

Breaking the waves: asymmetric random periodic features for low-bitrate kernel machines

Vincent Schellekens* and Laurent Jacques*

Abstract

Many signal processing and machine learning applications are built from evaluating a kernel on pairs of signals, *e.g.*, to assess the similarity of an incoming query to a database of known signals. This nonlinear evaluation can be simplified to a linear inner product of the random Fourier features of those signals: random projections followed by a periodic map, the complex exponential. It is known that a simple quantization of those features (corresponding to replacing the complex exponential by a different periodic map that takes binary values, which is appealing for their transmission and storage), distorts the approximated kernel, which may be undesirable in practice. Our take-home message is that when the features of only one of the two signals are quantized, the original kernel is recovered without distortion; its practical interest appears in several cases where the kernel evaluations are asymmetric by nature, such as a client-server scheme.

Concretely, we introduce the general framework of *asymmetric random periodic features*, where the two signals of interest are observed through random periodic features—random projections followed by a general periodic map, which is allowed to be different for both signals. We derive the influence of those periodic maps on the approximated kernel, and prove uniform probabilistic error bounds holding for all pair of signals picked in an infinite low-complexity set. Interestingly, our results allow the periodic maps to be discontinuous, thanks to a new mathematical tool, *i.e.*, the mean Lipschitz smoothness. We then apply this generic framework to semi-quantized kernel machines (where only one of the signals has quantized features and the other has classical random Fourier features), for which we show theoretically that the approximated kernel remains unchanged (with the associated error bound), and confirm the power of the approach with numerical simulations.

1 Introduction

Rather than to directly process high-dimensional signals, it is often more efficient to first summarize them to their main *features*. This assumes that these capture essential information for the considered processing, such as the proximity of any pair of signals. Mathematically, the signal summarization is modeled by a feature map φ from the signal space Σ to the feature (or embedding) space \mathcal{E} . This map φ transforms the representation of signals while encoding some aspects of their geometry; loosely speaking, this can be written as $\mathcal{D}_{\mathcal{E}}(\varphi(\mathbf{x}), \varphi(\mathbf{y})) \approx \mathcal{D}_{\Sigma}(\mathbf{x}, \mathbf{y})$ for any pair of signals $\mathbf{x}, \mathbf{y} \in \Sigma$, where \mathcal{D}_{Σ} is the preserved geometric quantity (such as an inner product or a distance), and $\mathcal{D}_{\mathcal{E}}$ is an evaluation procedure acting only on the signal features. This approach is useful whenever the features $\varphi(\mathbf{x})$ are easier to process with respect to some critical computational resource (*e.g.*, memory usage, computing time)—often at the price of an approximation error (as suggested by the approximation symbol $\mathcal{D}_{\mathcal{E}} \approx \mathcal{D}_{\Sigma}$ above). The map φ is most often than not a randomized function (drawn from a distribution). There are essentially three main ways to save computational resources with features: (i) leveraging *dimensionality reduction* (*i.e.*, $\varphi(\mathbf{x})$ is encoded by much less coefficients than the dimension of \mathbf{x}), such as in compressive

*E-mail: {vincent.schellekens, laurent.jacques}@uclouvain.be. ISPGROUP, ELEN/ICTEAM, UCLouvain (UCL), B1348 Louvain-la-Neuve, Belgium. VS and LJ are funded by Belgian National Science Foundation (F.R.S.-FNRS).

sensing techniques [1], where typically the “embedding” φ is linear and \mathcal{D}_Σ and $\mathcal{D}_\mathcal{E}$ are the Euclidean distances; (ii) using mappings that *linearize* the evaluation of an otherwise nonlinear quantity, such as random Fourier features (RFF) [2], where \mathcal{D}_Σ is a kernel κ and $\mathcal{D}_\mathcal{E}$ is simply the inner product; and (iii) *quantizing features*, where $\varphi(\mathbf{x})$ produces a quantized output that can be encoded with a highly reduced bitrate compared to the initial signal, allowing reducing the memory and/or transmission load, as well as paving the way for hardware-based procedures. This quantization objective is often combined with either (i), as in quantized compressive sensing [3–5], or (ii), as in one-bit universal embeddings [6]. As made clear below, our contributions also target the combination of (ii) with (iii).

In all those applications, it is almost always assumed that the available features for the two signals $\mathbf{x}, \mathbf{y} \in \Sigma$ come from the *same* feature map φ (we say the features are *symmetric*). However, we can legitimately wonder if removing this assumption, *i.e.*, accessing the signals through features $\varphi(\mathbf{x})$ and $\psi(\mathbf{y})$, where we have the freedom to set $\varphi \neq \psi$, can further reduce specific computational resources. The practical interest of this relaxation—that we call *asymmetric features*—arises when the two signals come from different sources (*i.e.*, when the setting is intrinsically asymmetric): for example, those sources might have different resources (such as memory or power) at their disposal, and will therefore benefit differently from techniques (i)-(iii).

In this work, we are interested in asymmetric features for linearizing kernel estimations, as explained in (ii). Anticipating over the detailed description of Sec. 2, we work with *random periodic features*, $\varphi(\mathbf{x}) := f(\Omega^T \mathbf{x} + \xi)$ and $\psi(\mathbf{y}) := g(\Omega^T \mathbf{y} + \xi)$ with Ω a random projection matrix, ξ a random dither, and f, g two periodic functions. We thus generalize the context of random Fourier features [2], where $f(\cdot) = g(\cdot) = \exp(i\cdot)$ (the complex exponential), by “*breaking the waves*” with possibly discontinuous, distinct functions f and g (as described in Sec. 4). We show how those features can be used to approximate shift-invariant kernels κ , *i.e.*, $\langle \varphi(\mathbf{x}), \psi(\mathbf{y}) \rangle \approx \kappa(\mathbf{x}, \mathbf{y})$, in expectation over the random quantities Ω, ξ . Our motivating use-case is to combine this approach with harsh quantization of some features, objective (iii), as we explain in the next paragraph. However, all our developments are generic, and of interest for any machine learning algorithm that processes or takes decisions from the local geometry of data.

Semi-binary kernel machines as a motivating application: If one of the periodic functions incorporates the quantization of the feature vector (say, the one-bit universal quantization, or square wave function $q : \mathbb{R} \rightarrow \{0, 1\}$ [7]; see Fig. 2), then that feature vector can be stored or transmitted (or both) much more efficiently than the usual (infinite-precision) random Fourier features. Consider for example a machine learning context, where a kernel method [8] such as a Support Vector Machine (SVM) [9] has been trained in advance on some dataset $X = \{\mathbf{x}_i\}_{i=1}^n \subset \Sigma$. To actually use this model for prediction on a new signal $\mathbf{x}' \in \Sigma$, the physical device that records it must either communicate with a server where the inference is performed remotely (Fig. 1a), or implement this model directly (Fig. 1b); in either case, this is an expensive operation whenever this device is under tight computational resources constraints, and quantization of feature vectors is potentially very helpful.

In the first scenario (inference done remotely on a server), we might quantize the feature vector of the query signal, $\varphi(\mathbf{x}') \in \{0, 1\}^m$ (but not the feature vectors of the dataset $\psi(\mathbf{x}_i)$). This allows to heavily reduce the bitrate when communicating this vector to the server, and even paves way for computing those features directly in hardware, *e.g.*, relying on voltage-controlled oscillators [10]. In the second context, we could conversely binarize the feature vectors of the dataset so that $\psi(\mathbf{x}_i) \in \{0, 1\}^m$ for all $\mathbf{x}_i \in X$, but not the incoming query vector $\varphi(\mathbf{x}')$ (it is even possible to encode only a subset of the dataset features for models that only need to access some entries, such as SVM with the support vectors). The advantage here is that the memory needed to store the model is heavily reduced, with additional computational benefits coming from the embedded processing of binary values. This idea has received significant attention in the literature, *e.g.*, following [11] for nearest-neighbor search.




Figure 1: Two motivating applications of our results (from Sec. 6, in green): combining one-bit universal features with usual random Fourier features yields the same desired kernel κ . (a) A “client” device records a “query” signal \mathbf{x}' , and transmits its quantized features $\mathbf{z}_q(\mathbf{x}')$, encoded efficiently as only m bits, to a “server” device that can evaluate the kernel similarity with the rest of a dataset from their usual n full-precision RFF $\{\mathbf{z}(\mathbf{x}_i)\}_{i=1}^n$. (b) A lightweight device implements a kernel method with very low memory requirements, only having to store $\{\mathbf{z}_q(\mathbf{x}_i)\}_{i=1}^n$ the n one-bit feature vectors of the dataset instead of the full-precision ones, provided the usual RFF $\mathbf{z}(\mathbf{x}')$ are used for the incoming query vectors.

