagonal kernel matrix with $M = 20$. Circle symbols denote the components of $[v_k]$ while the star symbols represent its symmetric partner $[v_{M-k+1}]$. Panels (a) to (d) correspond to $k = 1, 2, 3,$ and 4 , respectively.

in (11), which is expected from the squared exponential kernel. On the other hand, the prediction mean and variance are optimal for the used length scale $\ell/\Delta = 0.54$ because further increased ℓ/Δ shall lead to zero or negative eigenvalues in (11).

IV. MULTI-DIMENSIONAL GRID

For higher dimensional input $\mathbf{x} \in \mathbb{R}^d$, the squared exponential kernel function as an instance of a product kernel allows a compact and efficient representation of kernel matrix, namely [18]

$$K = K^{(1)} \otimes K^{(2)} \otimes \dots \otimes K^{(d)} \quad (14)$$

where the matrix element associated with the n th dimension is given by $K_{\mathbf{x}, \mathbf{x}'}^{(n)} = k(x_n, x'_n)$. Fig. 3 shows the regression results for data generated from the function $f(x, y) = \sin(4\pi x) \cos(4\pi y)$ using the 10×10 (middle) and 20×20 (right) grid, respectively, and $\ell = 0.03$.

V. KERNEL MATRIX RECONSTRUCTION

Thus far it has been seen how the eigen decomposition of kernel matrix using the standing wave basis in (12) leads to the GP prediction given one-dimensional (Figure 2) and two-dimensional grid inputs (Figure 3). Although the size of grid and the ensuing length Δ can be varied arbitrarily, the length scale always has a constraint. More precisely, demanding the positive semi-definiteness of kernel K , which is guaranteed if the min-

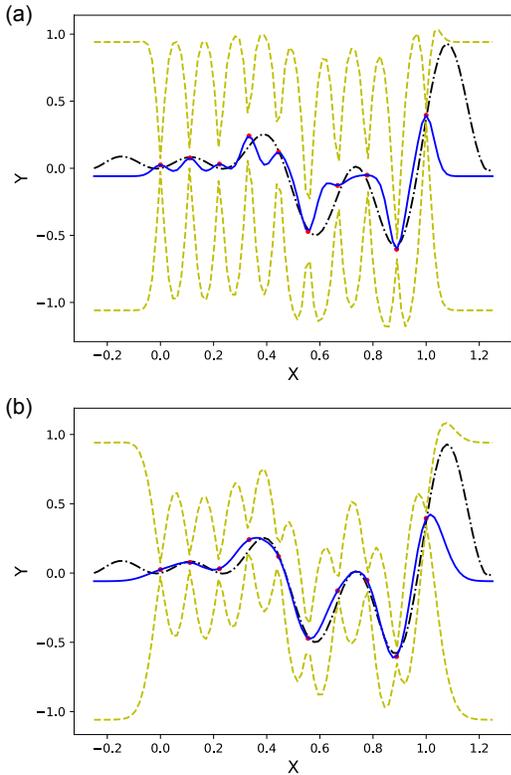


FIG. 2. Regression using on-grid data (red dots) generated by the function $f(x) = x \cos(2\pi x) \sin[4\pi(x + 0.03)]$ (black dashed line). Blue line and yellow dashed lines represent the prediction mean μ_* of SWD-GP and plus/minus one variance σ_* around the mean. The 10 grid points are uniformly placed in $[0, 1]$. The length scale are $\ell = 0.03$ and 0.06 , respectively, for (a) and (b).

imum eigenvalue λ in (11) is positive, leads to the constraint $\ell/\Delta < 1/\sqrt{2 \ln 2} \approx 0.85$. However, an exact GP in which every correlation term is retained is free of such constraint on ℓ . In the following, we show how to incrementally relax the constraint by reconstructing the kernel matrix.

As a motivation, let us examine Fig. 4(a) where the results from exact GP (green star) and tridiagonal SWD-GP (blue solid line) for data generated by $y = \sin(12\pi x) \cos(2\pi x)$ (black dashed line) using $\ell/\Delta = 0.57$ for a grid of 20 points in $[0, 1]$ are shown. First, we note that the GP and the present SWD-GP have identical results, suggesting that the SWD-GP is indeed a good approximation. Secondly, both results here are unable to capture the curvature feature near the extreme points, e.g. at $x = 1/8$ there is a subtle discrepancy against the true function. Besides the fact that there are not sufficiently many data points near these points, increasing the length scale may work, as suggested by the exact GP result (green star) in Fig. 4(b) where a longer $\ell/\Delta = 0.76$ is used. However, our simulation shows that directly applying the tridiagonal SWD-GP with $\ell/\Delta = 0.76$ does not

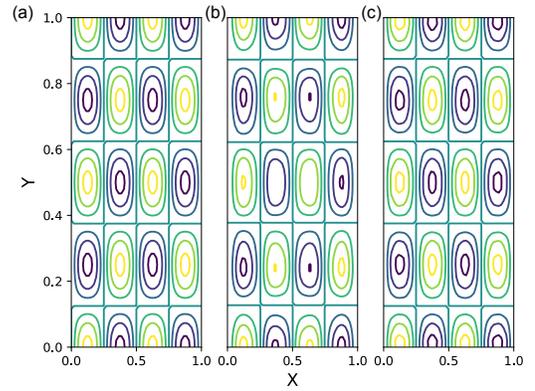


FIG. 3. Two-dimensional XY space grid used for regression with data generated by function $f(x, y) = \sin(4\pi x) \cos(4\pi y)$. The function values are represented by the color code. (a) The true function. (b) The regression result using 10×10 grid. (c) 20×20 grid.

generate as a smooth prediction mean as the case with $\ell/\Delta = 0.57$ does. This is because as the length scale approaches the upper bound, the smallest eigenvalue approaches zero and creates numerical instabilities.

