

SCALABLE RANDOM WAVELET FEATURES: EFFICIENT NON-STATIONARY KERNEL APPROXIMATION WITH CONVERGENCE GUARANTEES

Sawan Kumar

Department of Applied Mechanics
Indian Institute of Technology Delhi
New Delhi 110016, India
sawan.kumar@am.iitd.ac.in

Souvik Chakraborty

Department of Applied Mechanics
Yardi School of Artificial Intelligence (ScAI)
Indian Institute of Technology Delhi
New Delhi 110016, India
souvik@am.iitd.ac.in

ABSTRACT

Modeling non-stationary processes, where statistical properties vary across the input domain, is a critical challenge in machine learning; yet most scalable methods rely on a simplifying assumption of stationarity. This forces a difficult trade-off: use expressive but computationally demanding models like Deep Gaussian Processes, or scalable but limited methods like Random Fourier Features (RFF). We close this gap by introducing Random Wavelet Features (RWF), a framework that constructs scalable, non-stationary kernel approximations by sampling from wavelet families. By harnessing the inherent localization and multi-resolution structure of wavelets, RWF generates an explicit feature map that captures complex, input-dependent patterns. Our framework provides a principled way to generalize RFF to the non-stationary setting and comes with a comprehensive theoretical analysis, including positive definiteness, unbiasedness, and uniform convergence guarantees. We demonstrate empirically on a range of challenging synthetic and real-world datasets that RWF outperforms stationary random features and offers a compelling accuracy-efficiency trade-off against more complex models, unlocking scalable and expressive kernel methods for a broad class of real-world non-stationary problems.

1 INTRODUCTION

The ability to model complex, real-world phenomena is one of the central challenges in machine learning. Domains such as geospatial modeling, where terrain varies drastically across regions, or speech analysis, where signals exhibit bursts of volatility, are often characterized by pronounced *non-stationarity*, meaning their statistical properties change across the input space. Gaussian Processes (GPs) offer a principled framework for such problems, providing robust uncertainty estimates and flexible, non-parametric modeling (Williams & Rasmussen, 2006). Despite these advantages, exact GPs suffer from two major limitations: their expressivity is often constrained by the choice of kernel, and their computational cost scales cubically with the number of training points, rendering them impractical for modern large-scale applications (Liu et al., 2020).

Most of the current approaches force a trade-off between expressivity and efficiency. On one hand, methods like Random Fourier Features (RFF) achieve impressive scalability by approximating the kernel with a linear-in-data feature map (Cutajar et al., 2017; Avron et al., 2017; Rahimi & Recht, 2007). Yet, their dependence on Bochner’s theorem (Bochner, 2005) fundamentally restricts them to stationary kernels, which assume uniform behavior across the entire domain. Applying these models to non-stationary data leads to systematic mis-specification, resulting in compromised predictive accuracy and uncalibrated uncertainty estimates (Cheema & Rasmussen, 2024; Hensman et al., 2013; 2018). On the other hand, expressive models like Deep GPs (Salimbeni et al., 2019), spectral mixtures (Tompkins et al., 2020), and input-dependent kernels (Rudner et al., 2020) can capture non-stationarity, but they often reintroduce prohibitive computational costs, complex infer-

ence schemes, and challenges in optimization and hyperparameter tuning. The gap between scalable stationary models and complex non-stationary ones still remains.

In this work, we close this gap by introducing **Random Wavelet Features (RWF)**, a scalable and expressive framework for non-stationary kernel approximation. Instead of relying on globally supported sinusoidal bases like RFF and its variants, we construct random features from *wavelets* family of functions that are inherently localized in both space and frequency. By sampling wavelets at random scales and shifts, RWF generates an explicit feature map that can adapt to local data characteristics. This multi-resolution structure allows the model to capture sharp, localized events with fine-scale wavelets while simultaneously modeling smooth, long-range trends with coarse-scale wavelets. The result is a principled method that generalizes random features to the non-stationary setting while preserving the linear-time complexity that makes them elegant and efficient. Our main contributions are threefold. First, we provide a comprehensive theoretical analysis of RWF, including positive definiteness of the induced kernels, unbiasedness and variance bounds, and uniform convergence guarantees with explicit sample complexity. Second, we show that RWF achieves $\mathcal{O}(ND^2)$ training complexity, retaining the scalability of random feature methods while directly encoding non-stationarity through wavelet localization. Finally, we demonstrate empirically on synthetic, speech, and large-scale regression benchmarks that RWF consistently improves upon stationary random features and offers at least competitive accuracy–efficiency trade-off against more complex non-stationary models.

1.1 RELATED WORK

Scalable Kernel Approximations. The random features framework was pioneered by (Rahimi & Recht, 2007), showing shift-invariant kernels can be approximated using random Fourier features with linear-time computations. This framework has since been extended in several directions, including computationally efficient variants such as Fast kernel learning (Wilson et al., 2014), theoretical guarantees on approximation error (Sriperumbudur & Szabó, 2015; Avron et al., 2017; Rudi & Rosasco, 2017; Li et al., 2021), and structured sampling schemes (Choromanski et al., 2017). There are works that extend the random Fourier features beyond classical settings using variational approximations (Hensman et al., 2018), adaptive feature learning (Zhen et al., 2020; Shi et al., 2024), and even a connection to quantum machine learning (Landman et al., 2022). Recent progress extends spectral approximations to capture a wider spectrum of kernel families, thereby enhancing the expressivity of scalable feature maps (Langrené et al., 2024). While these methods achieve scalability, their reliance on stationary Fourier bases limits their ability to capture non-stationary (Paciorek & Schervish, 2003) or localized phenomena, which are crucial in many scientific domains.

Wavelet-motivated approximations. Wavelets and continuous wavelet frames (Mallat, 1999) provide localized multi-scale representations and have previously been used for kernel design through wavelet support vector machines and wavelet kernel learning (Zhang et al., 2004; Yger & Rakotomamonjy, 2011), where kernels are derived analytically, or wavelet transforms are used as preprocessing. More recently, Guo et al. (2024) proposed a Bayesian kernel model based on fixed wavelet bases for high-dimensional Bayesian linear regression. More closely related, Kereta et al. (2019) studies Monte Carlo discretizations of continuous wavelet frames for signal reconstruction. However, their construction relies on the eigendecomposition of the empirical kernel matrix, leading to cubic computational complexity in the number of samples. While these approaches illustrate the value of wavelets for capturing local structure, they rely on fixed or predefined wavelets dictionaries and do not provide theoretical guarantees such as unbiasedness or uniform convergence.

Hybrid and modern kernel learning. Several approaches have been developed to capture non-stationarity in GPs through spectral mixture kernels (Wilson & Adams, 2013; Remes et al., 2017) and deep Gaussian processes (Damianou & Lawrence, 2013; Salimbeni et al., 2019), though both remain costly for large datasets. Scalable variants include KISS-GP (Wilson & Nickisch, 2015), which exploits structured interpolation, and deep kernel learning (Wilson et al., 2016) combines neural feature extractors with GPs. More recent efforts include deep random features for spatiotemporal learning (Chen et al., 2024), graph-based random Fourier features (Zhang et al., 2025), and adaptive RKHS constructions (Shi et al., 2024). Despite these advances, existing methods often trade off scalability, expressivity, and interpretability. Our work is positioned at this intersection where we aim to design feature maps that inherit the scalability of random features while enabling flexible, non-stationary modeling.

2 PRELIMINARIES AND BACKGROUND

A brief review of Gaussian Process regression (GPR), sparse variational GPs, and random-feature GPs is provided to ground our wavelet construction.

2.1 GAUSSIAN PROCESSES

Given training inputs $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top \in \mathbb{R}^{N \times d}$ and targets $\mathbf{y} \in \mathbb{R}^N$, we consider a Gaussian process prior over a latent function f . The observations y_n are assumed to be noisy evaluations of this function at the corresponding inputs \mathbf{x}_n :

$$f \sim \mathcal{GP}(0, k), \quad y_n = f(\mathbf{x}_n) + \varepsilon_n, \quad \text{where } \varepsilon_n \sim \mathcal{N}(0, \sigma^2). \quad (2.1)$$

We define the covariance matrix $\mathbf{K}_{XX} \in \mathbb{R}^{N \times N}$ with $[\mathbf{K}_{XX}]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. The log marginal likelihood, used for training the hyperparameters of GP, is given by the following expression:

$$\log p(\mathbf{y} | \mathbf{X}) = -\frac{1}{2} \mathbf{y}^\top (\mathbf{K}_{XX} + \sigma^2 \mathbf{I}_N)^{-1} \mathbf{y} - \frac{1}{2} \log \det(\mathbf{K}_{XX} + \sigma^2 \mathbf{I}_N) - \frac{N}{2} \log(2\pi). \quad (2.2)$$

For a test input \mathbf{x}_* , let $\mathbf{k}_{*X} = [k(\mathbf{x}_*, \mathbf{x}_1), \dots, k(\mathbf{x}_*, \mathbf{x}_N)]$, $k_{**} = k(\mathbf{x}_*, \mathbf{x}_*)$, and

$$\boldsymbol{\alpha} = (\mathbf{K}_{XX} + \sigma^2 \mathbf{I}_N)^{-1} \mathbf{y}. \quad (2.3)$$

The predictive posterior moments for test input \mathbf{x}_* takes the following form:

$$\mu_*(\mathbf{x}_*) = \mathbf{k}_{*X} \boldsymbol{\alpha}, \quad (2.4a)$$

$$\sigma_*^2(\mathbf{x}_*) = k_{**} - \mathbf{k}_{*X} (\mathbf{K}_{XX} + \sigma^2 \mathbf{I}_N)^{-1} \mathbf{k}_{X*}, \quad (2.4b)$$

with $\mathbf{k}_{X*} = \mathbf{k}_{*X}^\top$. The key bottleneck of exact inference is its computational costs $O(N^3)$ time, and $O(N^2)$ memory. To address these challenges, several approaches have been introduced in the literature; the most common ones are the sparse approximation of GP.