For both of those examples, the main question that we seek to answer is to quantify the loss of accuracy (induced by quantization) as a function of the feature vector length m . More precisely, our goal is to obtain (probabilistic) guarantees on the decay of the kernel approximation error as a function of m , that hold uniformly for any pairwise comparisons of signals taken an infinite (but compact) set Σ . In this case, the main challenge lies in dealing with the discontinuous nature of the quantization operation—handling discontinuities is thus one of the key features of this paper.

1.1 Related work

Before detailing the elements of our approach, we find useful to mention a few related works, showing how they inspired us, and stressing their connections and differences with our contributions.

Quantization of (symmetric) random Fourier features: The construction of the general random periodic features considered in this work is instantiated in Sec. 6 to the case where the corresponding periodic map is the universal quantizer (or square wave function). This approach was introduced in [7, 12] as a binary map preserving local distances (*i.e.*, up to a given radius), the universal quantized embedding. Those features have subsequently been used for kernel methods in [6], which is similar to the framework we propose but where not one but both signal features being compared are quantized in a symmetric fashion, which distorts the kernel to be recovered. This line of work was further generalized in [13], where uniform guarantees are derived for generic periodic function (possibly discontinuous) instead of the one-bit universal quantizer, holding on infinite signal sets. This defines the random periodic features approach (see Sec. 2 for details) that we also consider; we provide an in-depth description of how our results relate to (and complement) those from [13] in Sec. 5.

Back to the particular problem of quantizing random Fourier features, another line of work [14] shows that a specific stochastic quantization hardly harms the generalization performances of RFF-based algorithms. The ultimate objective of this last work is, however, fairly different from ours: the authors seek to reduce the memory requirements *during training* by performing a more sophisticated quantization, and then use the full-precision RFF for the subsequent inference stage; on the other hand, our objective is to provide a simple quantization scheme to reduce the resources *during the inference stage*, without concerns for how the training was performed.

Asymmetric features and quantizations: The possibility to use asymmetric features has been explored for linear embeddings in [15], as an additional degree of freedom to minimize (in a data-dependent fashion) the average error of the distance estimation. In [16], weighted universal embeddings are used for distance estimation, where the weights depend upon one of the two signals (which makes the scheme asymmetric) to decrease the error on this estimation. Closer to our context, in [17], it is experimentally shown for a broad set of feature maps (such as Locality Sensitive Hashing, universal embeddings, and several variants of PCA) that quantizing the features of the dataset but not of the query (as in scenario Fig. 1b) significantly improves the performances compared to quantizing both features. Similarly, the authors of [18,19] recently considered linear random projections (with the same matrix) of two signals that have been quantized with different quantization levels.

Compressive learning: In a nutshell, compressive learning [20,21] aims at estimating the parameters of a distribution \mathcal{P} , or the parameters of a parametric distribution approximating it, from the sketching of an entire dataset of n signals $X = \{\mathbf{x}_i\}_{i=1}^n$ generated by \mathcal{P} , *i.e.*, such that $\mathbf{x}_i \sim_{\text{i.i.d.}} \mathcal{P}$. Given a random projection matrix $\mathbf{\Omega}$, the sketch $\mathbf{s}_X := \frac{1}{n} \sum_{i=1}^n \mathbf{z}(\mathbf{x}_i)$ is computed from the *pooling* (averaging) of the random Fourier features $\mathbf{z} : \mathbf{u} \rightarrow \exp(i\mathbf{\Omega}^\top \mathbf{u})$ (the exponential being computed componentwise onto vectors) of each dataset signal. For large value of n , this sketch estimates the characteristic function $\mathcal{A}(\mathcal{P}; \mathbf{\Omega}) := \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} \exp(i\mathbf{\Omega}^\top \mathbf{x})$ of \mathcal{P} over the “frequencies” supported by the rows of $\mathbf{\Omega}$. Therefore, under appropriate conditions, one can formulate an inverse problem aiming to estimate the parameters of \mathcal{P} by matching the characteristic function of a probing distribution $\hat{\mathcal{P}}$ (estimated over $\mathbf{\Omega}$) from $\mathbf{s}_X \approx \mathcal{A}(\mathcal{P}; \mathbf{\Omega})$. We considered in [22] the possibility to replace the random Fourier features used to build \mathbf{s}_X with a general (dithered) periodic function f , such as the universal quantizer, thus computing $\mathbf{s}'_X = \frac{1}{n} \sum_{i=1}^n \mathbf{z}_f(\mathbf{x}_i) = \frac{1}{n} \sum_{i=1}^n f(\mathbf{\Omega}^\top \mathbf{x}_i + \boldsymbol{\xi})$ with a random dither $\boldsymbol{\xi}$. While the dataset sketch is altered (*e.g.*, quantized with the universal quantizer), we showed that the estimation of the distribution parameters from the observed sketch is still accurate if we use the RFF (for the probing distribution), as if the sketch was not quantized, hence leading to an asymmetric scheme between the dataset sketching and the estimation procedure.

1.2 Paper organization

We provide in Sec. 2 several preliminary elements as well as important concepts of the relevant literature: random Fourier features and their (possibly quantized) extension to any periodic nonlinearity. We then start by analyzing how the kernel approached by asymmetric random periodic features behaves in expectation (in the asymptotic case), which is proved in Sec. 3. Our main results come in Sec. 4, where we prove uniform error bounds of the kernel approximation for infinite signal sets. In order to do so, we introduce a new tool, the mean Lipschitz smoothness property. Sec. 5 relates our approach to the context of geometry-preserving embedding (or coding) developed in [13], solving in the same time an error in the proof of one of their results. Next, we apply our general results of Sec. 4 to the semi-quantized setting motivated above in Sec. 6, and illustrate with numerical experiments in Sec. 7, before concluding in Sec. 8.

1.3 Notations

Vectors and matrices are denoted by bold symbols. The unit imaginary number is noted $i = \sqrt{-1}$. The real part, the imaginary part, and the complex conjugation of $a \in \mathbb{C}$ read $\Re(a)$, $\Im(a)$, and a^* , respectively. We will often consider bounded 2π -periodic functions $f, g : \mathbb{R} \rightarrow \mathbb{C}$ for which the 2-norm and the infinity norm read $\|f\|^2 = \frac{1}{2\pi} \int_0^{2\pi} |f(t)|^2 dt$ and $\|f\|_\infty := \sup_{t \in [0, 2\pi]} |f(t)|$, respectively, and the inner product of f and g is $\langle f, g \rangle = \frac{1}{2\pi} \int_0^{2\pi} f(t)g^*(t)dt$. For brevity and clarity, we will sometimes refer to a function using the “dot” notation, *e.g.*, $\exp(i\cdot)$ for the function $t \mapsto \exp(it) \in \mathbb{C}$ for $t \in \mathbb{R}$.

The ℓ_p -norm of a vector $\mathbf{u} \in \mathbb{R}^d$ reads $\|\mathbf{u}\|_p = (\sum_i |u_i|^p)^{1/p}$ for $p \geq 1$, with $\|\mathbf{u}\|_\infty = \max_i |u_i|$, and $\|\mathbf{u}\|_0 = |\text{supp } \mathbf{u}| = |\{i : u_i \neq 0\}|$. The unit ℓ_p -ball ($p \geq 1$) in dimension d is noted $\mathbb{B}_p^d := \{\mathbf{u} \in \mathbb{R}^d \mid \|\mathbf{u}\|_p \leq 1\}$, with the shorthand $\mathbb{B}^d = \mathbb{B}_2^d$. The cardinality of a finite set \mathcal{S} is $|\mathcal{S}|$, the Minkowski sum of two sets \mathcal{A} and \mathcal{B} is $\mathcal{A} + \mathcal{B} = \{a + b : a \in \mathcal{A}, b \in \mathcal{B}\}$, the index set in \mathbb{R}^d is $[d] := \{1, \dots, d\}$ for $d \in \mathbb{N}$, the identity matrix in \mathbb{R}^d is $\mathbf{I}_d \in \mathbb{R}^{d \times d}$, and the Kronecker delta $\delta_{k,k'}$ is defined as $\delta_{k,k'} = 1$ if $k = k'$ and $\delta_{k,k'} = 0$ otherwise. By abuse of notation, evaluating a scalar function $f : \mathbb{R} \rightarrow \mathbb{C}$ on a vector $\mathbf{u} \in \mathbb{R}^m$ means applying this function componentwise, *i.e.*, $f(\mathbf{u}) \in \mathbb{C}^m$ with $(f(\mathbf{u}))_j = f(u_j)$.