In order to extend the applicable range for ℓ/Δ , we are interested in restoring the next nearest neighbor (NNN) correlation terms proportional to α^4 in K , i.e. the matrix elements K_{ij} with $|i - j| = 2$. Inspection of (7) suggests that the eigenvalue acquires a new contribution so that the new eigenvalue $\lambda^{(1)}$ reads

$$\lambda^{(1)}/\sigma^2 = 1 + 2\alpha \cos \theta + 2\alpha^4 \cos 2\theta \quad (15)$$

while the eigenvectors are left intact. In fact, neither the eigenvalues nor the eigenvectors are exact because the correlation among the grid points is not symmetric near the boundary.

Instead of solving for the true eigenvalues and true eigenvectors of the kernel matrix K containing the NNN correlation, we reconstruct the kernel matrix which is still diagonalized by the eigen basis in (12) and has eigenvalues as in (15). We apply the following identity to reconstruct \tilde{K} ,

$$\sum_{k=1}^M \cos(q\theta_k) [\mathbf{v}_k \mathbf{v}_k^T]_{ij} = \frac{1}{2(M+1)} [a(i-j+q) + a(i-j-q) - a(i+j+q) - a(i+j-q)] \quad (16)$$

with the coefficient defined as,

$$a(p) \equiv \sum_{k=1}^M \cos \frac{kp\pi}{M+1} = \frac{-1 - (-1)^p}{2}, \quad (17)$$

for $p \in \mathbb{Z}_{2N+2} \setminus \{0\}$ and M otherwise. It is easy to see that the constant ($q = 0$) together with the first cosine term ($q = 1$) in (15) gives rise to the original tridiagonal matrix. As for the last term in (15) corresponding

to $q = 2$, the RHS of (16) associated with $i - j = \pm 2$ equal $\frac{1}{2}$, which leads to the desired NNN correlation term. However, special care is needed for the cases for $i + j - q = 0$ and $i + j + q = 2N + 2$, which, respectively, correspond to $(i, j) = (1, 1)$ and $(i, j) = (M, M)$ with $q = 2$. It is easily verified that the reconstructed kernel matrix elements corresponding to the end points in grid are $\tilde{K}_{11} = \tilde{K}_{MM} = K_{11} - \sigma^2\alpha^4$.

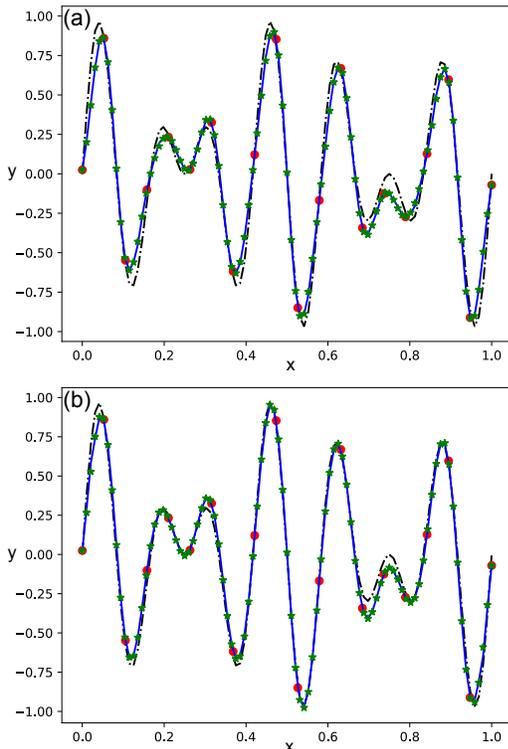


FIG. 4. Regression results from the SWD-GP (blue lines) and exact GP (green star symbols) in comparison with the true function (black-dashed lines). Red dots denote the training data points generated by $f(x) = \cos(2\pi x)\sin(12\pi x)$ (black dashed line) on the grid of 20 points in $[0, 1]$. (a) Fit is obtained by applying tridiagonal SWD-GP and use length scale $\ell = 0.03$. (b) Fit is obtained from the reconstructed pentadiagonal kernel matrix with $\ell = 0.04$.

In summary, the reconstructed pentadiagonal kernel matrix \tilde{K} , which contains the NNN correlation terms, reads,

$$\tilde{K}_{ij} = K_{ij} + \sigma^2\alpha^4 [\delta_{|i-j|,2} - \delta_{i,j}(\delta_{i,1} + \delta_{i,M})]. \quad (18)$$

As shown in Fig. 4(b), we employ the pentadiagonal \tilde{K} using $\ell/\Delta = 0.76$ to repeat the regression task, and we are able to reproduce the same result as the exact GP with same length scale. Moreover, including NNN correlation also improves the accuracy near the local extreme points.