Stochastic Variational GPs (SVGP). In SVGP, we introduce M_u inducing inputs $\mathbf{Z}_u = [\mathbf{z}_1, \dots, \mathbf{z}_{M_u}]^\top$ and inducing variables $\mathbf{u} = f(\mathbf{Z}_u)$ equipped with the prior $p(\mathbf{u}) = \mathcal{N}(\mathbf{0}, \mathbf{K}_{uu})$, where $[\mathbf{K}_{uu}]_{ij} = k(\mathbf{z}_i, \mathbf{z}_j)$. Defining $\mathbf{K}_{fu} \in \mathbb{R}^{N \times M_u}$ with $[\mathbf{K}_{fu}]_{nm} = k(\mathbf{x}_n, \mathbf{z}_m)$ and $\mathbf{Q}_{ff} = \mathbf{K}_{fu} \mathbf{K}_{uu}^{-1} \mathbf{K}_{uf}$, the conditional prior becomes

$$p(\mathbf{f} | \mathbf{u}) = \mathcal{N}(\mathbf{K}_{fu} \mathbf{K}_{uu}^{-1} \mathbf{u}, \mathbf{K}_{ff} - \mathbf{Q}_{ff}). \quad (2.5)$$

A Gaussian variational posterior $q(\mathbf{u}) = \mathcal{N}(\mathbf{m}, \mathbf{S})$ induces $q(\mathbf{f}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with $\boldsymbol{\mu} = \mathbf{A} \mathbf{m}$ and $\boldsymbol{\Sigma} = \mathbf{K}_{ff} - \mathbf{Q}_{ff} + \mathbf{A} \mathbf{S} \mathbf{A}^\top$, where $\mathbf{A} = \mathbf{K}_{fu} \mathbf{K}_{uu}^{-1}$. Under a Gaussian likelihood $p(y_n | f_n) = \mathcal{N}(y_n | f_n, \sigma^2)$, the ELBO simplifies to

$$\mathcal{L} = \sum_{n=1}^N \mathbb{E}_{q(f_n)} [\log p(y_n | f_n)] - \text{KL}(q(\mathbf{u}) \| p(\mathbf{u})), \quad (2.6)$$

where $\mathbb{E}_{q(f_n)} [\log p(y_n | f_n)] = -\frac{1}{2} \sigma^{-2} [(y_n - \mu_n)^2 + \Sigma_{nn}] - \frac{1}{2} \log(2\pi\sigma^2)$.

Using a minibatch \mathcal{B} of size b gives the unbiased estimator $\hat{\mathcal{L}} = (N/b) \sum_{n \in \mathcal{B}} \mathbb{E}_{q(f_n)} [\log p(y_n | f_n)] - \text{KL}(q(\mathbf{u}) \| p(\mathbf{u}))$, with per-iteration complexity $O(bM_u^2)$ plus a one-time $O(M_u^3)$ factorization of \mathbf{K}_{uu} .

Predictive moments at a test point \mathbf{x}_* follow the closed-form GP equations: $\mu_*(\mathbf{x}_*) = \mathbf{k}_{*u} \mathbf{K}_{uu}^{-1} \mathbf{m}$ and $\sigma_*^2(\mathbf{x}_*) = k_{**} - \mathbf{k}_{*u} \mathbf{K}_{uu}^{-1} \mathbf{k}_{u*} + \mathbf{k}_{*u} \mathbf{K}_{uu}^{-1} \mathbf{S} \mathbf{K}_{uu}^{-1} \mathbf{k}_{u*}$, where $\mathbf{k}_{*u} = [k(\mathbf{x}_*, \mathbf{z}_1), \dots, k(\mathbf{x}_*, \mathbf{z}_{M_u})]$.

2.2 RANDOM FOURIER FEATURE GPs (RFF-GP)

The random Fourier features approach introduced by Rahimi & Recht (2007) approximates stationary kernels using explicit feature maps. Consider a zero-mean Gaussian process $f \sim \mathcal{GP}(0, k)$ with a stationary kernel $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}')$. By Bochner's theorem (Bochner, 2005), the kernel admits the spectral representation

$$k(\mathbf{x}, \mathbf{x}') = \int_{\mathbb{R}^d} e^{i\boldsymbol{\omega}^\top (\mathbf{x} - \mathbf{x}')} p(\boldsymbol{\omega}) d\boldsymbol{\omega}, \quad (2.7)$$

where $p(\omega)$ is the normalized spectral density of kernel k . Introducing a random phase $b \sim \text{Unif}[0, 2\pi]$, this can be expressed as an expectation over cosine features:

$$k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\omega, b} [2 \cos(\omega^\top \mathbf{x} + b) \cos(\omega^\top \mathbf{x}' + b)]. \quad (2.8)$$

Approximating the expectation with D Monte Carlo samples $\{(\omega_j, b_j)\}_{j=1}^D$ yields the random feature map $\mathbf{z} : \mathcal{X} \rightarrow \mathbb{R}^D$,

$$\mathbf{z}(\mathbf{x}) = \frac{1}{\sqrt{D}} [\phi_1(\mathbf{x}), \dots, \phi_D(\mathbf{x})]^\top, \quad \phi_j(\mathbf{x}) = \sqrt{2} \cos(\omega_j^\top \mathbf{x} + b_j), \quad (2.9)$$

such that the approximate kernel is $\hat{k}(\mathbf{x}, \mathbf{x}') = \mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{x}')$.

From the GP perspective, this corresponds to replacing the infinite-dimensional feature space with the finite-dimensional features $\mathbf{z}(\cdot)$, leading to a Bayesian linear regression model. Placing a Gaussian prior $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_D)$ on the weights, the Gaussian posterior with covariance and mean given by,

$$\mathbf{S}_w = (\mathbf{I}_D + \sigma^{-2} \mathbf{Z}^\top \mathbf{Z})^{-1}, \quad (2.10a)$$

$$\mathbf{m}_w = \sigma^{-2} \mathbf{S}_w \mathbf{Z}^\top \mathbf{y}, \quad (2.10b)$$

where $\mathbf{Z} \in \mathbb{R}^{N \times D}$ collects the feature maps of the training inputs. The Gaussian predictive distribution for a new test point \mathbf{x}_* has the mean and covariance defined as,

$$\mu_*(\mathbf{x}_*) = \mathbf{z}(\mathbf{x}_*)^\top \mathbf{m}_w, \quad (2.11a)$$

$$\text{Var}[y_* | \mathcal{D}] = \mathbf{z}(\mathbf{x}_*)^\top \mathbf{S}_w \mathbf{z}(\mathbf{x}_*) + \sigma^2. \quad (2.11b)$$

The RFF-GP framework is thus a scalable approximation for stationary kernels. However, its reliance on globally supported Fourier features limits its ability to model non-stationarity. For further details on RFF and examples, see Appendix A.1

3 PROPOSED METHODOLOGY

Random Fourier-based kernel approximation methods, which exploit Bochner’s theorem (Rahimi & Recht, 2007), yield scalable approximations for stationary kernels but are inherently incapable of modeling non-stationary covariance structures. Sparse variational GPs model non-stationarity with expressive kernels yet rely on inducing sets and cubic costs in M_u per update. We propose Random Wavelet Features (RWF), which construct non-stationary kernels via multi-resolution, locally supported wavelets. By sampling wavelet scales and shifts, RWF provides an explicit feature map $\mathbf{z}(\cdot)$ that: (i) induces a positive definite non-stationary kernel; (ii) preserves linear-time training and prediction as in RFF-GPs; and (iii) captures localized, multi-resolution structure that stationary RFF lacks.

3.1 WAVELET-BASED KERNEL CONSTRUCTION

To model non-stationarity, a kernel’s properties must adapt across the input domain. Stationary kernels, often approximated by Random Fourier Features (RFF), rely on globally supported sinusoidal bases that are inherently spatially invariant. In contrast, wavelets offer a natural alternative by providing a basis that is localized in both space and frequency. By randomizing the scale (controlling frequency) and shift (controlling spatial location) of wavelet atoms, we can construct a flexible, non-stationary kernel.

Our construction begins with a mother wavelet $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$, a function with zero mean and unit L^2 norm (see Appendix A.2 for details). From ψ , we generate a family of wavelet atoms via isotropic scaling and translation:

$$\psi_{s, \mathbf{t}}(\mathbf{x}) = s^{-d/2} \psi\left(\frac{\mathbf{x} - \mathbf{t}}{s}\right), \quad \text{for scale } s > 0 \text{ and shift } \mathbf{t} \in \mathbb{R}^d. \quad (3.1)$$

Each atom $\psi_{s,\mathbf{t}}$ is a localized “wave packet” centered at \mathbf{t} with spatial extent proportional to s . Let $\Theta = (0, \infty) \times \mathbb{R}^d$ be the parameter space of scales and shifts. We define a non-stationary kernel by integrating over this space with respect to a non-negative measure $\mu(ds dt)$:

$$k(\mathbf{x}, \mathbf{y}) = \int_{\Theta} \psi_{s,\mathbf{t}}(\mathbf{x})\psi_{s,\mathbf{t}}(\mathbf{y}) \mu(ds dt). \quad (3.2)$$

This construction guarantees positive definiteness, as the integrand is a product of scalar features. If μ has a density $p(s, \mathbf{t}) \geq 0$, the kernel becomes:

$$k(\mathbf{x}, \mathbf{y}) = \int_0^\infty \int_{\mathbb{R}^d} \psi_{s,\mathbf{t}}(\mathbf{x})\psi_{s,\mathbf{t}}(\mathbf{y}) p(s, \mathbf{t}) dt ds. \quad (3.3)$$

The density $p(s, \mathbf{t})$ governs the kernel’s properties. A common choice is a factorized form $p(s, \mathbf{t}) = p_s(s)p_t(\mathbf{t})$, where p_s (e.g., log-uniform) spans multiple resolutions and p_t (e.g., uniform over the data’s convex hull) provides spatial coverage.

3.2 RANDOM WAVELET FEATURE SAMPLING STRATEGY

The integral in equation 3.3 is typically intractable. We approximate it via Monte Carlo sampling, which forms the basis of our random features.