We use the convention where the (d -dimensional) Fourier transform of a function $f : \mathbb{R}^d \rightarrow \mathbb{C}$ reads $\hat{f}(\boldsymbol{\omega}) = (\mathcal{F}f)(\boldsymbol{\omega}) := \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} e^{-i\mathbf{u}^\top \boldsymbol{\omega}} f(\mathbf{u}) d\mathbf{u}$, with inverse $(\mathcal{F}^{-1}\hat{f})(\mathbf{u}) := \int_{\mathbb{R}^d} e^{i\mathbf{u}^\top \boldsymbol{\omega}} \hat{f}(\boldsymbol{\omega}) d\boldsymbol{\omega}$. The same convention is used for the Fourier transform of finite measures on \mathbb{R}^d . The notation $\sim \mathcal{P}$ denotes that a random variable, vector or function is distributed according to the distribution \mathcal{P} . The uniform distribution on a set \mathcal{A} is noted $\mathcal{U}(\mathcal{A})$, and ‘‘i.i.d.’’ means ‘‘identically and independently’’ distributed. In all our developments, except if specified differently, $C, C', \dots, c, c', \dots > 0$ denote *universal* constants whose value may change from one instance to the other.

2 Preliminaries

We introduce here several fundamental concepts supporting our approach. We first precise the kind of signal space we consider, as well as how signals are compared through a kernel, before to briefly explain the principles sustaining the definition of random Fourier features (RFF). Next, we generalize RFF to any random periodic features for a family of bounded periodic functions.

2.1 Signals and kernels

In this work, we focus on signals belonging to a bounded signal space $\Sigma \subset \mathbb{R}^d$ having finite Kolmogorov η -entropy $\mathcal{H}_\eta(\Sigma)$ for any radius $\eta > 0$ [23]. This entropy, defined as $\mathcal{H}_\eta(\Sigma) := \log \mathcal{C}_\eta(\Sigma)$, is related to the covering number $\mathcal{C}_\eta(\Sigma)$ of Σ , the cardinality of the smallest finite subset of Σ that *covers* it with balls of radius η . Using the Minkowski sum, this means that

$$\mathcal{C}_\eta(\Sigma) := \min\{|\mathcal{S}| : \mathcal{S} \subset \Sigma \subset \mathcal{S} + \eta\mathbb{B}_2^d\},$$

which is finite for any compact set Σ .

The Kolmogorov entropy measures the intrinsic dimension of Σ in \mathbb{R}^d . In particular, $\mathcal{H}_\eta(\mathcal{V} \cap \mathbb{B}_2^d) \leq Cd' \log(1 + 1/\eta)$ for any subspace $\mathcal{V} \subset \mathbb{R}^d$ of dimension $d' \leq d$ [24], and the set of s -sparse vectors $\Sigma_s := \{\mathbf{x} \in \mathbb{R}^d, \|\mathbf{x}\|_0 \leq s\}$ restricted to the unit ball has entropy bounded by $\mathcal{H}_\eta(\Sigma_s \cap \mathbb{B}_2^d) \leq C \cdot s \log(d/s) \log(1 + 1/\eta)$, see for example [25]. Other bounds exist for, *e.g.*, the set of bounded group sparse signals [26], bounded low-rank matrices [27], or for specific low-dimensional manifolds [28].

At the heart of our study is the comparison of two signals through a *kernel*, *i.e.*, a function over pairs of signals $\kappa : \Sigma \times \Sigma \rightarrow \mathbb{C}$ (in the machine learning literature, kernels are usually real-valued). Typically, invoking the so-called ‘‘kernel trick’’ [9], κ represents the inner product between the input signals \mathbf{x}, \mathbf{x}' when they are mapped in some *implicit* feature space by an appropriate map $\phi : \Sigma \rightarrow \mathbb{H}$: $\kappa(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathbb{H}}$ for some Hilbert space \mathbb{H} . By definition of the inner product, the kernel κ must then necessarily be conjugate symmetric ($\kappa(\mathbf{x}, \mathbf{x}') = \kappa^*(\mathbf{x}', \mathbf{x})$), and positive definite (p.d.), *i.e.*, for any number n , $\sum_{i,j=1}^n c_i c_j^* \kappa(\mathbf{x}_i, \mathbf{x}_j) \geq 0$ for all $\mathbf{x}_1, \dots, \mathbf{x}_n \in \Sigma$ and $c_1, \dots, c_n \in \mathbb{C}$.

2.2 Random Fourier features

Random Fourier features (RFF) are implicitly built on *Bochner’s theorem* [29]. This theorem states that a shift-invariant continuous kernel $\kappa(\mathbf{x}, \mathbf{y}) = \kappa^\Delta(\mathbf{x} - \mathbf{y})$ (for some $\kappa^\Delta : \Sigma - \Sigma \rightarrow \mathbb{C}$) is positive definite if

and only if it is the (inverse) Fourier transform of a nonnegative finite measure Λ , *i.e.*,

$$\kappa \text{ positive definite} \quad \Leftrightarrow \quad \kappa^\Delta(\mathbf{u}) = (\mathcal{F}^{-1}\Lambda)(\mathbf{u}) = \int_{\mathbb{R}^d} e^{i\boldsymbol{\omega}^\top \mathbf{u}} d\Lambda(\boldsymbol{\omega}). \quad (1)$$

In particular, assuming *w.l.o.g.* the normalization $\kappa(\mathbf{x}, \mathbf{x}) = \kappa^\Delta(\mathbf{0}) = 1$, Λ is a probability distribution over \mathbb{R}^d , and the kernel can be written $\kappa^\Delta(\mathbf{u}) = \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} e^{i\boldsymbol{\omega}^\top \mathbf{u}}$. The key idea of random Fourier features [2] is thus to construct low-dimensional features $\mathbf{z}(\mathbf{x}), \mathbf{z}(\mathbf{y})$ whose inner product approximates the kernel $\kappa(\mathbf{x}, \mathbf{y})$ by Monte Carlo sampling of this expectation.

Definition 1 (Random Fourier features). *Let $\kappa(\mathbf{x}, \mathbf{y}) = \kappa^\Delta(\mathbf{x} - \mathbf{y})$ be a shift-invariant p.d. kernel, normalized such that $\kappa^\Delta(\mathbf{0}) = 1$, with Fourier transform $\Lambda = \mathcal{F}\kappa^\Delta$. Given a target dimension m , the associated “complex” random Fourier features are*

$$\mathbf{z}(\mathbf{x}) := \frac{1}{\sqrt{m}} \exp(i(\boldsymbol{\Omega}^\top \mathbf{x} + \boldsymbol{\xi})) \quad \in \mathbb{C}^m, \quad (2)$$

with random projections (or “frequencies”) $\boldsymbol{\Omega} := (\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_m) \in \mathbb{R}^{d \times m}$ generated as $\boldsymbol{\Omega} \sim \Lambda^m$, *i.e.*, with $\boldsymbol{\omega}_j \sim_{\text{i.i.d.}} \Lambda$ for $j \in [m]$, and a random dither $\boldsymbol{\xi} \in \mathbb{R}^m$ generated as $\boldsymbol{\xi} \sim \mathcal{U}^m([0, 2\pi])$, *i.e.*, with $\xi_j \sim_{\text{i.i.d.}} \mathcal{U}([0, 2\pi])$ for $j \in [m]$. We also define the “real” random Fourier features $\mathbf{z}_{\cos}(\mathbf{x})$ as $\Re[\mathbf{z}(\mathbf{x})]$, the real part of those features:

$$\mathbf{z}_{\cos}(\mathbf{x}) := \frac{1}{\sqrt{m}} \cos(\boldsymbol{\Omega}^\top \mathbf{x} + \boldsymbol{\xi}) \quad \in \mathbb{R}^m. \quad (3)$$

Remark. *The dither $\boldsymbol{\xi}$ was initially introduced in [2] when only the real RFF $\mathbf{z}_{\cos}(\mathbf{x})$ were used; in the (more widely used) complex case $\boldsymbol{\xi}$ is not necessary (see [30] for an in-depth comparison of the “real” versus “complex” random Fourier features). We still included it in this definition for the sake of consistency with Def. 2 below.*

By direct application of Bochner’s theorem, the inner product of RFF indeed approaches (in expectation over the draw of the frequencies $\boldsymbol{\Omega}$) the target kernel: $\mathbb{E}\langle \mathbf{z}(\mathbf{x}), \mathbf{z}(\mathbf{y}) \rangle = \kappa(\mathbf{x}, \mathbf{y})$. Moreover, for a finite feature dimension m , the error of the kernel approximation $\widehat{\kappa}(\mathbf{x}, \mathbf{y}) := \langle \mathbf{z}(\mathbf{x}), \mathbf{z}(\mathbf{y}) \rangle$ can be uniformly bounded (*i.e.*, bound the absolute error $|\widehat{\kappa}(\mathbf{x}, \mathbf{y}) - \kappa(\mathbf{x}, \mathbf{y})|$ for all values \mathbf{x}, \mathbf{y} in Σ), with high probability on the draw of $\boldsymbol{\Omega}$ (we work with different normalization choices, so the result we present here differs slightly from the initial bound [2, Claim 1]). Finer bounds can be found, among others, in [30, 31].