VI. LATENT GRID SWD-GP

A. Bayesian Data Projection

When applying to data (\mathbf{X}, \mathbf{y}) where the input points \mathbf{X} are randomly selected, we shall construct a latent grid \mathbf{X}_g so that the previous SWD-GP approach can be employed. The missing components are the corresponding function values \mathbf{g} evaluated at these grid points. Now we assume that \mathbf{g} is a set of random variable and we find the most probable values $\bar{\mathbf{g}}$ in the Bayesian approach. We first calculate the posterior distribution associated with \mathbf{g} ,

$$p(\mathbf{g}|\mathbf{X}_g, \mathbf{X}, \mathbf{y}) \propto \left(\prod_{i=1}^N p(y_i|\mathbf{x}_i, \mathbf{X}_g, \mathbf{g}) \right) p(\mathbf{g}|\mathbf{X}_g) \quad (19)$$

where we have assumed independence among the given data. The likelihood function follows the usual GP predictive distribution by assuming \mathbf{g} is known,

$$p(y|\mathbf{x}, \mathbf{X}, \mathbf{g}) = \mathcal{N}(K_{\mathbf{xg}}K_{\mathbf{gg}}^{-1}\mathbf{g}, \sigma_{\mathbf{x}}^2), \quad (20)$$

and the variance at \mathbf{x} is denoted by

$$\sigma_{\mathbf{x}}^2 = K_{\mathbf{xx}} - K_{\mathbf{xg}}K_{\mathbf{gg}}^{-1}K_{\mathbf{gx}} + \sigma_N^2 \quad (21)$$

with σ_N^2 being the variance of the observation noise. The prior has its zero mean distribution with the kernel matrix $K_{\mathbf{gg}}$,

$$p(\mathbf{g}|\mathbf{X}_g) = \mathcal{N}(0, K_{\mathbf{gg}}). \quad (22)$$

With some tedious calculations, the posterior distribution can be shown to be,

$$p(\mathbf{g}|\mathbf{X}_g, \mathbf{X}, \mathbf{y}) = \mathcal{N}(\bar{\mathbf{g}}, A^{-1}), \quad (23)$$

where the updated variance is encoded in

$$A^{-1} = K_{\mathbf{gg}}Q^{-1}K_{\mathbf{gg}}, \quad (24)$$

with the matrix

$$Q = K_{\mathbf{gg}} + K_{\mathbf{gx}}\Lambda^{-1}K_{\mathbf{xg}}. \quad (25)$$

The diagonal matrix $\Lambda = \text{diag}(\sigma_{\mathbf{x}_1}^2, \sigma_{\mathbf{x}_2}^2, \dots, \sigma_{\mathbf{x}_N}^2)$. Furthermore, the mean vector $\bar{\mathbf{g}}$ represents the most probable projected function values,

$$\bar{\mathbf{g}} = K_{\mathbf{gg}}Q^{-1}K_{\mathbf{gx}}\Lambda^{-1}\mathbf{y}. \quad (26)$$

Now, by marginalizing the latent function values \mathbf{g} , we may obtain the predictive distribution for the function value y_* associated with the test point \mathbf{x}_* ,

$$p(y_*|\mathbf{x}_*) = \mathcal{N}(K_{*\mathbf{g}}K_{\mathbf{gg}}^{-1}\bar{\mathbf{g}}, \sigma_*^2 + K_{*\mathbf{g}}Q^{-1}K_{\mathbf{g}*}), \quad (27)$$

which is consistent with the results presented in [19, 20].

B. Perturbation and Matrix Inversion

Unlike the grid data to which the inverted kernel matrix $K_{\mathbf{g}\mathbf{g}}^{-1}$ is readily applicable, here we have to invert the matrix Q in (25) in order to find the mean of predictive distribution in (27). To approach this problem, we shall eventually write Q in terms of the SW eigen basis of $K_{\mathbf{g}\mathbf{g}}$. Hence, the second term in (25) is treated as perturbation so that the eigenvalues in (11) as well as eigenvectors in (12) are modified. Here, to make the calculation more efficient and accessible, we regroup the M eigenvectors $\{v_i\}$ with $i = 1, 2, \dots, M$ into pairs $\{v_i, v_{M+1-i}\}$, and for ease of notation we use v and \bar{v} to represent the two eigenvectors in a given pair. Projecting Q onto the subspace spanned by v and \bar{v} , we arrive at the following eigenvalue equation for each pair, and the associated eigenvalue χ and eigenvector u can be obtained by solving,

$$\begin{pmatrix} \lambda + \epsilon & \delta \\ \delta & \bar{\lambda} + \bar{\epsilon} \end{pmatrix} u = \chi u. \quad (28)$$

The value ϵ is the first-order correction to unperturbed eigenvalue λ due to the term $v^T Q v$, while δ is the cross term from $\bar{v}^T Q v$. It can be shown that the updated eigenvalues and eigenvectors are given by,

$$\chi_{\pm} = \beta_{\pm} \pm \sqrt{\beta_{\pm}^2 + \delta^2}, \quad (29)$$

and

$$\begin{pmatrix} u_+ \\ u_- \end{pmatrix} = \begin{pmatrix} \cos \phi & \sin \phi \\ -\sin \phi & \cos \phi \end{pmatrix} \begin{pmatrix} v \\ \bar{v} \end{pmatrix} \quad (30)$$

with the rotation angle ϕ specified by

$$\tan \phi = \frac{\sqrt{\beta_-^2 + \delta^2} + \beta_-}{\sqrt{\beta_-^2 + \delta^2} - \beta_-}. \quad (31)$$

The parameter $\beta_{\pm} = [(\lambda + \epsilon) \pm (\bar{\lambda} + \bar{\epsilon})]/2$. Consequently, the matrix Q has the following decomposition,

$$Q^{-1} = \sum_{i=1}^M \frac{u_i u_i^T}{\chi_i}. \quad (32)$$

Next, we apply the the above approach to the data generated by the function $y = x \cos[8\pi(x + 0.15)] \cos(2\pi x)$ with input x randomly selected from $[0,1]$. Figure 5(a)–(b) and (c)–(d) show the results for the training number $|X| = 200$ and 1000, respectively. Panels (a) and (c) in Fig. 5 are obtained in first-order perturbation, or by setting $\delta = 0$ in (28). With a larger number of training points [Fig. 5(c) versus (a)], the most probable projected function values $\bar{\mathbf{g}}$ are closer to the true function values on the grid. Panels (b) and (d) correspond to a fuller second-order perturbation treatment by considering the eigenvector correction due to nonzero δ in (28). Incremental improvements, and in particular near the end points, can be seen by comparing (a) versus (b), or (c) versus (d).