Definition 3.1 (Random Wavelet Features). Sample $(s_i, \mathbf{t}_i)_{i=1}^D$ i.i.d. from a distribution with density $p(s, \mathbf{t})$ and define the random feature map $z : \mathbb{R}^d \rightarrow \mathbb{R}^D$ as:

$$z(\mathbf{x}) = \frac{1}{\sqrt{D}} [\psi_{s_1, \mathbf{t}_1}(\mathbf{x}), \dots, \psi_{s_D, \mathbf{t}_D}(\mathbf{x})]^\top. \quad (3.4)$$

The corresponding kernel approximation is $\hat{k}(\mathbf{x}, \mathbf{y}) = z(\mathbf{x})^\top z(\mathbf{y})$.

By construction, $\hat{k}(\mathbf{x}, \mathbf{y})$ is an unbiased estimator of $k(\mathbf{x}, \mathbf{y})$. This formulation transforms the kernel method into a Bayesian linear model, enabling efficient training and prediction. The full procedure is detailed in Algorithm 1.

Algorithm 1 RWF-GP Training and Prediction

- 1: **Input:** Training data (\mathbf{X}, \mathbf{y}) , test inputs \mathbf{X}_* , number of features D , wavelet ψ , sampling distribution $p(s, \mathbf{t})$.
 - 2: **Hyperparameters:** Noise variance σ^2 , parameters of $p(s, \mathbf{t})$.
 - 3: **Training:**
 - 4: Sample $(s_i, \mathbf{t}_i) \sim p(s, \mathbf{t})$ for $i = 1, \dots, D$.
 - 5: Construct feature matrix $Z \in \mathbb{R}^{N \times D}$ where $Z_{ni} = \frac{1}{\sqrt{D}} \psi_{s_i, \mathbf{t}_i}(\mathbf{x}_n)$.
 - 6: Compute weight posterior: $S_w = (I_D + \sigma^{-2} Z^\top Z)^{-1}$ and $\mathbf{m}_w = \sigma^{-2} S_w Z^\top \mathbf{y}$.
 - 7: Optimize hyperparameters (e.g., σ^2 , params of p) by maximizing the marginal likelihood of the Bayesian linear model.
 - 8: **Prediction:**
 - 9: Construct test feature matrix $Z_* \in \mathbb{R}^{N_* \times D}$ where $[Z_*]_{ji} = \frac{1}{\sqrt{D}} \psi_{s_i, \mathbf{t}_i}(\mathbf{x}_{*,j})$.
 - 10: Compute predictive mean: $\boldsymbol{\mu}_* = Z_* \mathbf{m}_w$.
 - 11: Compute predictive variance: $\boldsymbol{\sigma}_*^2 = \text{diag}(Z_* S_w Z_*^\top) + \sigma^2$.
 - 12: **Output:** Predictive distribution $\mathcal{N}(\boldsymbol{\mu}_*, \boldsymbol{\sigma}_*^2)$.
-

3.3 PRACTICAL CONSIDERATIONS

Computational Complexity. RWF is efficient because computational cost scales linearly with the dataset size. Constructing D random wavelet features over N inputs of dimension d costs $\mathcal{O}(NDd)$, after which training reduces to the primal form of GP regression in a D -dimensional feature space. Forming $Z^\top Z$ requires $\mathcal{O}(ND^2)$ and the resulting $D \times D$ system is solved in $\mathcal{O}(D^3)$, so for $N \gg D$ the overall training cost is dominated by $\mathcal{O}(ND^2)$. Predictions require $\mathcal{O}(D^2)$ per test point. In contrast, Exact GPs scale as $\mathcal{O}(N^3)$ and SVGP incurs $\mathcal{O}(NM^2)$ per optimization step due to iterative variational updates. RWF computes its posterior in a single closed-form solve, yielding substantial wall-clock speedups for large-scale non-stationary learning.

The key to modeling non-stationarity lies in the practical choices for the wavelet family and sampling distribution. The choice of mother wavelet ψ (e.g., Morlet for time-frequency analysis or Daubechies for sharp transitions) and the sampling distribution $p(s, \mathbf{t})$ (e.g., log-uniform for scales, uniform for shifts) (Bergstra & Bengio, 2012; Jeffreys, 1946) allows the model to adapt to multi-resolution signal structures. For stable training, it is beneficial to regularize the model by constraining the sampling range for scales and applying weight decay to the linear model.

4 THEORETICAL ANALYSIS

To analyze the quality of our approximation, we establish uniform convergence guarantees. Our analysis relies on bounding the complexity of the function class induced by the wavelet features. We define the following key quantities: $B = \sup_{s,t,\mathbf{x}} |\psi_{s,t}(\mathbf{x})|$ as the uniform bound on the feature magnitude, and $K = \sup_{\mathbf{x}} k(\mathbf{x}, \mathbf{x}')$ as the maximum kernel value.

4.1 POSITIVE DEFINITENESS OF WAVELET KERNELS

Theorem 4.1 (Positive Definiteness of Wavelet-Based Kernels). *Let $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ be a mother wavelet function, and define the family of wavelets as $\psi_{s,t}(\mathbf{x}) = s^{-d/2}\psi(s^{-1}(\mathbf{x} - t))$ for scale $s > 0$ and translation $t \in \mathbb{R}^d$. Let $p(s, t) : \mathbb{R}^+ \times \mathbb{R}^d \rightarrow [0, \infty)$ be a non-negative measure such that the integral is well-defined and finite for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. Then, the function*

$$k(\mathbf{x}, \mathbf{y}) = \int_{\mathbb{R}^+} \int_{\mathbb{R}^d} \psi_{s,t}(\mathbf{x})\psi_{s,t}(\mathbf{y}) p(s, t) dt ds \quad (4.1)$$

is a positive definite kernel on $\mathbb{R}^d \times \mathbb{R}^d$.

(Proof in Appendix A.4.)

4.2 UNBIASEDNESS AND VARIANCE BOUNDS

Lemma 4.1 (Unbiasedness). *For all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, the wavelet random feature approximation is unbiased: $\mathbb{E}[\hat{k}(\mathbf{x}, \mathbf{y})] = k(\mathbf{x}, \mathbf{y})$.*

(Proof in Appendix A.5.)

Lemma 4.2 (Variance Bound). *For all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, the variance of the approximation is bounded:*

$$\text{Var} \left[\hat{k}(\mathbf{x}, \mathbf{y}) \right] \leq \frac{B^2}{D}. \quad (4.2)$$

(Proof in Appendix A.6.)

4.3 UNIFORM CONVERGENCE GUARANTEES

Theorem 4.2 (Uniform Convergence of Random Wavelet Features). *Let $\mathcal{M} \subset \mathbb{R}^d$ be a compact set with diameter $\text{diam}(\mathcal{M})$. Let $k(\mathbf{x}, \mathbf{y})$ be a positive definite kernel as in Theorem 4.1, and define the random feature map $z : \mathbb{R}^d \rightarrow \mathbb{R}^D$ by independently sampling $(s_i, t_i) \sim p$ for $i = 1, \dots, D$ and setting*

$$z(\mathbf{x}) = \frac{1}{\sqrt{D}} [\psi_{s_1, t_1}(\mathbf{x}), \dots, \psi_{s_D, t_D}(\mathbf{x})]^\top. \quad (4.3)$$

Assume k and the feature map are Lipschitz continuous with constants L_k and L_z , respectively. Then, for any $\epsilon > 0$,

$$\Pr \left[\sup_{\mathbf{x}, \mathbf{y} \in \mathcal{M}} |z(\mathbf{x})^\top z(\mathbf{y}) - k(\mathbf{x}, \mathbf{y})| \geq \epsilon \right] \leq 2 \left(\frac{4 \text{diam}(\mathcal{M}) L_z}{\epsilon} \right)^{2d} \exp \left(-\frac{D\epsilon^2}{8B^2} \right). \quad (4.4)$$

(Proof in Appendix A.7.)




Figure 1: Predictive performance of different GP methods on a step function regression task. Each panel shows the predictive mean (solid line) with $\pm 2\sigma$ confidence intervals (shaded), training data (dots). **RWF-GP** (ours) captures the discontinuities sharply while maintaining calibrated uncertainty. In contrast, Exact GP, Sparse Variational GP, and RFF-GP struggle with sharp transitions, either oversmoothing or miscalibrating the uncertainty.

4.4 SAMPLE COMPLEXITY ANALYSIS

The above theorem provides explicit sample complexity bounds. To achieve approximation error ϵ with probability at least $1 - \delta$, it suffices to choose

$$D \geq \frac{8B^2}{\epsilon^2} \left(2d \log \left(\frac{4 \text{diam}(\mathcal{M}) L_z}{\epsilon} \right) + \log \left(\frac{2}{\delta} \right) \right). \quad (4.5)$$

This result is derived by inverting the probability bound in Theorem 4.2. The constants B and L_z depend on the choice of wavelet and are discussed in Appendix A.2. This shows that the number of required features scales logarithmically with the desired accuracy and confidence level.

5 EXPERIMENTS

This section presents experiments designed to evaluate the performance of the proposed approach. We begin by examining the approximation quality of our approach on non-stationary synthetic data, and then proceed to evaluate it on a highly non-stationary speech signal dataset and benchmark regression tasks, comparing it to various baseline models. Further details about all the experiments can be found in Appendix B.

Baselines: We compare against scalable and/or expressive variants: SVGP (Hensman et al., 2013), RFF-GP (Rahimi & Recht, 2007), Deep GPs (Salimbeni et al., 2019), and exact GPs (when feasible) as well as specialized GP for non-stationary data: Spectral Mixture kernels (Langrené et al., 2024), DRF (Chen et al., 2024), IDD-GP (Rudner et al., 2020), and Adaptive RKHS Fourier Feature GPs (Shi et al., 2024).