Proposition 1 (Uniform kernel approximation error for RFF). *Let Σ be a compact set, and $\mathbf{z}(\mathbf{x})$ be the RFF defined above. Assume that there exists an associated constant C_Λ , such that*

$$\mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} |\boldsymbol{\omega}^\top \mathbf{a}| \leq C_\Lambda \|\mathbf{a}\|_2, \quad \forall \mathbf{a} \in \mathbb{R}^d. \quad (4)$$

Provided that, for $\epsilon > 0$,

$$m \geq C\epsilon^{-2} \mathcal{H}_{c\epsilon/C_\Lambda}(\Sigma),$$

the kernel approximation $\widehat{\kappa}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{z}(\mathbf{x}), \mathbf{z}(\mathbf{y}) \rangle$ has error uniformly bounded by

$$|\widehat{\kappa}(\mathbf{x}, \mathbf{y}) - \kappa(\mathbf{x}, \mathbf{y})| \leq \epsilon, \quad \forall \mathbf{x}, \mathbf{y} \in \Sigma,$$

with probability exceeding $1 - C'e^{-c'm\epsilon^2}$.

Proof. This version of the RFF approximation error is obtained as a particular case of our Prop. 5; see [2] for the initial result. \square

The constant C_Λ defined in (4) characterizes the smoothness of the kernel (if the kernel is smoother, it exhibits less high-frequency content, and C_Λ will be lower). In most of the RFF literature, this constant is bounded by the Cauchy-Schwarz inequality as $C_\Lambda = \mathbb{E}_{\omega \sim \Lambda} \|\omega\|_2$. Then, one can (as done in [2]) further bound $\mathbb{E}_{\omega \sim \Lambda} \|\omega\|_2 \leq \sigma_\Lambda$ where σ_Λ^2 is the second moment of Λ , equivalent to the kernel curvature at the origin, *i.e.*, $C_\Lambda^2 \leq \sigma_\Lambda^2 := \mathbb{E}_{\omega \sim \Lambda} \|\omega\|_2^2 = \nabla^2 \kappa^\Delta|_{\mathbf{u}=\mathbf{0}}$, with ∇^2 the Laplacian operator. However, for specific distributions, using the Cauchy-Schwarz inequality in high dimension leads to a loose bound of C_Λ . For example, if the covariance matrix of ω is upper bounded by $\tilde{\sigma}_\Lambda^2 \mathbf{I}_d$ for some $\tilde{\sigma}_\Lambda^2 > 0$ (if, *e.g.*, $\mathcal{F}^{-1}\Lambda$ is the Gaussian RBF (“radial basis function”) kernel with radius $1/\tilde{\sigma}_\Lambda^2$, Λ is isotropic [32], or if each component of ω are i.i.d. with variance bounded by $\tilde{\sigma}_\Lambda^2$) then

$$(\mathbb{E}_{\omega \sim \Lambda} |\omega^\top \mathbf{a}|)^2 \leq \mathbb{E}_{\omega \sim \Lambda} |\omega^\top \mathbf{a}|^2 = \mathbf{a}^\top (\mathbb{E}_{\omega \sim \Lambda} \omega \omega^\top) \mathbf{a} \leq \mathbf{a}^\top \cdot \tilde{\sigma}_\Lambda^2 \mathbf{I}_d \cdot \mathbf{a} = \tilde{\sigma}_\Lambda^2 \|\mathbf{a}\|_2^2.$$

In this case we obtain $C_\Lambda = \tilde{\sigma}_\Lambda$, while Cauchy-Schwarz gives $\sqrt{d} \cdot \tilde{\sigma}_\Lambda$, hence overestimating the constant by a factor \sqrt{d} .

Example 1. Consider the simple case where the signals of interest have an ℓ_2 -norm smaller than 1 and lie in a union of S subspaces of \mathbb{R}^d , each with dimension s . This signal space model encompasses, for instance, $\Sigma = \mathbb{B}_2^d$ (that is, where $S = 1$ and $s = d$), the set of bounded s -sparse signals in \mathbb{R}^d for which $S = \binom{d}{s} \leq (\frac{ed}{s})^s$ and each subspace (one per fixed sparse signal support) has dimension s , or more advanced models with structured sparsity [26, 33]. For such a model, the Kolmogorov entropy is bounded by $Cs \cdot \log(\frac{1}{\eta}) \leq \mathcal{H}_\eta(\Sigma) \leq C's \cdot \log(1 + \frac{2}{\eta}) + \log S$ (see, *e.g.*, [34, Lemma 10]). Assume that we target the usual Gaussian kernel with unit bandwidth, hence $C_\Lambda = 1$. In this case, the RFF kernel approximation error is uniformly bounded over Σ , with high probability, provided that the number of features satisfies $m \geq C\epsilon^{-2}(s \log(\frac{1}{c\epsilon}) + \log S)$. For instance, for bounded s -sparse signals we need $m \geq C\epsilon^{-2}s \log(\frac{ed}{cs\epsilon})$.

2.3 Random periodic features

A crucial generalization to RFF has been proposed in [13], where the complex exponential is replaced by a *generic periodic function* f . We refer to this approach as *random periodic features* (RPF). Without loss of generality, we make the following normalization assumptions throughout this work: f has period given by 2π , is centered (zero mean), and takes (absolute) values bounded by one. We note this compactly as

$$f \in \text{PF}, \quad \text{with} \quad \text{PF} := \{f : \mathbb{R} \rightarrow \mathbb{C} \mid f \text{ is } 2\pi\text{-periodic, } \int_0^{2\pi} f(t)dt = 0, \|f\|_\infty \leq 1\}.$$

Functions of PF can be expressed as a Fourier series of the following form

$$f(t) = \sum_{k \in \mathbb{Z}} F_k e^{ikt}, \quad \text{where} \quad F_k := \frac{1}{2\pi} \int_0^{2\pi} f(t) e^{-ikt} dt. \quad (5)$$

Note that $f \in \text{PF}$ implies $F_0 = 0$ (because f is centered) and $|F_k| \leq 1$ (because f is bounded).

Definition 2 (Random periodic features). *Let f be a generic periodic function, normalized such that $f \in \text{PF}$, and Λ a probability distribution on \mathbb{R}^d . Given a target dimension m , the associated random periodic features are*

$$\mathbf{z}_f(\mathbf{x}) := \frac{1}{\sqrt{m}} f(\mathbf{\Omega}^\top \mathbf{x} + \boldsymbol{\xi}) \in \mathbb{C}^m, \quad (6)$$

with a $d \times m$ random projection matrix $\mathbf{\Omega} := (\omega_1, \dots, \omega_m) \sim \Lambda^m$, and a random dither $\boldsymbol{\xi} \sim \mathcal{U}^m([0, 2\pi])$.

Remark. As the complex exponentiation satisfies $\exp(i \cdot) \in \text{PF}$, this definition includes the classical random Fourier features, with $\mathbf{z}(\mathbf{x}) = \mathbf{z}_{\exp(i \cdot)}(\mathbf{x})$. The real RFF $\mathbf{z}_{\cos}(\mathbf{x})$ are also a particular case of this definition.

The geometry induced by such generic features can be characterized the inner product $\hat{\kappa}_{f,f}(\mathbf{x}, \mathbf{y}) := \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_f(\mathbf{y}) \rangle$. As explained by the following result (adapted from [13, Theorem 4.4]), this product is associated with a modified kernel $\kappa_{f,f}(\mathbf{x}, \mathbf{y})$ (the rationale for these notations is clarified in the next section).




Figure 2: (Left) The solid black curve represents the universal quantization function $q(t)$ (with $q \in \text{PF}$) defined in (8). Up to a shift and rescaling, this function corresponds to the *least significant bit* of a standard uniform scalar quantizer. In dashed green, we display the related integrand $I_\delta(t)$, with $\delta = 0.35$. This quantity refers to the proof of the mean smoothness (Def. 3) of q in Prop. 6 (Sec. 6). (Right) When drawing Ω from a Gaussian distribution $\Lambda = \mathcal{N}(0, \mathbf{I}_d)$, the associated RFF recover the Gaussian kernel κ (in black), but the RPF with universal quantization approximate a “distorted” kernel $\kappa_{q,q}$ (dashed red) whose almost linear behavior close to the origin is explained by (16).