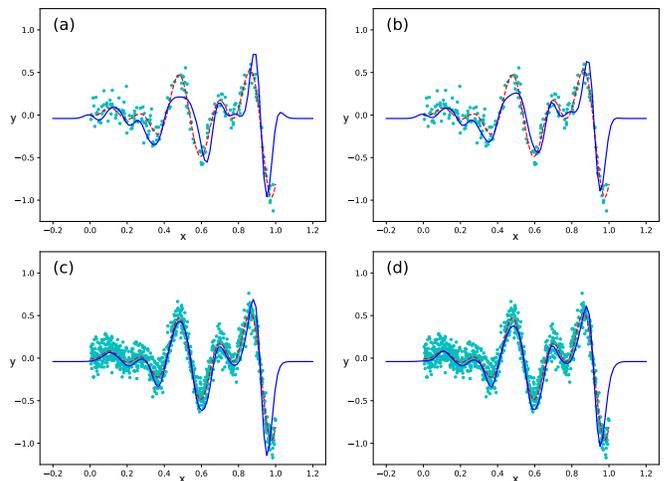


FIG. 5. Illustration of latent grid SWD-GP 1st order $\delta = 0$ in (28) and 2nd order ($\delta > 0$) perturbation using data generated from function $f(x) = x \cos[8\pi(x + 0.15)] \cos(2\pi x)$ (red dashed line) plus noise. A latent grid of 20 points and $\ell = 0.03$ are used, and the prediction mean is denoted by the blue solid line. (a) First-order perturbation with 200 data points. (b) Second-order perturbation with same parameters as (a). (c) Same as (a) but with 1000 data points. (d) Same as (b) but the number of data points increased to 1000.

The above results suggest that the inference through Eq. (26) using the approximated matrix inversion is effective for larger $|X|$ and noise parameter σ_N/σ of order 10^{-1} . To get insight into why it works, we shall show that the second matrix in Eq. (25) can be approximated as a linear combination of identity and super/sub diagonal matrices,

$$[K_{\mathbf{g}\mathbf{x}} \Lambda^{-1} K_{\mathbf{x}\mathbf{g}}]_{ij} = a \delta_{ij} + b \delta_{|i-j|=1} + R_{ij}, \quad (33)$$

where a and b stand for the averages of diagonal and super/sub diagonal matrix elements, respectively. The last term R denotes the residual matrix. Therefore, the first two terms in Eq. (33) commute with $K_{\mathbf{g}\mathbf{g}}$ while R is *small* if the variations among the diagonal and super/sub diagonal matrix elements are weak. We demonstrate in Fig. 6(a)–(c) that larger values of σ_N/σ result in weaker variation among the matrix elements. This is further confirmed in panels (d) and (e) where the ratio of standard deviation over mean along diagonal [Fig. 6(d)] and super/sub diagonal [Fig. 6(e)] lines are sufficiently small for $\sigma_N/\sigma > 0.1$.

VII. SIMULATIONS

We compare the method presented here with three existing, popular methods—exact GP, fully independent training conditional GP (FITC-GP) [19, 20], and Kernel Interpolation for Scalable Structured GP (KISS-GP) [21]. The main quantities of interest are data efficiency—the amount of data needed to reach a certain accuracy—

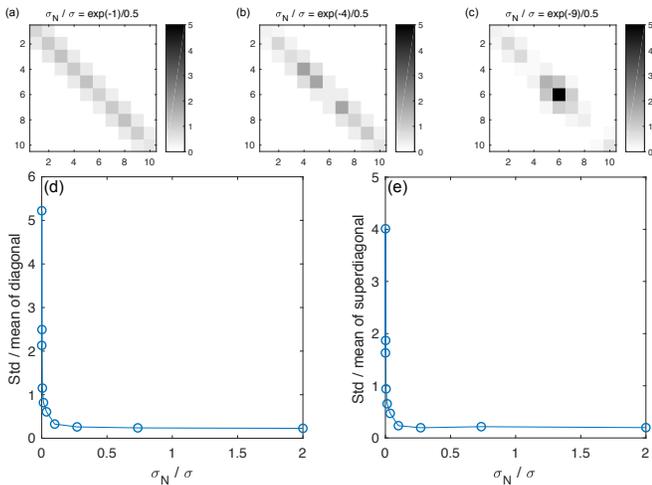


FIG. 6. Variation of matrix elements of $K_{\mathbf{g}\mathbf{x}}\Lambda^{-1}K_{\mathbf{x}\mathbf{g}}$, the second matrix in Eq. (25). Upper panel: results for grid size $|\mathbf{X}_{\mathbf{g}}| = 10$ and number of training points $|\mathbf{X}| = 200$ with different noise parameters σ_N/σ . From (a) to (c), decreasing σ_N/σ 's leads to stronger variations among the matrix elements along the diagonal, superdiagonal, and subdiagonal. Lower panel: Using $|\mathbf{X}_{\mathbf{g}}| = 50$ and $|\mathbf{X}| = 1000$, the ratio of standard deviation over the mean for matrix elements along the diagonal (d) and sub/super diagonal (e) against σ_N/σ is obtained. A critical value $\sigma_N/\sigma = 0.1$ is suggested for valid perturbation.