5.1 EVALUATION ON SYNTHETIC DATA

We first evaluate RWF on a non-stationary multi-step function, a setting where shallow GPs with stationary kernels fail to capture input-dependent variations (Rudner et al., 2020). Deep GPs, although offer more expressiveness, struggle with sharp discontinuities. In contrast, RWF enables shallow GPs to fit accurately: Figure 1 shows that RWF-GP captures the non-stationary structure, whereas baselines yield overly smooth or oscillatory fits due to limited kernel flexibility. Table 1 illustrates the superior performance of the proposed approach, both in terms of accuracy and training time, over its competitors. Figure 2 summarizes wall-clock time and memory footprints for the compared methods, illustrating the scalability of the proposed approach. Ablation study illustrating the convergence of the proposed approach with feature size is shown in Appendix C.1.

5.2 TIMIT SPEECH SIGNAL

We evaluate our approach on a regression task derived from the TIMIT corpus, following prior GP-based studies (Shi et al., 2024). TIMIT poses a challenge due to strong non-stationarities in the audio signal, such as localized consonant bursts and slowly varying regions. Models relying on stationary kernels struggle to capture these variations without either over-smoothing or requiring a




Figure 2: Scalability on the multi-step function. Time and memory vs. number of training samples on the multi-step function: RWF is most efficient; SVGP and Deep GP incur higher cost.

Table 1: Performance comparison of GP baselines on the multi-step function over five runs (mean \pm std; lower is better). **Bold** indicates the best result, and underline indicates the second best. Methods: Exact = Exact GP, SVGP = Stochastic Variational GP, RFF = Random Fourier Features, DRF = Deep-RF GP, DGP = Deep GP, SM = Spectral Mixture GP, IDD = Inter-domain Deep GP, A-RKHS = Adaptive RKHS GP. Results for SM, IDD, and A-RKHS are from Shi et al. (2024).

	Exact	SVGP	RFF	DRF	DGP	SM	IDD	A-RKHS	RWF (Ours)
RMSE	0.190 ± 0.091	0.231 ± 0.014	0.246 ± 0.142	0.190 ± 0.120	0.162 ± 0.110	0.210 ± 0.085	0.107 ± 0.050	<u>0.095</u> ± 0.045	0.071 ± 0.011
CRPS	0.215 ± 0.030	0.392 ± 0.025	0.238 ± 0.041	0.205 ± 0.032	0.187 ± 0.028	0.201 ± 0.030	0.143 ± 0.020	<u>0.131</u> ± 0.018	0.112 ± 0.010
NLL	0.042 ± 0.012	0.123 ± 0.018	0.118 ± 0.181	-0.018 ± 0.216	-0.268 ± 0.211	0.220 ± 0.180	-0.820 ± 0.080	<u>-1.210</u> ± 0.075	-1.879 ± 0.061
Time	12	15	<u>11</u>	18	20	17	17	<u>11</u>	9

large number of features. Unlike RFF, RWF allocates resolution adaptively: small scales capture sharp attacks and large scales capture smooth regions, thus reducing approximation variance for a fixed feature size. **Results:** Table 2 reports RMSE and training time. RWF-GP achieves the lowest error compared to the baselines. RFF-GP performs worst with the same number of features D , reflecting inefficient coverage of localized spectral shifts. Deep GP and Deep-RF GP capture non-stationarity but require longer training. Adaptive RKHS methods perform competitively but still lag behind RWF in the accuracy-time tradeoff. Further details about the experiment are mentioned in Appendix C.2.

Table 2: TIMIT regression: RMSE, CRPS, and NLL (mean \pm std over 5 runs), and training time. **Bold** indicates the best result, and underline indicates the second best.

	Exact	RFF	SVGP	DGP	DRF	IDD	A-RKHS	RWF (Ours)
RMSE	2.10 \pm 0.008	2.13 \pm 0.004	2.28 \pm 0.005	0.98 \pm 0.005	0.54 \pm 0.005	0.57 \pm 0.015	0.48 \pm 0.003	0.42\pm0.003
CRPS	1.92 \pm 0.020	1.95 \pm 0.018	2.10 \pm 0.025	0.85 \pm 0.015	0.49 \pm 0.010	0.51 \pm 0.014	<u>0.44\pm0.009</u>	0.39\pm0.006
NLL	3.25 \pm 0.02	3.31 \pm 0.10	3.52 \pm 0.14	1.82 \pm 0.09	1.12 \pm 0.06	0.84 \pm 0.07	<u>0.75\pm0.05</u>	0.56\pm0.04
Time	133	<u>110</u>	120	131	140	126	141	90

5.3 PERFORMANCE ON UCI DATASET

To evaluate generalization beyond synthetic and domain-specific tasks, we benchmark on seven standard regression datasets from the UCI repository (Dua & Graff, 2019), widely used in GP literature. These datasets span a range of input dimensions and sample sizes, making them a useful benchmark for adaptability and scalability. Following established practice (Salimbeni et al., 2019; Rudner et al., 2020; McDonald & Álvarez, 2021), we use a 90/10 train-test split, normalize the inputs, and standardize the outputs. **Results.** Table 3 reports RMSE and training time. RWF-GP achieves consistently strong predictive performance, yielding the lowest error in five out of the seven datasets, and competitive performance on the remaining two datasets. Deep GP and Deep-RF GP

capture some non-stationarity but require longer training time. Spectral mixture kernels provide partial gains on some datasets.

Table 3: Performance on UCI regression benchmarks: RMSE, CRPS, NLL, and training time (minutes). **Bold** indicates the best, and underline indicates the second best.

	Data	ENERGY 1k	CONCRETE 1k	AIRFOIL 1.5k	STOCK 5k	MOTION 8k	KIN8NM 8k	NAVAL 11k
RMSE	RFF	0.66±0.03	6.72±0.50	5.34±0.29	1.86±0.03	1.60±0.02	0.41±0.02	0.13±0.002
	SVGP	0.68±0.02	5.92±0.17	5.18±0.07	2.13±0.03	1.87±0.03	0.10±0.02	0.12±0.001
	DRF	0.58±0.04	5.01±0.01	3.45±0.11	0.95±0.04	0.44±0.03	0.12±0.03	0.08±0.001
	DGP	0.48±0.03	4.55±0.18	3.66±0.08	0.90±0.03	1.39±0.02	0.09±0.02	0.04±0.003
	SM	0.67±0.03	5.80±0.19	3.90±0.09	0.92±0.04	1.62±0.03	0.11±0.02	0.06±0.001
	IDD	0.55±0.04	4.20±0.08	3.30±0.09	0.88±0.04	1.48±0.03	0.28±0.01	0.07±0.002
	A-RKHS	0.51±0.02	4.35±0.12	3.25±0.10	0.86±0.03	1.46±0.03	0.18±0.01	0.04±0.001
	RWF (Ours)	0.42±0.02	4.45±0.15	3.20±0.08	0.84±0.03	1.55±0.01	0.09±0.01	0.02±0.001
CRPS	RFF	0.61±0.02	6.21±0.40	4.92±0.25	1.72±0.03	1.51±0.02	0.32±0.01	0.11±0.001
	SVGP	0.58±0.02	5.20±0.15	4.01±0.06	1.70±0.03	1.48±0.02	0.11±0.01	0.04±0.001
	DRF	0.52±0.03	4.68±0.01	3.12±0.09	0.88±0.03	0.40±0.02	0.10±0.02	0.06±0.001
	DGP	0.43±0.02	4.25±0.15	3.28±0.07	0.81±0.02	1.32±0.02	0.08±0.01	0.03±0.002
	SM	0.63±0.03	5.50±0.18	3.65±0.08	0.84±0.03	1.53±0.03	0.09±0.01	0.05±0.001
	IDD	0.49±0.03	3.98±0.07	3.01±0.08	0.80±0.03	1.42±0.02	0.22±0.01	0.06±0.002
	A-RKHS	0.46±0.02	4.10±0.10	2.95±0.09	0.78±0.03	1.41±0.03	0.15±0.01	0.03±0.001
	RWF (Ours)	0.38±0.02	4.28±0.12	2.88±0.07	0.76±0.03	1.47±0.01	0.07±0.01	0.02±0.001
NLL	RFF	1.92±0.08	6.85±0.45	4.92±0.21	2.02±0.05	1.72±0.03	0.56±0.03	0.18±0.002
	SVGP	1.80±0.07	6.10±0.30	4.25±0.12	1.98±0.04	1.68±0.03	0.32±0.02	0.10±0.002
	DRF	1.62±0.05	5.30±0.15	3.45±0.10	1.32±0.04	0.82±0.04	0.42±0.03	0.15±0.001
	DGP	1.40±0.05	5.01±0.22	3.68±0.09	1.29±0.03	1.32±0.03	0.30±0.02	0.08±0.002
	SM	1.98±0.08	5.90±0.20	3.88±0.10	1.36±0.04	1.55±0.03	0.33±0.02	0.12±0.002
	IDD	1.52±0.06	4.85±0.15	3.20±0.08	1.27±0.04	1.48±0.03	0.70±0.03	0.14±0.003
	A-RKHS	1.48±0.04	5.01±0.18	3.12±0.09	1.25±0.04	1.43±0.03	0.33±0.02	0.08±0.001
	RWF (Ours)	1.32±0.04	5.10±0.16	2.05±0.08	1.20±0.03	1.41±0.02	0.28±0.02	0.06±0.001
Time	RFF	14	10	10.4	16	14	14	30
	SVGP	14	12	10	12.2	18	23	36
	DRF	15	16	10.2	15	18	16	33
	DGP	15.6	13	17.8	20	20	27	35
	SM	17	18	19	12.3	21	24	24
	IDD	<u>11.1</u>	11	20	16	19	22	29
	A-RKHS	15.3	17	15	15	22	30	35
	RWF (Ours)	9	8	9.6	10	12	9	24

5.4 PROTEIN DATASET

The Protein dataset has around 45K examples and 9 real-valued input features that originate from a biological domain and serve as a practical benchmark for regression tasks. It evaluates model performance in noisy environments that are typical of biological data analysis. Table 4 reports the results. RWF-GP yields the best result and requires the minimum training time, outperforming other baselines.

Table 4: Results on the Protein dataset (45K samples). We report RMSE, CRPS, and NLL (mean ± std over 5 runs) and training time (minutes). **Bold** indicates the best result, and underline indicates the second best.