Proposition 2 (Kernel from symmetric RPF). *The inner product of random periodic features (6) approaches, on average, a kernel $\kappa_{f,f}(\mathbf{x}, \mathbf{y}) := \mathbb{E}\langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_f(\mathbf{y}) \rangle$ that is shift-invariant and given by*

$$\kappa_{f,f}(\mathbf{x}, \mathbf{y}) = \sum_{k \in \mathbb{Z}} |F_k|^2 \kappa^\Delta(k(\mathbf{x} - \mathbf{y})) =: \kappa_{f,f}^\Delta(\mathbf{x} - \mathbf{y}), \quad (7)$$

where $\kappa^\Delta(\mathbf{u}) = (\mathcal{F}^{-1}\Lambda)(\mathbf{u})$ is the shift-invariant kernel associated with the distribution of Ω in the RPF.

Proof. This version is obtained as a particular case of our Prop. 3; see [13] for the initial result. \square

The modified kernel $\kappa_{f,f}$ is thus a *scale mixture* of the initial kernel κ (that is approached by the “classical” RFF), where the weight of scale k is given by $|F_k|^2$. In the non-asymptotic case, the authors of [13] show that, for all pairs of vectors taken in a finite set Σ of size N , $\widehat{\kappa}_{f,f}(\mathbf{x}, \mathbf{y})$ quickly concentrates around $\kappa_{f,f}(\mathbf{x}, \mathbf{y})$ when m is large compared to $\log N$; the deviation error scaling as $O(\sqrt{\log N/m})$ when m increases. Our result in Prop. 5 provides a uniform approximation bound valid for infinite sets.

Random periodic features were introduced as a general theoretical framework to analyze the so-called universal quantization embeddings [7]; those binary embeddings encode the local distances (*i.e.*, the distances below a given threshold) on an efficiently small number of bits. This embedding relies on the “one-bit universal quantization” given by $\mathcal{Q}_\Delta : \mathbb{R} \rightarrow \{0, 1\} : t \mapsto \mathcal{Q}_\Delta(t) = 1$ if $(2k - 1) \leq t/\Delta \leq 2k$ for any $k \in \mathbb{Z}$ and 0 otherwise. It can be interpreted as the least significant bit of a usual, plain scalar quantizer with stepsize Δ , and visualized as a square wave: see Fig. 2, left. Here, we will for convenience use q instead, its normalized equivalent in PF,

$$q(t) := \text{sign} \circ \cos(t) = \sum_{k \in \mathbb{Z}} Q_k e^{ikt}, \quad \text{with coefficients } Q_k = \begin{cases} \frac{2}{k\pi} (-1)^{(k-1)/2} & \text{if } k \text{ odd,} \\ 0 & \text{if } k \text{ even.} \end{cases} \quad (8)$$

Using the universal quantization as periodic nonlinearity is appealing because $\mathbf{z}_q(\mathbf{x}) \in \{-1, +1\}^m$, which can thus be encoded/transmitted by only m bits. However, as predicted by (7), the approximated kernel is modified, as illustrated for the Gaussian kernel Fig. 2, right. Moreover, proving uniform kernel approximation bounds (as in Prop. 1) for infinite sets Σ is specially challenging when the nonlinearity f presents discontinuities (which is the case when $f = q$, for example). In [13], the authors introduced a formalism (the T -part Lipschitz functions) to deal with this problem and to obtain uniform approximation bounds on infinite signal sets for the universal embeddings. As we explain in Sec. 5, the proof relying on this approach is however wrong, which motivates us to introduce another tool, the mean Lipschitz smoothness, to deal with discontinuous maps.

3 Expected kernel (asymptotic case)

Following the considerations of the Introduction, let us now consider the *asymmetric features* setting where a pair of signals of interest, $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, are available only through their random periodic features, $\mathbf{z}_f(\mathbf{x})$ and $\mathbf{z}_g(\mathbf{y})$, as defined in (6). Those features are allowed to result from different periodic maps $f, g \in \text{PF}$, but the preceding projection $\mathbf{\Omega}$ and dithering $\boldsymbol{\xi}$ are kept identical.

In this section, we characterize the properties of the *expected kernel* yielded by the expectation, over the draw of $\mathbf{\Omega}$ and $\boldsymbol{\xi}$, of the following ‘‘asymmetric’’ inner product:

$$\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) := \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle. \quad (9)$$

This asymmetric RPF kernel is defined from

$$\begin{aligned} \kappa_{f,g}(\mathbf{x}, \mathbf{y}) &:= \mathbb{E}_{\mathbf{\Omega}, \boldsymbol{\xi}} \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle \\ &= \frac{1}{m} \sum_{j=1}^m \mathbb{E}_{\boldsymbol{\omega}_j, \xi_j} f(\boldsymbol{\omega}_j^\top \mathbf{x} + \xi_j) g^*(\boldsymbol{\omega}_j^\top \mathbf{y} + \xi_j) \\ &= \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda, \xi \sim \mathcal{U}([0, 2\pi])} f(\boldsymbol{\omega}^\top \mathbf{x} + \xi) g^*(\boldsymbol{\omega}^\top \mathbf{y} + \xi). \end{aligned} \quad (10)$$

In the two bottom lines, we used the fact that $\boldsymbol{\omega}_j$ (and ξ_j) are independently and identically distributed, for all j , with $\boldsymbol{\omega}_j$ and ξ_j mutually independent. Remark that by the law of large numbers, $\kappa_{f,g}$ thus corresponds to the kernel that the asymmetric inner product $\widehat{\kappa}_{f,g}$ approximates when we let the feature space dimension m grow to infinity.

Proposition 3 (Expected kernel for asymmetric periodic random features). *Let \mathbf{z}_f and \mathbf{z}_g be random periodic features, associated with functions $f, g \in \text{PF}$, frequencies $\boldsymbol{\omega}_j \sim_{\text{i.i.d.}} \Lambda = \mathcal{F}\kappa^\Delta$ and $\xi_j \sim_{\text{i.i.d.}} \mathcal{U}([0, 2\pi])$. For any pair $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, the expected kernel $\kappa_{f,g}(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{\mathbf{\Omega}, \boldsymbol{\xi}} \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle$ satisfies*

$$\kappa_{f,g}(\mathbf{x}, \mathbf{y}) = \sum_{k \in \mathbb{Z}} F_k G_k^* \kappa^\Delta(k(\mathbf{x} - \mathbf{y})) =: \kappa_{f,g}^\Delta(\mathbf{x} - \mathbf{y}). \quad (11)$$

Although here expanded as an infinite series, this kernel is bounded $|\kappa_{f,g}| \leq 1$ since $f, g \in \text{PF}$.

Proof. Starting from the last line of (10), and decomposing f and g as their Fourier series,

$$\begin{aligned} \kappa_{f,g}(\mathbf{x}, \mathbf{y}) &= \mathbb{E}_{\boldsymbol{\omega}, \xi} \sum_{k \in \mathbb{Z}} \sum_{k' \in \mathbb{Z}} F_k G_{k'}^* e^{ik(\boldsymbol{\omega}^\top \mathbf{x} + \xi)} e^{-ik'(\boldsymbol{\omega}^\top \mathbf{y} + \xi)} \\ &= \sum_{k, k'} F_k G_{k'}^* \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} e^{i\boldsymbol{\omega}^\top (k\mathbf{x} - k'\mathbf{y})} \mathbb{E}_{\xi \sim \mathcal{U}([0, 2\pi])} e^{i(k - k')\xi} \\ &= \sum_{k, k'} F_k G_{k'}^* \kappa^\Delta(k\mathbf{x} - k'\mathbf{y}) \delta_{k, k'} \\ &= \sum_k F_k G_k^* \kappa^\Delta(k(\mathbf{x} - \mathbf{y})), \end{aligned} \quad (12)$$

where in the third line we used Bochner’s theorem (1) and the orthogonality of complex exponentials on one period: $\frac{1}{2\pi} \int_0^{2\pi} e^{ikt} e^{-ik't} dt = \delta_{k, k'}$. \square

Example 2. *As will be further developed in Sec. 6, when $f(\cdot) = \cos(\cdot)$ and $g(\cdot) = q(\cdot)$ the universal quantization defined in (8), we observe that the expected kernel is (up to a proportionality constant) exactly the ‘‘initial’’ kernel approximated by the RFF, i.e., $\kappa_{\cos, q}(\mathbf{x}, \mathbf{y}) = \frac{2}{\pi} \kappa(\mathbf{x}, \mathbf{y})$.*

The dither $\boldsymbol{\xi}$ plays here a crucial role: it cancels out (in expectation) the ‘‘cross-terms’’ in (12), each related to $F_k G_{k'}^* \kappa^\Delta(k\mathbf{x} - k'\mathbf{y}) = F_k G_{k'}^* \kappa(k\mathbf{x}, k'\mathbf{y})$, that have different scales $k \neq k'$ for \mathbf{x} and \mathbf{y} . As a consequence, the expected kernel is—as any kernel should be—conjugate symmetric, i.e., $\kappa_{f,g}(\mathbf{x}, \mathbf{y}) = \kappa_{f,g}^*(\mathbf{y}, \mathbf{x})$, despite the asymmetry of its empirical approximation, i.e., $\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) \neq \widehat{\kappa}_{f,g}^*(\mathbf{y}, \mathbf{x})$. The dithering can thus be thought of as a means to symmetrize, through expectation, the kernel associated with the asymmetric features inner product.