and run time. We apply all methods to a toy regression problem, where the x-values are drawn uniformly random from $[0, 1]$, and the corresponding y-values are given by $y = \sin(\frac{5\pi}{x+0.1}) + \text{Normal}(0, 0.2^2)$. Accuracy is measured by standardized mean-square error (SMSE). Run time is measured on a 2015 model of MacBook with 1.1 GHz dualcore Intel Core M and 8GB 1600MHz LPDDR3. All methods are implemented in MATLAB 2017b. We largely followed an example² in the documentation of the GPML toolbox for using the existing methods. For the implementation of our latent grid standing wave decomposition GP (LGSWD-GP), we used only first-order correction on the eigenvalues (i.e., setting $\delta = 0$ in (28)), since Fig. 5 suggests that the improvement from adding second-order correction on the eigenvectors is relatively small compared to increasing the number of training points. Our code is available at: TBA.

For all methods, the squared-exponential kernel is used. The kernel's overall variance, σ^2 , is fixed at 0.5^2 , and the variance of likelihood is set to the true value, 0.2^2 . Exact GP and KISS-GP are allowed to optimize for the kernel's length scale, and FITC-GP is allowed to optimize for the location of the inducing points in addition to the kernel's length scale. These hyper-parameter optimization are limited to 20 iterations. Note that for our LGSWD-GP, the length scale can be determined without

invoking hyper-parameter optimization (see Sec. V). The number of inducing points are fixed at 300 throughout for FITC-GP, KISS-GP, and LGSWD-GP. The inducing grid in KISS-GP and LGSWD-GP are evenly spaced in $[0, 1]$. The test points are 500 evenly spaced points in $[0, 1]$. Example fits are shown in Figure 7.

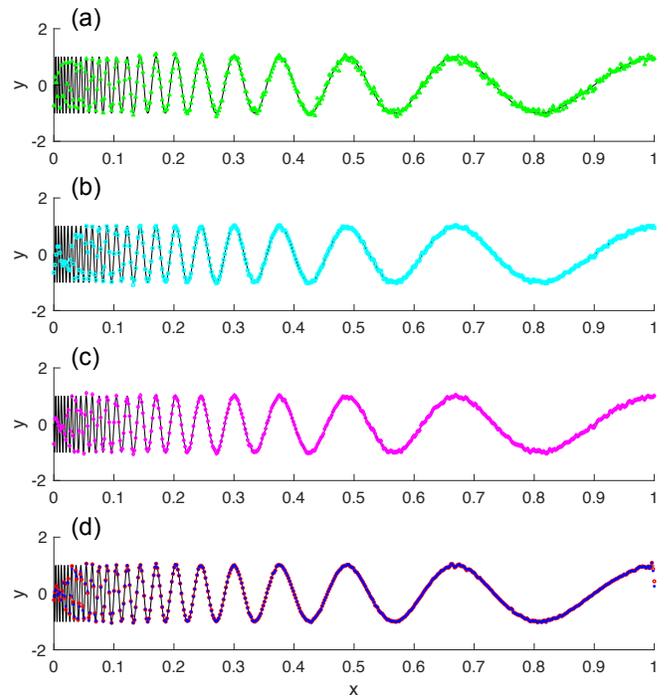


FIG. 7. Example fits of (a) exact GP with $N = 3000$, (b) FITC-GP with $N = 10^4$, (c) KISS-GP with $N = 10^4$, and (d) our LGSWD-GP with $N = 5 \times 10^5$, where N is the number of training data points. The different N used in each method corresponds to the largest N used for that method in Fig. 8. Black line represents the true underlying function; markers represent the mean of the predictive distribution at the test points. For (d), red circles and blue crosses correspond to tridiagonal and pentadiagonal reconstruction, respectively.

Figure 8 shows the accuracy and run time as a function of number of training points for the different methods. From the error plot, we see that our method has a lower data efficiency compared to existing methods; that is, compared to those methods, the LGSWD-GP needs more data to achieve the same accuracy. Nevertheless, the runtime plot shows that the poorer data-efficiency is offset by the radical increase in computational efficiency. The LGSWD-GP is roughly two orders of magnitude faster than FITC-GP and KISS-GP and is capable of reaching or even surpassing these methods in accuracy in the same amount of time. A trade-off between speed and accuracy is expected. What is interesting is that these seemingly bold tridiagonal and pentadiagonal kernel approximations provides enough data efficiency to make the trade off favorable.

Next, we analyze the accuracy by breaking the test region into five equi-sized regions, i.e.,

² www.gaussianprocess.org/gpml/code/matlab/doc/demoGrid1d.m

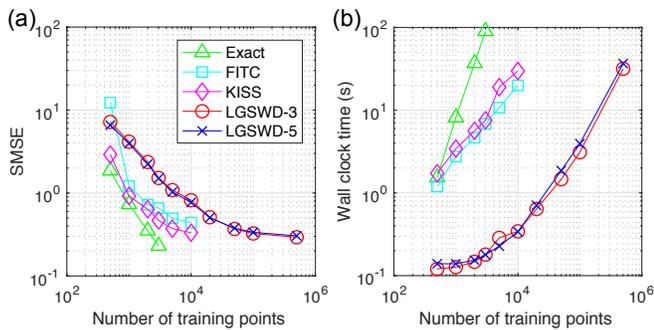


FIG. 8. SMSE error (a) and run time (b) as a function of number of training points for exact GP, FITC-GP, KISS-GP, and LGSWD-GP. LGSWD-3 and LGSWD-5 correspond to the tridiagonal and pentadiagonal approximations, respectively.