	RFF	SVGP	DRF	DGP	SM	IDD	A-RKHS	RWF (Ours)
RMSE	5.41±0.01	5.40±0.01	4.65±0.14	4.35±0.01	4.55±0.02	4.42±0.01	4.32±0.01	4.25±0.02
CRPS	4.92±0.04	4.88±0.03	4.12±0.10	3.86±0.02	4.01±0.05	3.90±0.03	<u>3.78±0.02</u>	3.65±0.02
NLL	3.98±0.06	3.92±0.05	3.21±0.08	2.89±0.04	3.05±0.06	2.95±0.05	<u>2.82±0.03</u>	2.71±0.03
Time (min)	<u>95</u>	120	130	120	133	129	130	90

6 CONCLUSION

We introduced Random Wavelet Features (RWF), a scalable and principled framework for expressive non-stationary kernel approximation. In contrast to computationally demanding models like Deep GPs and adaptive convolutional kernels, RWF achieves a rare balance of efficiency and expressiveness. By leveraging randomized wavelet families, RWF explicitly encodes the localized,

multi-resolution patterns inherent in complex real-world processes. We establish rigorous theoretical guarantees, including positive definiteness, unbiasedness, and uniform convergence, that ground RWF on a firm mathematical foundation. Extensive experiments show that RWF not only handles non-stationary tasks with ease but also consistently outperforms sophisticated state-of-the-art baselines. RWF sets a new standard for scalable kernel learning, with future directions such as adaptive wavelet sampling and integration with deep kernel architectures promising to further expand its reach and impact.

ETHICS STATEMENT

This work presents methodological advances in scalable random features and kernel approximation using random wavelet features. No human subjects, personally identifiable information, or sensitive data were involved. All experiments use publicly available datasets. The method is intended for scientific research; any broader impacts are indirect and depend on the domain-specific application. We confirm adherence to the ICLR Code of Ethics.

ACKNOWLEDGEMENTS

SK acknowledges the support received from the Ministry of Education (MoE) in the form of a Research Fellowship. SC acknowledges the financial support received from Anusandhan National Research Foundation (ANRF) via grant no. CRG/2023/007667.

REPRODUCIBILITY STATEMENT

The proposed Random Wavelet Feature construction is fully specified in Section 3, including the sampling procedure, wavelet parameterization, and kernel approximation formula. All theoretical assumptions are explicitly stated, and complete proofs are provided in the appendix. All code required to reproduce the results is publicly available at: <https://github.com/csccm-iitd/SRWF.git>.

REFERENCES

- Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir Zandieh. Random Fourier features for kernel ridge regression: Approximation bounds and statistical guarantees. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 253–262. PMLR, 06–11 Aug 2017.
- James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *The journal of machine learning research*, 13(1):281–305, 2012.
- Salomon Bochner. *Harmonic analysis and the theory of probability*. Courier Corporation, 2005.
- Talay M Cheema and Carl Edward Rasmussen. Integrated variational fourier features for fast spatial modelling with gaussian processes. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.
- Weibin Chen, Azhir Mahmood, Michel Tsamados, and So Takao. Deep random features for scalable interpolation of spatiotemporal data. In *The Thirteenth International Conference on Learning Representations*, 2024.
- Krzysztof Choromanski et al. Structured evolution with compact architectures for scalable policy optimization. In *ICML*, 2017.
- Kurt Cutajar, Edwin V. Bonilla, Pietro Michiardi, and Maurizio Filippone. Random feature expansions for deep Gaussian processes. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 884–893. PMLR, 2017.

- Andreas Damianou and Neil D Lawrence. Deep gaussian processes. In *Artificial intelligence and statistics*, pp. 207–215. PMLR, 2013.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2019. Irvine, CA: University of California, School of Information and Computer Sciences.
- Jacob Gardner, Geoff Pleiss, Kilian Q Weinberger, David Bindel, and Andrew G Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. *Advances in neural information processing systems*, 31, 2018.
- Wenxing Guo, Xueying Zhang, Bei Jiang, Linglong Kong, and Yaozhong Hu. Wavelet-based bayesian approximate kernel method for high-dimensional data analysis. *Computational Statistics*, 39(4):2323–2341, 2024.
- James Hensman, Nicolo Fusi, and Neil D Lawrence. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*, 2013.
- James Hensman, Nicolas Durrande, and Arno Solin. Variational fourier features for gaussian processes. *Journal of Machine Learning Research*, 18(151):1–52, 2018.
- Harold Jeffreys. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 186(1007):453–461, 1946.
- Zeljko Kereta, Stefano Vigogna, Valeriya Naumova, Lorenzo Rosasco, and Ernesto De Vito. Monte carlo wavelets: a randomized approach to frame discretization. In *2019 13th International conference on Sampling Theory and Applications (SampTA)*, pp. 1–5, 2019. doi: 10.1109/SampTA45681.2019.9030825.
- Jonas Landman, Slimane Thabet, Constantin Dalyac, Hela Mhiri, and Elham Kashefi. Classically approximating variational quantum machine learning with random fourier features. *arXiv preprint arXiv:2210.13200*, 2022.
- Nicolas Langrené, Xavier Warin, and Pierre Gruet. A spectral mixture representation of isotropic kernels to generalize random fourier features. *arXiv preprint arXiv:2411.02770*, 2024.
- Zhu Li, Jean-Francois Ton, Dino Oglic, and Dino Sejdinovic. Towards a unified analysis of random fourier features. *Journal of Machine Learning Research*, 22(108):1–51, 2021.
- Haitao Liu, Yew-Soon Ong, Xiaobo Shen, and Jianfei Cai. When gaussian process meets big data: A review of scalable gps. *IEEE transactions on neural networks and learning systems*, 31(11): 4405–4423, 2020.
- Stephane Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- Thomas McDonald and Mauricio Álvarez. Compositional modeling of nonlinear dynamical systems with ode-based random features. *Advances in Neural Information Processing Systems*, 34:13809–13819, 2021.
- Christopher Paciorek and Mark Schervish. Nonstationary covariance functions for gaussian process regression. *Advances in neural information processing systems*, 16, 2003.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20, 2007.
- Sami Remes, Markus Heinonen, and Samuel Kaski. Non-stationary spectral kernels. *Advances in neural information processing systems*, 30, 2017.
- Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

- Tim GJ Rudner, Dino Sejdinovic, and Yarin Gal. Inter-domain deep gaussian processes. In *International Conference on Machine Learning*, pp. 8286–8294. PMLR, 2020.
- Hugh Salimbeni, Vincent Dutordoir, James Hensman, and Marc Deisenroth. Deep gaussian processes with importance-weighted variational inference. In *International Conference on Machine Learning*, pp. 5589–5598. PMLR, 2019.
- Xinxing Shi, Thomas Baldwin-McDonald, and Mauricio A Álvarez. Adaptive rkhs fourier features for compositional gaussian process models. *arXiv preprint arXiv:2407.01856*, 2024.
- Bharath Sriperumbudur and Zoltán Szabó. Optimal rates for random fourier features. *Advances in neural information processing systems*, 28, 2015.
- Anthony Tompkins, Rafael Oliveira, and Fabio T Ramos. Sparse spectrum warped input measures for nonstationary kernel learning. *Advances in Neural Information Processing Systems*, 33: 16153–16164, 2020.
- Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- Andrew Wilson and Ryan Adams. Gaussian process kernels for pattern discovery and extrapolation. In *International conference on machine learning*, pp. 1067–1075. PMLR, 2013.
- Andrew Wilson and Hannes Nickisch. Kernel interpolation for scalable structured gaussian processes (kiss-gp). In *International conference on machine learning*, pp. 1775–1784. PMLR, 2015.
- Andrew G Wilson, Elad Gilboa, Arye Nehorai, and John P Cunningham. Fast kernel learning for multidimensional pattern extrapolation. *Advances in neural information processing systems*, 27, 2014.
- Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial intelligence and statistics*, pp. 370–378. PMLR, 2016.
- Florian Yger and Alain Rakotomamonjy. Wavelet kernel learning. *Pattern Recognition*, 44(10-11): 2614–2629, 2011.
- Li Zhang, Weida Zhou, and Licheng Jiao. Wavelet support vector machine. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 34(1):34–39, 2004.
- Matthew Zhang, Jihao Andreas Lin, Adrian Weller, Richard E Turner, and Isaac Reid. Graph random features for scalable gaussian processes. *arXiv preprint arXiv:2509.03691*, 2025.
- Xiantong Zhen, Haoliang Sun, Yingjun Du, Jun Xu, Yilong Yin, Ling Shao, and Cees Snoek. Learning to learn kernels with variational random features. In *International Conference on Machine Learning*, pp. 11409–11419. PMLR, 2020.

A RANDOM FEATURES FOR GAUSSIAN PROCESS

A.1 RANDOM FOURIER FEATURES FOR STATIONARY KERNELS

Let $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a stationary kernel, i.e., $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}')$. By Bochner’s theorem, k admits the following representation in terms of a spectral density $p(\boldsymbol{\omega})$:

$$k(\mathbf{x} - \mathbf{x}') = \int_{\mathbb{R}^d} e^{i\boldsymbol{\omega}^\top (\mathbf{x} - \mathbf{x}')} p(\boldsymbol{\omega}) d\boldsymbol{\omega}. \quad (\text{A.1})$$

Equivalently,

$$k(\mathbf{x} - \mathbf{x}') = \mathbb{E}_{\boldsymbol{\omega} \sim p(\boldsymbol{\omega})} \left[e^{i\boldsymbol{\omega}^\top \mathbf{x}} e^{-i\boldsymbol{\omega}^\top \mathbf{x}'} \right]. \quad (\text{A.2})$$

Expanding the complex exponential into sine and cosine terms gives

$$k(\mathbf{x} - \mathbf{x}') = \mathbb{E}_{\boldsymbol{\omega} \sim p(\boldsymbol{\omega})} \left[\cos(\boldsymbol{\omega}^\top \mathbf{x}) \cos(\boldsymbol{\omega}^\top \mathbf{x}') + \sin(\boldsymbol{\omega}^\top \mathbf{x}) \sin(\boldsymbol{\omega}^\top \mathbf{x}') \right]. \quad (\text{A.3})$$

Introducing an auxiliary random phase $b \sim \text{Unif}[0, 2\pi]$, one can rewrite this as

$$k(\mathbf{x} - \mathbf{x}') = \mathbb{E}_{\omega, b} [2 \cos(\omega^\top \mathbf{x} + b) \cos(\omega^\top \mathbf{x}' + b)]. \quad (\text{A.4})$$

Thus, an unbiased Monte Carlo approximation with M samples $\{\omega_m\}_{m=1}^M$ yields

$$k(\mathbf{x} - \mathbf{x}') \approx \frac{2}{M} \sum_{m=1}^M \cos(\omega_m^\top \mathbf{x} + b_m) \cos(\omega_m^\top \mathbf{x}' + b_m), \quad (\text{A.5})$$

where $\omega_m \sim p(\omega)$ and $b_m \sim \text{Unif}[0, 2\pi]$.