For the same reason, the dithering ensures that the expected kernel remains shift-invariant; Prop. 3 provides $\kappa_{f,g}(\mathbf{x}, \mathbf{y}) = \kappa_{f,g}^\Delta(\mathbf{x} - \mathbf{y})$, where κ^Δ in (11) is the kernel related to the frequency sampling pattern Λ . The expected kernel is thus a *scale mixture*, a linear combination of copies of κ^Δ , scaled (actually contracted) by an integer factor k (which is non-zero, since $F_0 = G_0 = 0$), and weighted by coefficients $F_k G_k^*$. In general, we expect this scale mixture $\kappa_{f,g}^\Delta$ to be narrower than the initial kernel κ^Δ (or more spread out in the frequency domain).

In general, however, the positive definiteness of κ does not imply that $\kappa_{f,g}$ is p.d., since taking, for instance, $g = -f$ (i.e., $F_k G_k^* = -|F_k|^2 < 0$) induces that $\kappa_{f,g}(\mathbf{x}, \mathbf{x}) = -\sum_k |F_k|^2 \kappa^\Delta(\mathbf{0}) < 0$ for all $\mathbf{x} \in \mathbb{R}^d$. Whether $\kappa_{f,g}$ is a positive definite kernel or not depends on the phase synchronization between the Fourier coefficients of f and g . A sufficient condition for $\kappa_{f,g}$ to be p.d. is to ensure that $F_k G_k^* \in \mathbb{R}^+$ for all k , as verified by taking $f = g$ in Sec. 5, or the combination $f = q$, $g(\cdot) = \cos(\cdot)$ in Sec. 6 and Sec. 7.

Remark. In light of (11), we could decide to normalize our approach differently. Assuming that $f, g \in \text{PF}$ are not orthogonal, i.e., $\langle f, g \rangle \neq 0$, we can define, for $\mathbf{x}, \mathbf{y} \in \Sigma$, the normalized kernels

$$\tilde{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) := \frac{1}{\langle f, g \rangle} \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle, \quad \dot{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) := \frac{1}{\langle f, g \rangle} \kappa_{f,g}(\mathbf{x}, \mathbf{y}), \quad \dot{\kappa}_{f,g}^\Delta(\mathbf{u}) := \frac{1}{\langle f, g \rangle} \kappa_{f,g}^\Delta(\mathbf{u}). \quad (13)$$

Since $\kappa^\Delta(\mathbf{0}) = 1$ and, from (11), $\langle f, g \rangle = \sum_k F_k G_k^* = \kappa_{f,g}^\Delta(\mathbf{0})$, (13) ensures that, for any $\mathbf{x} \in \Sigma$, $\mathbb{E} \tilde{\kappa}_{f,g}(\mathbf{x}, \mathbf{x}) = \dot{\kappa}_{f,g}(\mathbf{x}, \mathbf{x}) = \dot{\kappa}_{f,g}^\Delta(\mathbf{0}) = 1 = \kappa^\Delta(\mathbf{0})$. Without guaranteeing that $\dot{\kappa}_{f,g}$ is p.d., this normalization prevents the counterexample $f = -g$ to lead to a kernel with negative value on the origin. For clarity, we do not base our following developments on $\dot{\kappa}_{f,g}(\mathbf{x}, \mathbf{y})$ but we will refer to this useful quantity in Sec. 6 when, for $f = q$ and $g(\cdot) = \cos(\cdot)$, we will need to compare $\dot{\kappa}_{f,g}^\Delta$ to the RFF kernel κ^Δ .

Let us now provide an alternative expression of the expected kernel $\kappa_{f,g}(\mathbf{x}, \mathbf{y}) = \kappa_{f,g}^\Delta(\mathbf{x} - \mathbf{y})$, that will prove to be useful in the next section.

Lemma 1. Define the correlation h between f and g ,

$$h(t) := (f * \bar{g})(t) = \frac{1}{2\pi} \int_0^{2\pi} f(\tau) g^*(\tau - t) d\tau, \quad (14)$$

where $\bar{g}(t) := g^*(-t)$ denotes the conjugate reverse of g , and $*$ the convolution operator on $[0, 2\pi]$. The expected (shift-invariant) kernel $\kappa_{f,g}^\Delta$ can be expressed by

$$\kappa_{f,g}^\Delta(\mathbf{u}) = \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} h(\boldsymbol{\omega}^\top \mathbf{u}) = \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} (f * \bar{g})(\boldsymbol{\omega}^\top \mathbf{u}). \quad (15)$$

Proof. By the convolution theorem, the Fourier series coefficients of h are given by $H_k = F_k G_k^*$. The result follows from plugging this fact into the proof of Prop. 3. \square

Lemma 1 can be interpreted as an expansion similar to the one Bochner's theorem provides: whereas the initial kernel $\kappa^\Delta(\mathbf{u}) = \int e^{i\boldsymbol{\omega}^\top \mathbf{u}} d\Lambda(\boldsymbol{\omega})$ can be expressed in a basis that is a family complex exponentials $e^{i\boldsymbol{\omega}^\top \mathbf{u}}$ with "coordinates" given by Λ , $\kappa_{f,g}^\Delta(\mathbf{u}) = \int h(\boldsymbol{\omega}^\top \mathbf{u}) d\Lambda(\boldsymbol{\omega})$ can be expressed in a basis that is the family of functions $\{\mathbf{u} \rightarrow h(\boldsymbol{\omega}^\top \mathbf{u}) : \boldsymbol{\omega} \in \mathbb{R}^d\}$, another type of 2π -periodic functions (replacing the complex exponential $\exp(i\cdot)$ with $h(\cdot)$).

As a side note, when $\Lambda \sim \mathcal{N}(0, \mathbf{I}_d)$ and q is the square wave representing the universal quantization, Lemma 1 allows us to easily explain the linear slope of $\kappa_{q,q}^\Delta$ at the origin (see Fig. 2, left). Indeed, in this case h is the autocorrelation of q , the triangular wave

$$h(t) = (q * \bar{q})(t) = \max(1 - \frac{|t'|}{\pi}, \frac{|t'|}{\pi} - 1), \quad \text{with } t' := t \bmod 2\pi.$$

Therefore, defining $\tilde{h}(t) := 1 - |t|/\pi$, since $h(t) = \tilde{h}(t)$ for $|t| < \pi$, and $0 \leq h(t) - \tilde{h}(t) \leq 2(|t| - \pi)/\pi$ for $|t| \geq \pi$, we find

$$\kappa_{q,q}^\Delta(\mathbf{u}) = \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} h(|\boldsymbol{\omega}^\top \mathbf{u}|) = \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} \tilde{h}(|\boldsymbol{\omega}^\top \mathbf{u}|) + R(\mathbf{u}, \Lambda),$$

with $R(\mathbf{u}, \Lambda) := \mathbb{E}_{\omega \sim \Lambda} [h(|\omega^\top \mathbf{u}|) - \tilde{h}(|\omega^\top \mathbf{u}|)] \geq 0$. Since $\omega^\top \mathbf{u} \sim \|\mathbf{u}\|g$ with $g \sim \mathcal{N}(0, 1)$, and $\mathbb{E}|g| = \sqrt{2/\pi}$, we find

$$\begin{aligned} R(\mathbf{u}, \Lambda) &\leq \frac{2}{\sqrt{2\pi}} \int_{\pi/\|\mathbf{u}\|}^{+\infty} \frac{2}{\pi} (r\|\mathbf{u}\| - \pi) e^{-\frac{r^2}{2}} dr = \frac{c}{\|\mathbf{u}\|} \int_{\pi}^{+\infty} (s - \pi) e^{-\frac{s^2}{2\|\mathbf{u}\|^2}} ds \\ &= \frac{c}{\|\mathbf{u}\|} \int_0^{+\infty} s e^{-\frac{s^2 + \pi^2 + 2\pi s}{2\|\mathbf{u}\|^2}} ds \leq \frac{c}{2} e^{-\frac{\pi^2}{2\|\mathbf{u}\|^2}} \mathbb{E}[\|\mathbf{u}\||g|] = c'\|\mathbf{u}\| e^{-\frac{\pi^2}{2\|\mathbf{u}\|^2}}. \end{aligned}$$

Moreover, $\mathbb{E}_{\omega \sim \Lambda} h(|\omega^\top \mathbf{u}|) = 1 - \frac{\|\mathbf{u}\|}{\pi} \mathbb{E}_{\omega \sim \Lambda} |g| = 1 - \frac{\sqrt{2}}{\pi^{3/2}} \|\mathbf{u}\|$, which finally proves that

$$|\kappa_{q,q}^\Delta(\mathbf{u}) - (1 - \frac{\sqrt{2}}{\pi^{3/2}} \|\mathbf{u}\|)| = O(\|\mathbf{u}\| e^{-\frac{\pi^2}{2\|\mathbf{u}\|^2}}). \quad (16)$$

This shows that for $\|\mathbf{u}\| \ll \pi$, $\kappa_{q,q}^\Delta(\mathbf{u}) \asymp 1 - \frac{\sqrt{2}}{\pi^{3/2}} \|\mathbf{u}\|$.