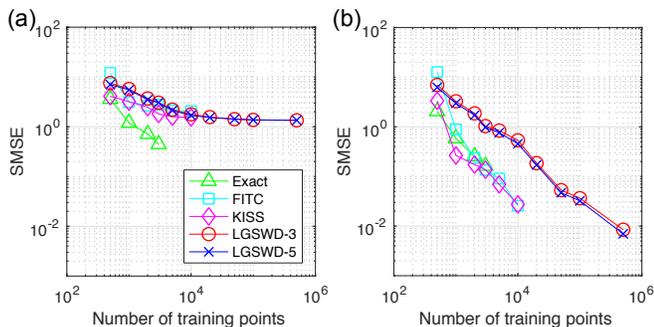


FIG. 9. Accuracy by regions. SMSE is calculated for the 100 test points in regions (a) $x = [0, 0.2]$ and (b) $x = (0.2, 0.4)$. In theory, all the training points affect the prediction at every test point; thus, the x-axis is the same as that in Fig. 8. But effectively, only roughly one fifth of the training points (those in the test region) contributed to the fit. The error plots for the remaining regions $(0.4, 0.6)$, $(0.6, 0.8)$, and $(0.8, 1]$ all look similar to (b).

$[0, 0.2), (0.2, 0.4), \dots, (0.8, 1]$. Note that the effective length scale of the underlying function is different in each region (see Fig. 7). Let the effective length scale in region r be ξ_r and the number of training points in ξ_r be $n_r = N \frac{\xi_r}{L}$. The ratio Δ/ξ_r and n_r are useful in characterizing the accuracy behavior of GP. First, all methods do relatively poorly when n_r is small (Fig. 9(a)) and relatively well when n_r is large enough (Fig. 9(b)), which is just a restatement of classical sampling theory. Second, methods that depend on inducing points reach a plateau when Δ/ξ_r is relatively large (Fig. 9(a)) and thus are limited in data efficiency. Third, when Δ/ξ_r becomes sufficiently small (Fig. 9(b)), this limitation no longer exists. Based on these observations, we recommend using exact GP when both Δ/ξ_r and n_r are sufficiently large, FITC-GP or KISS-GP when Δ/ξ_r is sufficiently small and n_r is sufficiently large, and LGSWD-GP when Δ/ξ_r is sufficient small and n_r is large. In other words, among the methods shown, our method is the most accurate and time-efficient one in

the regime where the latent grid is relatively fine and the number of training points is large.

VIII. TIME COMPLEXITY ANALYSIS

For N training data points on the grid, (4) and (13) gives the predictive mean for SWD-GP:

$$\mu_* = \sum_{i=1}^N \frac{1}{\lambda_i} (\mathbf{k}_{*v_i}^T v_i^T \mathbf{y}). \quad (34)$$

Since the second inner product has time complexity $O(N)$, the sum has a total time complexity of $O(N^2)$. The space complexity of SWD-GP is $O(N)$. For M inducing points on a grid and N training data points, from Sec. VI one can show that the predictive mean for LGSWD-GP is

$$\mu_* = \sum_{i=1}^M \frac{1}{\chi_i} (k_{*g} u_i) (u_i^T K_{g\mathbf{x}} \Lambda^{-1} \mathbf{y}), \quad (35)$$

where $\chi_i = \lambda_i + u_i^T K_{g\mathbf{x}} \Lambda^{-1} K_{\mathbf{x}g} u_i$. The time complexity for computing χ_i is $O(pN)$, where p is the number of bands used in the approximation ($p = 3$ for tridiagonal and $p = 5$ for pentadiagonal); thus, the sum gives a total time complexity of $O(MpN)$. The same time complexity is found for computing Λ^{-1} and Q^{-1} . The space complexity of LGSWD-GP is $O(pN)$ with the bottleneck being at representing $K_{g\mathbf{x}}$.

Next, for clarity of comparison, we show how our method differs from other methods in terms of the other methods' bottleneck operation: Exact GP requires composing the kernel matrix and inverting it. The bottleneck operation is the inverse, which has time complexity $O(N^3)$ for N training points. If the data lie on a grid, the SWD method can compose the inverted kernel analytically via (13), although in practice, as shown by (34), we never need to explicitly represent the inverted kernel matrix for computing the predictive mean.

The bottleneck operation of FITC-GP is the matrix multiplication $K_{g\mathbf{x}} \Lambda^{-1} K_{\mathbf{x}g}$ in (25), which has time complexity $O(M^2N)$ for M inducing points and N training points. If the inducing points lie on a grid, that multiplication in LGSWD-GP with p -diagonal kernel approximation has time complexity $O(MpN)$, because each row in $K_{\mathbf{x}g}$ has at most p non-zero elements. Note that one can choose not to explicitly compute this particular matrix product as it is absorbed into the χ_i 's in (35). In addition, whenever a matrix inverse operation is required, such as in (21), LGSWD-GP benefits from being able to construct the inverse directly.

KISS-GP, which also uses inducing points on a grid, replaces all the operations in computing the predictive mean in (27) by solving a linear system $\tilde{K}_{*x} \tilde{K}_{xx}^{-1} \mathbf{y}$ through linear conjugate gradients, where $\tilde{K}_{ab} \approx K_{ag} K_{gg}^{-1} K_{gb}$ [21]. The optimization has time complexity $O(SqN)$ because the product $\tilde{K}_{xx} \mathbf{y}$ can be computed

in $O(qN)$ given the sparse approximation made by KISS-GP and because the optimization can be fixed to S steps. The contribution of q and k in KISS-GP correspond to the contribution of p and M in LGSWD-GP, respectively: p and q stem from sparse approximations, while M and a good choice for S depend on the number of contributing eigenvalues.