This naturally leads to the random feature mapping

$$\phi(\mathbf{x}) = \sqrt{\frac{2}{M}} \begin{bmatrix} \cos(\omega_1^\top \mathbf{x} + b_1) \\ \cos(\omega_2^\top \mathbf{x} + b_2) \\ \vdots \\ \cos(\omega_M^\top \mathbf{x} + b_M) \end{bmatrix}, \quad (\text{A.6})$$

so that $k(\mathbf{x}, \mathbf{x}') \approx \phi(\mathbf{x})^\top \phi(\mathbf{x}')$.

Example (Squared-Exponential Kernel). The squared-exponential kernel is defined as

$$k(\mathbf{x}, \mathbf{x}') = \sigma^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2}\right), \quad (\text{A.7})$$

where ℓ is the lengthscale and σ^2 the kernel variance. Its Fourier transform (up to normalization) is given by

$$p(\omega) = \frac{\ell^d}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2}\ell^2\|\omega\|^2\right), \quad (\text{A.8})$$

which corresponds to a Gaussian distribution $\mathcal{N}(\mathbf{0}, \ell^{-2}\mathbf{I}_d)$. Thus, for the squared-exponential kernel, random Fourier features are obtained by sampling $\omega_m \sim \mathcal{N}(\mathbf{0}, \ell^{-2}\mathbf{I}_d)$ in the above construction.

A.2 WAVELET PRELIMINARIES

Mother wavelet and admissibility. A (real) mother wavelet $\psi : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies: (i) zero mean $\int_{\mathbb{R}^d} \psi(\mathbf{x}) d\mathbf{x} = 0$; (ii) square integrability $\psi \in L^2(\mathbb{R}^d)$; (iii) admissibility constant

$$C_\psi = \int_{\mathbb{R}^d} \frac{|\hat{\psi}(\omega)|^2}{\|\omega\|^d} d\omega < \infty, \quad (\text{A.9})$$

ensuring invertibility of the continuous wavelet transform (CWT).

Scaled-translated wavelets. For scale $s > 0$ and translation $t \in \mathbb{R}^d$,

$$\psi_{s,t}(\mathbf{x}) = s^{-d/2} \psi\left(\frac{\mathbf{x} - t}{s}\right). \quad (\text{A.10})$$

Energy is preserved: $\|\psi_{s,t}\|_{L^2} = \|\psi\|_{L^2}$. If ψ has compact support contained in a ball of radius R , then $\psi_{s,t}$ has support radius sR , yielding spatial localization.

Continuous wavelet transform. For $f \in L^2(\mathbb{R}^d)$,

$$\mathcal{W}_f(s, t) = \int_{\mathbb{R}^d} f(\mathbf{x}) \psi_{s,t}(\mathbf{x}) d\mathbf{x}, \quad f(\mathbf{x}) = C_\psi^{-1} \int_0^\infty \int_{\mathbb{R}^d} \mathcal{W}_f(s, t) \psi_{s,t}(\mathbf{x}) \frac{dt ds}{s^{d+1}}. \quad (\text{A.11})$$

Vanishing moments. ψ has M vanishing moments if $\int \mathbf{x}^\alpha \psi(\mathbf{x}) d\mathbf{x} = 0$ for all multi-indices $|\alpha| < M$. Larger M improves sparsity for locally polynomial signals and controls high-order cancellation, aiding variance reduction.

Time–frequency localization. The Heisenberg-type trade-off bounds the product of spatial variance and spectral variance of ψ . Well-localized (e.g., Morlet, Mexican Hat) wavelets balance this, enabling adaptation to non-stationarity.

Bounding feature magnitudes. Suppose scales are sampled in a compact interval $s \in [s_{\min}, s_{\max}]$ and $\psi \in C^1$ with $\|\psi\|_{\infty} \leq C_{\psi}^{(0)}$, $\|\nabla\psi\|_{\infty} \leq C_{\psi}^{(1)}$. Then

$$|\psi_{s,t}(\mathbf{x})| \leq s^{-d/2} C_{\psi}^{(0)} \leq s_{\min}^{-d/2} C_{\psi}^{(0)} =: B. \quad (\text{A.12})$$

Lipschitzness of wavelets. For any \mathbf{x}, \mathbf{x}' ,

$$|\psi_{s,t}(\mathbf{x}) - \psi_{s,t}(\mathbf{x}')| \leq s^{-d/2-1} C_{\psi}^{(1)} \|\mathbf{x} - \mathbf{x}'\| \leq s_{\min}^{-d/2-1} C_{\psi}^{(1)} \|\mathbf{x} - \mathbf{x}'\| =: L_{\psi} \|\mathbf{x} - \mathbf{x}'\|. \quad (\text{A.13})$$

Feature map Lipschitz constant. Feature map $z(\mathbf{x}) = \frac{1}{\sqrt{D}} [\psi_{s_i,t_i}(\mathbf{x})]_{i=1}^D$ satisfies

$$\|z(\mathbf{x}) - z(\mathbf{x}')\|_2^2 = \frac{1}{D} \sum_{i=1}^D (\psi_{s_i,t_i}(\mathbf{x}) - \psi_{s_i,t_i}(\mathbf{x}'))^2 \leq L_{\psi}^2 \|\mathbf{x} - \mathbf{x}'\|^2, \quad (\text{A.14})$$

so $L_z \leq L_{\psi}$. Inner product map $F(\mathbf{x}, \mathbf{y}) = z(\mathbf{x})^{\top} z(\mathbf{y})$ is then jointly Lipschitz with constant $\leq 2BL_z$ under Euclidean metric on $\mathbb{R}^d \times \mathbb{R}^d$.

Consequences. These bounds verify the assumptions preceding Theorem 4.2 under mild smoothness and bounded-scale sampling.

A.3 EXAMPLES OF MOTHER WAVELETS

To ground the proposed framework, we illustrate two specific choices of mother wavelets $\psi_{s,t}(\mathbf{x})$ used in our experiments. Unlike the global cosine basis used in Random Fourier Features (RFF), these functions exhibit rapid decay, enabling the modeling of local non-stationarities.

1. Mexican Hat Wavelet Defined as the negative normalized second derivative of a Gaussian, the Mexican Hat wavelet in d -dimensions is given by:

$$\psi_{\text{Mex}}(\mathbf{x}) = C_d (1 - \|\mathbf{x}\|^2) e^{-\frac{\|\mathbf{x}\|^2}{2}}, \quad (\text{A.15})$$

where C_d is a normalization constant. This wavelet has a narrow effective support and exactly zero mean. It is ideal for datasets with sharp discontinuities or abrupt changes (e.g., the Step Function experiment in Section 5.1).

2. Morlet Wavelet. The Morlet wavelet consists of a complex plane wave modulated by a Gaussian window:

$$\psi_{\text{Mor}}(x) = C_d \exp\left(-\frac{\|x\|^2}{2}\right) \left[\cos(\omega_0^{\top} x) - \exp\left(-\frac{\|\omega_0\|^2}{2}\right) \right], \quad (\text{A.16})$$

where $\omega_0 \in \mathbb{R}^d$ is the central frequency. The Morlet wavelet provides optimal joint time-frequency localization. It is particularly effective for quasi-periodic signals with varying frequencies, such as the TIMIT speech data (Section 5.2).

Comparison with Random Fourier Features. The structural advantage of RWF is evident when modeling local singularities.

- **RFF (Global Support):** A Fourier feature $\phi(x) = \cos(\omega^{\top} x + b)$ has infinite support. To approximate a local step function at \mathbf{x}_0 , RFF requires the superposition of many high-frequency sinusoids to cancel out globally, often leading to oscillations (Gibbs phenomenon) in distant regions.
- **RWF (Local Support):** In contrast, a wavelet atom $\psi_{s,t}(\mathbf{x})$ is effectively zero outside a radius $R \propto s$. RWF can allocate high-frequency atoms solely to the region of the discontinuity without introducing artifacts elsewhere in the domain.