4 Approximation error analysis (non-asymptotic case)

In the practical setting where the vectors $\mathbf{z}_f(\mathbf{x})$ and $\mathbf{z}_g(\mathbf{y})$ are to be quickly processed or stored in memory, their size m must be as small as possible. On the other hand, setting m too small hurts the empirical estimation $\hat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y})$ of the expected kernel $\kappa_{f,g}(\mathbf{x}, \mathbf{y}) = \mathbb{E} \hat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y})$. To understand this trade-off, we are thus interested in a probabilistic bound for the (absolute) kernel approximation error $|\hat{\kappa}_{f,g} - \kappa_{f,g}|$, as a function of the RPF dimension m . We give here an answer to this question under generic assumptions, and show how to apply it in a concrete situation—for asymmetric kernel estimation with one-bit quantized RFF—in Sec. 6.

4.1 Non-uniform approximation error

Ultimately, we want to obtain a (probabilistic) bound for the kernel approximation error that holds uniformly over all $\mathbf{x}, \mathbf{y} \in \Sigma$. First bounding the error for one *fixed* pair (\mathbf{x}, \mathbf{y}) is often used as an easier intermediary step. This is provided by the following proposition.

Proposition 4 (Non-uniform kernel approximation error from asymmetric periodic random features). *For two functions $f, g \in \text{PF}$, let $\mathbf{z}_f, \mathbf{z}_g$ be random periodic features associated with frequencies Ω and a dither ξ . For any fixed pair $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^d$, the inner product $\hat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle$ concentrates, in probability over the draw of $\Omega \sim \Lambda^m, \xi \sim \mathcal{U}^m([0, 2\pi])$, around $\kappa_{f,g}(\mathbf{x}, \mathbf{y}) = \mathbb{E}_{\Omega, \xi} \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle$ as*

$$\mathbb{P}[|\hat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f,g}(\mathbf{x}, \mathbf{y})| \leq \epsilon] \geq 1 - 2e^{-m\epsilon^2/2}. \quad (17)$$

Proof. We rewrite $\langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle = \frac{1}{m} \sum_j Z_j$, with the random variables $Z_j := f(\omega_j^\top \mathbf{x} + \xi_j) g^*(\omega_j^\top \mathbf{y} + \xi_j)$. The Z_j variables are i.i.d., have mean $\kappa_{f,g}(\mathbf{x}, \mathbf{y})$ by definition of the expected kernel (10), and are bounded by $|Z_j| \leq \|f\|_\infty \|g\|_\infty \leq 1$ (because $f, g \in \text{PF}$). The result follows by Hoeffding's inequality. \square

4.2 Uniform approximation error

We now want to extend the error bound in Prop. 4 to hold not only for one fixed pair (\mathbf{x}, \mathbf{y}) but simultaneously over all pairs $(\mathbf{x}, \mathbf{y}) \in \Sigma \times \Sigma$; this is called a *uniform bound*. The classical argument invoked in this type of proofs (e.g., [2, 25, 35]) goes as follows. If Σ is a finite set (of finite cardinality $|\Sigma|$), the uniform bound is obtained by applying a union bound over $|\Sigma|^2$ instances of Prop. 4 (one for each pair in $\Sigma \times \Sigma$). In the case where $\Sigma \subset \mathbb{R}^d$ is an infinite but compact set, the strategy is to bound the approximation error on a finite set Σ_η that covers Σ by balls of some radius $\eta > 0$, then to extend

this bound by some notion of continuity (smoothness) over the η -balls. We then obtain a bound which holds over $\Sigma_\eta + \eta\mathbb{B}_2^d \supseteq \Sigma$, which concludes the proof.

In our setting, the last step of this proof technique would ideally use Lipschitz continuity; we say that a function $f : \mathbb{R} \rightarrow \mathbb{C}$ is *Lipschitz continuous* with constant L_f if, for all $t, t' \in \mathbb{R}$, $|f(t) - f(t')| \leq L_f |t - t'|$, which is equivalent to

$$\forall t \in \mathbb{R}, \forall \delta > 0, \quad \sup_{r \in [-\delta, \delta]} \{|f(t+r) - f(t)|\} \leq L_f \cdot \delta. \quad (18)$$

However, this strategy fails when any of the maps f or g is not Lipschitz continuous (*e.g.*, when they present discontinuities, such as the “square wave” universal quantization map q from (8)). To be able to include such maps in our analysis, we must define a more permissive notion of smoothness, just as the T -part Lipschitz property defined in [13] (but without the limitations explained in Sec. 5). In this work, we rather introduce the concept of *mean Lipschitz smoothness property* for periodic function in PF. Intuitively, a periodic function is smooth in the mean Lipschitz sense if its largest local deviation is small *on average*.

Definition 3 (Mean Lipschitz property). *Let $f : \mathbb{R} \rightarrow \mathbb{C}$ be a generic periodic function (here w.l.o.g. assumed of period 2π). We say it is mean Lipschitz smooth with mean Lipschitz constant L_f^μ if for all radii $\delta \in (0, \pi]$, the average maximum deviation of f in $[-\delta, \delta]$ is bounded by $L_f^\mu \cdot \delta$:*

$$\mathbb{E}_{t \sim \mathcal{U}([0, 2\pi])} \sup_{r \in [-\delta, \delta]} \{|f(t+r) - f(t)|\} = \frac{1}{2\pi} \int_0^{2\pi} \sup_{r \in [-\delta, \delta]} \{|f(t+r) - f(t)|\} dt \leq L_f^\mu \cdot \delta. \quad (19)$$

The mean Lipschitz property (that we will refer to as “mean smoothness” to avoid confusion with the usual Lipschitz continuity when necessary) can truly be understood as the Lipschitz continuity after an averaging. It is reminiscent of the mean modulus of continuity from [36,37] but where the order of the supremum and averaging operations are reversed (which is less restrictive). If f is Lipschitz continuous with Lipschitz constant L_f , then it has necessarily also the mean smoothness property with constant $L_f^\mu \leq L_f$. However, it is possible that $L_f^\mu \ll L_f$ (if the large slopes of f are concentrated on a small portion of $[0, 2\pi]$), and discontinuous function can have a finite L_f^μ constant—for example, the square wave q representing the universal quantization is mean smooth with constant $L_q^\mu = \frac{4}{\pi}$ (see Prop. 6), although it is not a Lipschitz continuous function. Leaving the detailed proof to Sec. 6, the trick is to observe that the integrand $I_\delta(t) = \sup_{r \in [-\delta, \delta]} \{|f(t+r) - f(t)|\}$ is supported on an interval whose length is proportional to δ , as shown Fig. 2, left. Moreover, the convolution of any PF function with a mean smooth PF function yields a Lipschitz continuous one.

Lemma 2. *Given two functions $f, g \in \text{PF}$, among which f is mean smooth with constant L_f^μ , their convolution $(f * g)$ is Lipschitz continuous with constant $L_{f * g} \leq L_f^\mu$.*

Proof. Re-writing (18) for $(f * g)(t) = \frac{1}{2\pi} \int_0^{2\pi} f(t - \tau)g(\tau) d\tau$ gives, since $\|g\|_\infty \leq 1$,

$$\begin{aligned} \sup_{r \in [-\delta, \delta]} \{|(f * g)(t+r) - (f * g)(t)|\} &= \sup_{|r| \leq \delta} \left| \frac{1}{2\pi} \int_0^{2\pi} [f(t+r-\tau) - f(t-\tau)] g(\tau) d\tau \right| \\ &\leq \frac{1}{2\pi} \sup_{|r| \leq \delta} \int_0^{2\pi} |f(t+r-\tau) - f(t-\tau)| d\tau \\ &\leq \frac{1}{2\pi} \int_0^{2\pi} \sup_{|r| \leq \delta} |f(\tau'+r) - f(\tau')| d\tau' \leq L_f^\mu \delta. \end{aligned}$$

□

In other words, the convolution of two PF functions, among whom one of them is mean Lipschitz, is “smoother” than its factors, a property that comes from the convolution itself (which is to be put in

correspondence with the fact that for f differentiable and g discontinuous, $f * g$ is differentiable). In particular, if *both* f and g are mean smooth with constants L_f^μ and L_g^μ respectively, their *correlation* $h = f * \bar{g}$ is Lipschitz with $L_h \leq \min(L_f^\mu, L_g^\mu)$. Coming back to our setting, this fact allows us (using Lemma 1) to characterize the Lipschitz continuity of the expected kernel $\kappa_{f,g}^\Delta$. With that, we have all the tools to prove our main result, a uniform bound on the kernel approximation error obtained with possibly discontinuous (but mean smooth) maps.