Lastly, we mention that there are a couple of ways to further speed up SWD-GP and LGSWD-GP. First, as a result of the sine functions appeared in $\{v_k\}_{k=1}^N$, we may rewrite the mean and variance in (4) and (5) in terms of Discrete Fourier Transformed vector for the training data $\{y_i\}_{i=1}^N$,

$$[\text{FFT}_\theta(\mathbf{y})]_j = \sum_{k=1}^N e^{ik\theta_j} y_k, \quad (36)$$

where $\theta_j = j\pi/(N+1)$. As such, the mean function becomes

$$\mu_* = \frac{2}{N+1} \sum_{i=1}^N \frac{\mathcal{I}m[\text{FFT}_\theta(k_*)]_i \mathcal{I}m[\text{FFT}_\theta(\mathbf{y})]_i}{\lambda_i}, \quad (37)$$

and the predicting variance is

$$\sigma_*^2 = \sigma^2 - \frac{2}{N+1} \sum_{i=1}^N \frac{\{\mathcal{I}m[\text{FFT}_\theta(k_*)]_i\}^2}{\lambda_i}. \quad (38)$$

For SWD-GP, this reduces the time complexity from $O(N^2)$ to $O(N \log N)$, analogous to how Toeplitz structure reduces the time complexity of matrix multiplication. For LGSWD-GP, although this will not help reduce the bottleneck time complexity, it can help reduce many of the operations involving matrix products with the eigenvectors. Second, the sum in the eigen-decomposition of (12) and (32) can be computed independently and is directly amenable to parallelization.

IX. RELATED WORK

Previous works have also taken advantage of structured data to increase the computational efficiency of GP. For input data that lie on a grid, [22, 23] sidestep the matrix inversion by solving $K^{-1}\mathbf{y}$ with gradient-based methods. Importantly, the optimization is fast because Toeplitz / circulant structure can be exploited for fast matrix multiplication. These fast multiplications take on forms similar to (36). For input data that do not lie on a grid, [21] (KISS-GP) introduces a latent grid, finds a sparse, approximate representation of $K_{\mathbf{x}\mathbf{g}}$ in terms of $K_{\mathbf{g}\mathbf{g}}$, and makes inference using the same kind of optimization problem in [22, 23]. The sparsification in KISS-GP and that in our LGSWD-GP share the same flavor in ignoring the diminishing long-range correlations among data points. Recently, [24] further extended this optimization approach to efficiently handle high dimensional

data by applying the Nyström approximation for eigen-decomposition and by exploiting properties of Kronecker and Khatri-Rao products for fast matrix multiplication. Overall, our method and these methods all take advantage of the grid structure to form fast matrix multiplication; however, our method differs in that it overcomes the matrix inversion bottleneck via analytic diagonalization rather than optimization.

Variational approaches are another popular method for GP inference. Variational methods turn the inference of the predictive mean and variance into an optimization problem [25]. Variational GP often provides more accurate prediction than FITC [26] and has been made progressively faster through a series of development—from stochastic variational inference [27], to distributed variational inference [28], to asynchronous distributed variational inference [29]—to handle billions of input data. Similar to our work and [24], [30] uses grid inducing points with stochastic variational inference, which allows added efficiencies in computation via the Kronecker and Khatri-Rao products.

X. DISCUSSION

In this paper we present an analytic approach to invert a class of kernels for GP regression with grid inputs through standing wave decomposition (SWD). This class of kernels is roughly Toeplitz³ and thus approximates all translational invariant kernels with grid inputs. Here we show results for tridiagonal Toeplitz kernels, which can be thought of as approximating a square exponential kernel by keeping only the nearest neighbor correlation among data points, and a pentadiagonal kernel, which includes the next nearest neighbor correlation. We demonstrate that the approach can be extended to approximate product kernels in higher dimensions. Lastly, we apply the approach to a latent grid of inducing points so as to handle training data that do not lie on a grid.

From the perspective of signal processing, one may regard the grid as a collection of input points at which the function is sampled. Thus, the grid representation of data and function is sufficient for prediction and reconstruction as long as the underlying functions have a finite Nyquist rate.

Our SWD approach is unique in that the inversion of the (approximated) kernel matrix is done analytically. This analytic inverse, along with the sparse matrices induced by keeping only the most important correlations among data points, makes this approach computationally very efficient (Sec. VIII). Furthermore, the optimal length scale hyperparameter in the SWD approach can be determined without the need to invoke optimization

³ The tridiagonal kernel is exactly Toeplitz. For pentadiagonal kernel, (18) shows which elements are modified and by how much. The procedure is readily extendable to more diagonals.

(Sec. V). When applied to training data on a grid, we see that the predictions of SWD-GP are very similar to those of exact GP if the two use the same length scale (Fig. 4), although vastly faster. For data that do not lie on a grid, our results show that LGSWD-GP can be both faster and more accurate than existing approximation GP methods when the number of training points is large (Figs. 8 and 9). We expect that this advantage will be magnified in situations where a highly precise latent representation is not crucial, such as GP classification. The analytical form of the eigen-decomposition also allows straight-forward parallelization of our approach, which we leave to future work. Other future directions

include extending the approach to kernel approximations with longer-range correlation, to other types kernel such as periodic kernels, and to latent grid in high dimensions.

ACKNOWLEDGMENTS

This material is based on research sponsored by the Air Force research Laboratory and DARPA under agreement number FA8750-17-2-0146 to P.S. and S.Y. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.