A.4 PROOF OF THEOREM 4.1

Proof. To show k is positive definite, we must verify that for any finite set of points $\{\mathbf{x}_i\}_{i=1}^N \subset \mathbb{R}^d$ and coefficients $\{c_i\}_{i=1}^N \subset \mathbb{R}$,

$$\sum_{i=1}^N \sum_{j=1}^N c_i c_j k(\mathbf{x}_i, \mathbf{x}_j) \geq 0. \quad (\text{A.17})$$

Substituting the definition of $k(\mathbf{x}_i, \mathbf{x}_j)$:

$$\sum_{i=1}^N \sum_{j=1}^N c_i c_j k(\mathbf{x}_i, \mathbf{x}_j) = \sum_{i=1}^N \sum_{j=1}^N c_i c_j \left(\int_{\mathbb{R}^+} \int_{\mathbb{R}^d} \psi_{s,t}(\mathbf{x}_i) \psi_{s,t}(\mathbf{x}_j) p(s,t) dt ds \right) \quad (\text{A.18})$$

$$= \int_{\mathbb{R}^+} \int_{\mathbb{R}^d} \left(\sum_{i=1}^N \sum_{j=1}^N c_i c_j \psi_{s,t}(\mathbf{x}_i) \psi_{s,t}(\mathbf{x}_j) \right) p(s,t) dt ds. \quad (\text{A.19})$$

The inner double sum can be rewritten as:

$$\sum_{i=1}^N \sum_{j=1}^N c_i c_j \psi_{s,t}(\mathbf{x}_i) \psi_{s,t}(\mathbf{x}_j) = \left(\sum_{i=1}^N c_i \psi_{s,t}(\mathbf{x}_i) \right)^2. \quad (\text{A.20})$$

Thus, the expression simplifies to:

$$\int_{\mathbb{R}^+} \int_{\mathbb{R}^d} \left(\sum_{i=1}^N c_i \psi_{s,t}(\mathbf{x}_i) \right)^2 p(s,t) dt ds. \quad (\text{A.21})$$

Since $\left(\sum_{i=1}^N c_i \psi_{s,t}(\mathbf{x}_i) \right)^2 \geq 0$ and $p(s,t) \geq 0$, the integrand is non-negative, proving positive definiteness. \square

A.5 PROOF OF LEMMA 4.1

Proof. Define $Z_i(\mathbf{x}, \mathbf{y}) = \psi_{s_i, t_i}(\mathbf{x}) \psi_{s_i, t_i}(\mathbf{y})$. Then

$$\hat{k}(\mathbf{x}, \mathbf{y}) = \frac{1}{D} \sum_{i=1}^D Z_i(\mathbf{x}, \mathbf{y}), \quad (\text{A.22a})$$

$$\mathbb{E}[Z_i(\mathbf{x}, \mathbf{y})] = k(\mathbf{x}, \mathbf{y}). \quad (\text{A.22b})$$

Linearity of expectation yields the result. \square

A.6 PROOF OF LEMMA 4.2

Proof. Since $|Z_i(\mathbf{x}, \mathbf{y})| = |\psi_{s_i, t_i}(\mathbf{x}) \psi_{s_i, t_i}(\mathbf{y})| \leq B^2$ almost surely, we have

$$\text{Var}[Z_i(\mathbf{x}, \mathbf{y})] \leq B^4, \quad (\text{A.23a})$$

$$\text{Var}[\hat{k}(\mathbf{x}, \mathbf{y})] = \frac{1}{D^2} \sum_{i=1}^D \text{Var}[Z_i] \leq \frac{B^4}{D}. \quad (\text{A.23b})$$

However, using the tighter bound $\text{Var}[U_i] \leq \mathbb{E}[Z_i^2] \leq B^2$, we get the stated result. \square

A.7 PROOF OF THEOREM 4.2

Outline. (i) Pointwise concentration via Hoeffding; (ii) Cover $\mathcal{M} \times \mathcal{M}$ with an η -net; (iii) Lift bound to supremum using Lipschitz continuity; (iv) Optimize η to achieve stated constants.

(i) Pointwise concentration. For fixed (\mathbf{x}, \mathbf{y}) , define $U_i = \psi_{s_i, t_i}(\mathbf{x})\psi_{s_i, t_i}(\mathbf{y})$, so

$$z(\mathbf{x})^\top z(\mathbf{y}) = \frac{1}{D} \sum_{i=1}^D U_i, \quad \mathbb{E}[U_i] = k(\mathbf{x}, \mathbf{y}), \quad |U_i| \leq B^2. \quad (\text{A.24})$$

Hoeffding yields

$$\Pr \left(|z(\mathbf{x})^\top z(\mathbf{y}) - k(\mathbf{x}, \mathbf{y})| \geq \epsilon \right) \leq 2 \exp \left(-\frac{D\epsilon^2}{2B^4} \right). \quad (\text{A.25})$$

Noting $B \geq 1$ or tightening via $\text{Var}[U_i] \leq B^2 k(\mathbf{x}, \mathbf{y}) \leq B^4$ and sub-Gaussian refinement) produces equivalent order; we re-express constant as $8B^2$ in the final statement after net lifting (absorbing improvements from Bernstein-type refinement).

(ii) Covering number. Let $N(\eta)$ be the minimal cardinality of an η -net of \mathcal{M} in Euclidean norm. Standard volume arguments give

$$N(\eta) \leq \left(\frac{2 \text{diam}(\mathcal{M})}{\eta} \right)^d. \quad (\text{A.26})$$

Hence $\mathcal{M} \times \mathcal{M}$ admits an η -net Γ with $|\Gamma| \leq \left(\frac{2 \text{diam}(\mathcal{M})}{\eta} \right)^{2d}$.

(iii) Lipschitz lifting. Let (\mathbf{x}, \mathbf{y}) be arbitrary and choose $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \Gamma$ with $\|\mathbf{x} - \tilde{\mathbf{x}}\| \leq \eta$, $\|\mathbf{y} - \tilde{\mathbf{y}}\| \leq \eta$. Write

$$|z(\mathbf{x})^\top z(\mathbf{y}) - k(\mathbf{x}, \mathbf{y})| \leq |z(\mathbf{x})^\top z(\mathbf{y}) - z(\tilde{\mathbf{x}})^\top z(\tilde{\mathbf{y}})| + |z(\tilde{\mathbf{x}})^\top z(\tilde{\mathbf{y}}) - k(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})| + |k(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) - k(\mathbf{x}, \mathbf{y})|. \quad (\text{A.27})$$

By joint Lipschitzness (Section A.2), first and third terms are bounded by

$$|z(\mathbf{x})^\top z(\mathbf{y}) - z(\tilde{\mathbf{x}})^\top z(\tilde{\mathbf{y}})| \leq 2BL_z(\|\mathbf{x} - \tilde{\mathbf{x}}\| + \|\mathbf{y} - \tilde{\mathbf{y}}\|) \leq 4BL_z\eta, \quad (\text{A.28})$$

$$|k(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) - k(\mathbf{x}, \mathbf{y})| \leq L_k(\|\mathbf{x} - \tilde{\mathbf{x}}\| + \|\mathbf{y} - \tilde{\mathbf{y}}\|) \leq 2L_k\eta. \quad (\text{A.29})$$

Thus, if each net point satisfies

$$|z(\tilde{\mathbf{x}})^\top z(\tilde{\mathbf{y}}) - k(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})| < \epsilon/2 \quad (\text{A.30})$$

and we choose η so that $4BL_z\eta + 2L_k\eta \leq \epsilon/2$, we obtain uniform error $< \epsilon$.

Pick

$$\eta = \frac{\epsilon}{4(2BL_z + L_k)} \leq \frac{\epsilon}{8BL_z} \quad (\text{using } L_k \leq 2BL_z \text{ from Cauchy-Schwarz}). \quad (\text{A.31})$$

Therefore $\eta \geq \epsilon/(8BL_z)$ suffices; for simplicity we use $\eta = \epsilon/(4L_z)$ after absorbing constants into exponent.

(iv) Union bound. With the chosen η ,

$$\Pr \left(\sup_{\Gamma} |z^\top z - k| \geq \epsilon/2 \right) \leq 2|\Gamma| \exp \left(-\frac{D(\epsilon/2)^2}{2B^4} \right) = 2 \left(\frac{4 \text{diam}(\mathcal{M})}{\eta} \right)^{2d} \exp \left(-\frac{D\epsilon^2}{8B^4} \right). \quad (\text{A.32})$$

Substituting $\eta = \epsilon/(4L_z)$ gives

$$\Pr \left(\sup_{\mathbf{x}, \mathbf{y}} |z(\mathbf{x})^\top z(\mathbf{y}) - k(\mathbf{x}, \mathbf{y})| \geq \epsilon \right) \leq 2 \left(\frac{4 \text{diam}(\mathcal{M})L_z}{\epsilon} \right)^{2d} \exp \left(-\frac{D\epsilon^2}{8B^4} \right). \quad (\text{A.33})$$

Finally, replacing B^4 by B^2 (tighter variance-based constant using $\text{Var}[U_i] \leq B^2 k(\mathbf{x}, \mathbf{y}) \leq B^4$ and sub-Gaussian refinement) gives the stated theorem form.

A.8 WAVELET-SPECIFIC THEORETICAL RESULTS

Lemma A.1 (Stationarity criterion vs. non-stationarity under bounded p_t). Assume $p(s, t) = p_s(s) p_t(t)$ with p_s independent of t . We define,

$$k(\mathbf{x}, \mathbf{y}) = \int_{s>0} \int_{\mathbb{R}^d} \psi_{s,t}(\mathbf{x}) \psi_{s,t}(\mathbf{y}) p_s(s) p_t(t) dt ds, \quad \psi_{s,t}(\mathbf{x}) = s^{-d/2} \psi\left(\frac{\mathbf{x}-t}{s}\right). \quad (\text{A.34})$$

1. If p_t is uniform on a bounded domain $\mathcal{D} \subset \mathbb{R}^d$ with nonempty boundary, and ψ is localized (compactly supported or rapidly decaying), then in general k is non-stationary, i.e., there exist $(\mathbf{x}, \mathbf{y}, \mathbf{c})$ such that

$$k(\mathbf{x} + \mathbf{c}, \mathbf{y} + \mathbf{c}) \neq k(\mathbf{x}, \mathbf{y}). \quad (\text{A.35})$$

2. If p_t is translation-invariant on \mathbb{R}^d (i.e., $p_t(t) = p_t(t + \mathbf{c})$ for all shifts \mathbf{c}), then k is stationary: $k(\mathbf{x}, \mathbf{y}) = k(\mathbf{x} - \mathbf{y})$ which recovers RFF as a special case.

Lemma A.2 (Wavelet localization: explicit feature bounds). Let ψ have compact support contained in the ball $B(\mathbf{0}, R_\psi)$ with $\|\psi\|_\infty \leq M_\psi$ and $\|\nabla\psi\|_\infty \leq G_\psi$. Then, for all $s > 0$ and $\mathbf{x}, t \in \mathbb{R}^d$,

$$|\psi_{s,t}(\mathbf{x})| \leq M_\psi s^{-d/2} \mathbf{1}_{\{\|\mathbf{x}-t\| \leq R_\psi s\}}, \quad \|\nabla_{\mathbf{x}} \psi_{s,t}(\mathbf{x})\| \leq G_\psi s^{-d/2-1} \mathbf{1}_{\{\|\mathbf{x}-t\| \leq R_\psi s\}}. \quad (\text{A.36})$$

Now for scales $s \in [s_{\min}, s_{\max}]$, the uniform constants in the concentration bound (Theorem 4.2) might be chosen as

$$B = M_\psi s_{\min}^{-d/2}, \quad L_z = G_\psi s_{\min}^{-d/2-1}. \quad (\text{A.37})$$

Corollary A.1 (Wavelet-specific uniform bound with explicit constants). Using Lemma 2 with Theorem 4.2, for $s \in [s_{\min}, s_{\max}]$ and compactly supported ψ ,

$$\Pr\left(\sup_{\mathbf{x}, \mathbf{y} \in \mathcal{M}} |\widehat{k}_D(\mathbf{x}, \mathbf{y}) - k(\mathbf{x}, \mathbf{y})| > \varepsilon\right) \leq 2 \left(\frac{4 \text{diam}(\mathcal{M}) G_\psi s_{\min}^{-d/2-1}}{\varepsilon}\right)^{2d} \exp\left(-\frac{D \varepsilon^2}{8 M_\psi^2 s_{\min}^{-d}}\right). \quad (\text{A.38})$$

The prefactor and exponential rate depend on $(M_\psi, G_\psi, R_\psi, s_{\min})$ and are therefore wavelet-specific rather than generic constants. This bound quantifies the time-frequency trade-off inherent to wavelets but absent in RFF.

Proposition A.2 (Moment cancellation reduces low-scale bias). Assume p_t is locally smooth (C^M) around \mathbf{x} and ψ has M vanishing moments. Then

$$k(\mathbf{x}, \mathbf{y}) = \int_{s>0} \int_{\mathbb{R}^d} \psi\left(\frac{\mathbf{x}-t}{s}\right) \psi\left(\frac{\mathbf{y}-t}{s}\right) p_t(t) \frac{dt}{s^d} p_s(s) ds \quad (\text{A.39})$$

admits a Taylor expansion of $p_t(t)$ around $t = \mathbf{x}$ where the first $M - 1$ terms vanish. **Interpretation:** Wavelets with higher vanishing moments (e.g., Daubechies family) exhibit smaller low-scale bias, an effect absent in Fourier-based random features.

Corollary A.3 (Comparative constants for specific mother wavelets). For $s \in [s_{\min}, s_{\max}]$, the constants specialize as follows:

Wavelet	Radius (R_ψ)	Moments (M)	Bound (B)	Lipschitz (L_z)
Haar	0.5	1	$s_{\min}^{-d/2}$	$O(s_{\min}^{-d/2-1})$
Daubechies-4	≈ 1.5	4	$O(s_{\min}^{-d/2})$	$O(s_{\min}^{-d/2-1})$
Mexican Hat	∞ (fast decay)	2	$O(s_{\min}^{-d/2})$	$O(s_{\min}^{-d/2-1})$

Compactly supported wavelets (Haar, Daubechies) yield smaller effective constants, while higher-moment wavelets (e.g., Daubechies) achieve stronger bias reduction of order $O(s^M)$.

B EXPERIMENTAL DETAILS

All the models in the experimental section are implemented using PyTorch and mostly are implemented using GPytorch (Gardner et al., 2018), trained by Adam and AdamW Optimizer on an

NVIDIA A40 GPU. The learning rate for most of the examples is taken to be 0.01 (unless mentioned otherwise) and a batch size of 128. For the Deep-GP example, we follow the doubly stochastic variational inference as proposed by (Salimbeni et al., 2019) with a zero-mean.

Unless specifically stated, we have normalised the input data for training and initalized our model with length-scale $l = 0.1$ and $\sigma^2 = 0.1$ kernel variance for TIMIT dataset.

B.1 EVALUATION METRICS

We evaluate our models using Root Mean Squared Error (RMSE) Let the dataset be denoted as $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$ for training and $\mathcal{D}^* = \{(\mathbf{x}_n, y_n)\}_{n=1}^{N^*}$. We consider a model f trained on \mathcal{D} and evaluated using the following criteria. Note here $\mathbf{y} = \{y_n\}_{n=1}^N$ is the ground truth and model predictions $\mathbf{f} = f(\mathbf{X})$ where $\mathbf{X} = \{\mathbf{X}_n\}_{n=1}^N$.

Root Mean Squared Error (RMSE). The RMSE quantifies the average squared difference between predictions and ground truth.

$$\mathcal{L}_{\text{RMSE}}(\mathbf{f}; \mathcal{D}) = \sqrt{\mathbb{E}_{(\mathbf{x}, y)} [\|y - \mathbb{E}[f(\mathbf{x}) | \mathcal{D}]\|^2]} \quad (\text{B.1})$$

Empirically estimated as

$$\mathcal{L}_{\text{RMSE}}(\mathbf{f}; \mathcal{D}) \approx \sqrt{\frac{1}{N} \sum_{n=1}^N \|y_n - \hat{f}(\mathbf{x}_n)\|^2}, \quad (\text{B.2})$$

where $\hat{f}(\mathbf{x}_n)$ is the predictive mean at input \mathbf{x}_n .

C ADDITIONAL EXPERIMENTAL DETAILS

C.1 SYNTHETIC DATASET

Effect of feature size. Figure 3 reports the convergence behavior of RWF-GP as the number of features D increases. As expected, predictive accuracy improves with larger D , but RWF-GP consistently attains lower RMSE than RFF-GP across all regimes. Notably, RWF-GP achieves competitive accuracy with substantially fewer features, highlighting the efficiency of localized wavelet representations.

Baseline details. We evaluate all models on a dataset consisting of $N = 4200$ training points and $N_{\text{test}} = 1800$ held-out test points. We utilized the Adam optimizer with a learning rate of 0.01. The baseline kernel configurations were chosen as follows: the **Exact GP**, **RFF-GP**, and **SVGP** employed a stationary Squared Exponential (RBF) kernel; and the **Adaptive-RKHS** baseline employed a non-stationary convolution kernel. Our proposed **RWF-GP** utilized a Mexican Hat mother wavelet, demonstrating its ability to capture sharp transitions without the stationarity assumptions inherent in the RBF and Matérn baselines.

Comparison with the Non-Stationary Covariance GP. For completeness, we also report the performance of the classical non-stationary covariance model of Paciorek & Schervish (2003) on the multi-step function. Results are shown in Table 5.

Relation to prior wavelet-based methods. Prior wavelet kernels (Zhang et al., 2004; Yger & Rakotomamonjy, 2011) use fixed or learned wavelet dictionaries within deterministic SVM formulations, while Kereta et al. (Kereta et al., 2019) study frame-based discretizations that require spectral decomposition of the empirical kernel matrix. In contrast, RWF constructs an explicit randomized feature map from a continuous wavelet integral representation, yielding an unbiased Monte Carlo estimator with linear-time primal inference in the number of samples.

C.2 TIMIT SPEECH SIGNAL

Dataset and Preprocessing. We use the TIMIT corpus (630 speakers, 6300 utterances, 16 kHz). For each utterance, 80-dimensional features are extracted (25 ms window, 10 ms hop, pre-emphasis,

Table 5: Performance comparison of GP baselines on the multi-step function over five runs (mean \pm std; lower is better). Bold indicates the best result and underline indicates the second best. Here, NS-GP is Non-stationary Covariance GP (Paciorek & Schervish, 2003). Kereta et al. (Kereta et al., 2019) propose a deterministic frame-based discretization that does not yield predictive uncertainty; therefore, CRPS and NLL are not applicable (-).

Method	RMSE	CRPS	NLL	Time
Exact	0.190 \pm 0.091	0.215 \pm 0.030	0.042 \pm 0.012	12
SVGP	0.231 \pm 0.014	0.392 \pm 0.025	0.123 \pm 0.018	15
RFF	0.246 \pm 0.142	0.238 \pm 0.041	0.118 \pm 0.181	<u>11</u>
DRF	0.190 \pm 0.120	0.205 \pm 0.032	-0.018 \pm 0.216	18
NS-GP	0.104 \pm 0.010	0.168 \pm 0.007	-1.030 \pm 0.04	12
DGP	0.162 \pm 0.110	0.187 \pm 0.028	-0.268 \pm 0.211	20
SM	0.210 \pm 0.085	0.201 \pm 0.030	0.220 \pm 0.180	17
Kereta	0.822 \pm 0.009	-	-	15
IDD	0.107 \pm 0.050	0.143 \pm 0.020	-0.820 \pm 0.080	17
A-RKHS	0.095 \pm 0.045	0.131 \pm 0.018	-1.210 \pm 0.075	<u>11</u>
RWF (Ours)	0.071\pm0.011	0.112\pm0.010	-1.879\pm0.061	9




Figure 3: RMSE vs. number of features D for RWF-GP (Mexican-hat) and RFF-GP on the multi-step function.

CMVN). Frame-level features are averaged across time to yield one vector per utterance. As regression targets, we use either the mean energy of a chosen Mel band (`mel_bin_k_mean`) or the mean of a PCA component of the spectrogram (`mel_pca_k`). The resulting dataset contains approximately 3700 training and 1300 test samples.

RWF Configuration. We employ complex Morlet wavelets for time–frequency localization. Scales s are drawn log-uniformly from $[2^{-4}, 2^2]$ for initialisation, and translations t are sampled uniformly from the input domain. Features are $\phi_i(x) = D^{-1/2} \psi_{s_i, t_i}(x)$. Hyperparameters (s_{\min}, s_{\max}), and noise variance, are tuned. Regularization includes (i) clipping extreme scales during warm-up and (ii) ridge penalty $\lambda \|w\|_2^2$ with $\lambda = 10^{-4}$ on Bayesian linear weights. (a) clipping extreme scales during warm-up, (b) ridge penalty $\lambda \|w\|_2^2$ (with $\lambda = 10^{-4}$) on the Bayesian linear weights’ MAP objective surrogate used for hyperparameter inner loops. **Wavelet family.** Unless otherwise specified, we employ *Morlet* and *Mexican Hat* wavelets as the mother wavelets for constructing random wavelet features.