-
- [1] C. E. Rasmussen and C. K. I. Williams. *Gaussian Process for Machine Learning*. MIT press, Cambridge, MA, 2006.
 - [2] Radford M Neal. Monte carlo implementation of gaussian process models for bayesian regression and classification. *arXiv preprint physics/9701026*, 1997.
 - [3] Christopher KI Williams. Computing with infinite networks. In *Advances in neural information processing systems*, pages 295–301, 1997.
 - [4] Atsuto Seko, Tomoya Maekawa, Koji Tsuda, and Isao Tanaka. Machine learning with systematic density-functional theory calculations: Application to melting temperatures of single-and binary-component solids. *Physical Review B*, 89(5):054303, 2014.
 - [5] Albert P Bartók, Mike C Payne, Risi Kondor, and Gábor Csányi. Gaussian approximation potentials: The accuracy of quantum mechanics, without the electrons. *Physical review letters*, 104(13):136403, 2010.
 - [6] Tridib Sadhu, Mathieu Delorme, and Kay Jörg Wiese. Generalized arcsine laws for fractional brownian motion. *Physical review letters*, 120(4):040603, 2018.
 - [7] A Baker, B Audit, SC-H Yang, J Bechhoefer, A Arneodo, et al. Inferring where and when replication initiates from genome-wide replication timing data. *Physical review letters*, 108(26):268101, 2012.
 - [8] A Baker and J Bechhoefer. Inferring the spatiotemporal dna replication program from noisy data. *Physical Review E*, 89(3):032703, 2014.
 - [9] Zoheyr Doctor, Ben Farr, Daniel E Holz, and Michael Pürrer. Statistical gravitational waveform models: What to simulate next? *Physical Review D*, 96(12):123011, 2017.
 - [10] Pablo Morales-Álvarez, Adrián Pérez-Suay, Rafael Molina, and Gustau Camps-Valls. Remote sensing image classification with large-scale gaussian processes. *IEEE Transactions on Geoscience and Remote Sensing*, 2017.
 - [11] M. Titsias and N. Lawrence. Bayesian gaussian process latent variable model. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 844–851, 2010.
 - [12] Christopher KI Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. In *Advances in neural information processing systems*, pages 682–688, 2001.
 - [13] Siddhartha Das, George Siopsis, and Christian Weedbrook. Continuous-variable quantum gaussian process regression and quantum singular value decomposition of nonsparse low-rank matrices. *Physical Review A*, 97(2):022315, 2018.
 - [14] Zhikuan Zhao, Jack K Fitzsimons, and Joseph F Fitzsimons. Quantum assisted gaussian process regression. *arXiv preprint arXiv:1512.03929*, 2015.
 - [15] Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical review letters*, 103(15):150502, 2009.
 - [16] R. A. Marcus. Tight-binding approximation for semi-infinite solids. application of a transform method and of delta function normalization. *The Journal of Chemical Physics*, 98(7):5604–5611, 1993.
 - [17] Robert M Gray et al. Toeplitz and circulant matrices: A review. *Foundations and Trends® in Communications and Information Theory*, 2(3):155–239, 2006.
 - [18] Yunus Saatçi. *Scalable inference for structured Gaussian process models*. PhD thesis, University of Cambridge, 2012.
 - [19] E. Snelson and Z. Ghahramani. Sparse gaussian processes using pseudo-inputs. In *Advances in neural information processing systems*, pages 1257–1264, 2006.
 - [20] Joaquín Quiñero-Candela and Carl Edward Rasmussen. A unifying view of sparse approximate gaussian process regression. *Journal of Machine Learning Research*, 6(Dec):1939–1959, 2005.
 - [21] Andrew Wilson and Hannes Nickisch. Kernel interpolation for scalable structured gaussian processes (kiss-gp). In *International Conference on Machine Learning*, pages 1775–1784, 2015.
 - [22] John P Cunningham, Krishna V Shenoy, and Maneesh Sahani. Fast gaussian process methods for point process intensity estimation. In *Proceedings of the 25th international conference on Machine learning*, pages 192–199. ACM, 2008.
 - [23] Richard E Turner. *Statistical models for natural sounds*. PhD thesis, UCL (University College London), 2010.
 - [24] Trefor W Evans and Prasanth B Nair. Scalable gaussian processes with grid-structured eigenfunctions (gp-grief). 2017.
 - [25] Michalis Titsias. Variational learning of inducing variables in sparse gaussian processes. In *Artificial Intelligence and Statistics*, pages 567–574, 2009.
 - [26] Matthias Bauer, Mark van der Wilk, and Carl Edward Rasmussen. Understanding probabilistic sparse gaussian

- process approximations. In *Advances in neural information processing systems*, pages 1533–1541, 2016.
- [27] James Hensman, Nicolo Fusi, and Neil D Lawrence. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*, 2013.
- [28] Yarin Gal, Mark Van Der Wilk, and Carl Edward Rasmussen. Distributed variational inference in sparse gaussian process regression and latent variable models. In *Advances in Neural Information Processing Systems*, pages 3257–3265, 2014.
- [29] Hao Peng, Shandian Zhe, Xiao Zhang, and Yuan Qi. Asynchronous distributed variational gaussian process for regression. In *International Conference on Machine Learning*, pages 2788–2797, 2017.
- [30] Thomas Nickson, Tom Gunter, Chris Lloyd, Michael A Osborne, and Stephen Roberts. Blitzkriging: Kronecker-structured stochastic gaussian processes. *arXiv preprint arXiv:1510.07965*, 2015.