Proposition 5 (Uniform kernel approximation error from asymmetric periodic random features). *Let Σ be a compact set and $f, g \in \text{PF}$ periodic functions with finite mean smoothness constants L_f^μ and L_g^μ , respectively, and let $C_\Lambda < \infty$ such that $\mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} |\boldsymbol{\omega}^\top \mathbf{a}| \leq C_\Lambda \|\mathbf{a}\|_2$ for all \mathbf{a} (the kernel smoothness constant). For all error level $\epsilon > 0$, provided the feature dimension is larger than*

$$m \geq 128 \cdot \frac{1}{\epsilon^2} \cdot \mathcal{H}_{\epsilon/c}(\Sigma), \quad (20)$$

with the constant $c = 4C_\Lambda(L_f^\mu + L_g^\mu + 2\min(L_f^\mu, L_g^\mu))$, the following kernel approximation bounds holds uniformly:

$$|\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f,g}(\mathbf{x}, \mathbf{y})| \leq \epsilon, \quad \forall \mathbf{x}, \mathbf{y} \in \Sigma, \quad (21)$$

with probability exceeding $1 - 3 \exp(-\frac{m\epsilon^2}{64})$.

Proof. With Σ_η a finite optimal η -covering of Σ , any $\mathbf{x}' \in \Sigma$ (resp. $\mathbf{y}' \in \Sigma$) can be written $\mathbf{x}' = \mathbf{x} + \mathbf{r}_x$ (resp. $\mathbf{y}' = \mathbf{y} + \mathbf{r}_y$) for centers $\mathbf{x}, \mathbf{y} \in \Sigma_\eta$, and $\mathbf{r}_x, \mathbf{r}_y \in \eta\mathbb{B}_2^d$. The proof proceeds by defining three events $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ from which Prop. 5 follows, and then by bounding the failure probability of their joint occurrence. First, for any covering center $\mathbf{x} \in \Sigma_\eta$, we can expect that the set of m functions $h_j^f : \eta\mathbb{B}_2^d \rightarrow \mathbb{C}$ defined for $j \in [m]$ as $h_j^f(\mathbf{r}; \mathbf{x}) := f(\boldsymbol{\omega}_j^\top \mathbf{x} + \boldsymbol{\omega}_j^\top \mathbf{r} + \xi_j)$ contains, on average, few ‘‘variations’’ over the η -ball. More precisely, defining the largest variation of $h_j^f(\mathbf{r}; \mathbf{x})$ over the η -ball as

$$H_j^f(\eta; \mathbf{x}) := \sup_{\mathbf{r} \in \eta\mathbb{B}_2^d} |h_j^f(\mathbf{r}; \mathbf{x}) - h_j^f(\mathbf{0}; \mathbf{x})| = \sup_{\mathbf{r} \in \eta\mathbb{B}_2^d} |f(\boldsymbol{\omega}_j^\top \mathbf{x} + \boldsymbol{\omega}_j^\top \mathbf{r} + \xi_j) - f(\boldsymbol{\omega}_j^\top \mathbf{x} + \xi_j)|,$$

we first assume that, given $\epsilon_1 > 0$, the event \mathcal{E}_1 holds, with

$$\mathcal{E}_1 : \sup_{\mathbf{x} \in \Sigma_\eta} \frac{1}{m} \sum_{j=1}^m H_j^f(\eta; \mathbf{x}) \leq L_f^\mu \eta C_\Lambda + \epsilon_1.$$

Similarly for g , we define the event \mathcal{E}_2 such that

$$\mathcal{E}_2 : \sup_{\mathbf{y} \in \Sigma_\eta} \frac{1}{m} \sum_{j=1}^m H_j^g(\eta; \mathbf{y}) \leq L_g^\mu \eta C_\Lambda + \epsilon_2.$$

Next, we suppose that the kernel approximation has error bounded by ϵ_3 for all the covering centers $\mathbf{x}, \mathbf{y} \in \Sigma_\eta$, i.e.,

$$\mathcal{E}_3 : \sup_{\mathbf{x}, \mathbf{y} \in \Sigma_\eta} |\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f,g}(\mathbf{x}, \mathbf{y})| \leq \epsilon_3.$$

Under those events, we establish a deterministic bound for all \mathbf{x}', \mathbf{y}' using a chain of triangle inequalities:

$$|\widehat{\kappa}_{f,g}(\mathbf{x}', \mathbf{y}') - \kappa_{f,g}(\mathbf{x}', \mathbf{y}')| = |\widehat{\kappa}_{f,g}(\mathbf{x} + \mathbf{r}_x, \mathbf{y} + \mathbf{r}_y) - \kappa_{f,g}(\mathbf{x} + \mathbf{r}_x, \mathbf{y} + \mathbf{r}_y)| \leq \delta_1 + \delta_2 + \delta_3 + \delta_4,$$

where the error terms are defined as

$$\begin{aligned} \delta_1 &:= |\widehat{\kappa}_{f,g}(\mathbf{x} + \mathbf{r}_x, \mathbf{y} + \mathbf{r}_y) - \widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y} + \mathbf{r}_y)|, & \delta_3 &:= |\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f,g}(\mathbf{x}, \mathbf{y})|, \\ \delta_2 &:= |\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y} + \mathbf{r}_y) - \widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y})|, & \delta_4 &:= |\kappa_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f,g}(\mathbf{x} + \mathbf{r}_x, \mathbf{y} + \mathbf{r}_y)|. \end{aligned}$$

First, we observe that, thanks to \mathcal{E}_1 :

$$\begin{aligned} \delta_1 &= \frac{1}{m} \left| \sum_{j=1}^m \left[f(\boldsymbol{\omega}_j^\top (\mathbf{x} + \mathbf{r}_x) + \xi_j) - f(\boldsymbol{\omega}_j^\top \mathbf{x} + \xi_j) \right] g^*(\boldsymbol{\omega}_j^\top \mathbf{y}' + \xi_j) \right| \\ &\leq \|g\|_\infty \cdot \frac{1}{m} \sum_{j=1}^m \left| f(\boldsymbol{\omega}_j^\top \mathbf{x} + \boldsymbol{\omega}_j^\top \mathbf{r}_x + \xi_j) - f(\boldsymbol{\omega}_j^\top \mathbf{x} + \xi_j) \right| \leq \eta L_f^\mu C_\Lambda + \epsilon_1. \end{aligned}$$

Similarly, using \mathcal{E}_2 , we get for δ_2 :

$$\delta_2 = \frac{1}{m} \left| \sum_{j=1}^m f(\boldsymbol{\omega}_j^\top \mathbf{x} + \xi_j) \left[g^*(\boldsymbol{\omega}_j^\top \mathbf{y} + \boldsymbol{\omega}_j^\top \mathbf{r}_y + \xi_j) - g^*(\boldsymbol{\omega}_j^\top \mathbf{y} + \xi_j) \right] \right| \leq \eta L_g^\mu C_\Lambda + \epsilon_2.$$

Regarding δ_3 , we directly get from \mathcal{E}_3 :

$$\delta_3 \leq \sup_{\mathbf{x}, \mathbf{y} \in \Sigma_\eta} |\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f,g}(\mathbf{x}, \mathbf{y})| \leq \epsilon_3.$$

Finally, for δ_4 , denoting $\mathbf{u} := \mathbf{x} - \mathbf{y}$, $\mathbf{r}_u := \mathbf{r}_x - \mathbf{r}_y \in 2\eta\mathbb{B}_2^d$ and using Lemma 1 (also recall that $h = f * \bar{g}$ is Lipschitz continuous with $L_h \leq \min(L_f^\mu, L_g^\mu)$ by Lemma 2), as well as the definition of C_Λ in (4),

$$\begin{aligned} \delta_4 &= \left| \kappa_{f,g}^\Delta(\mathbf{u}) - \kappa_{f,g}^\Delta(\mathbf{u} + \mathbf{r}_u) \right| = \left| \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} h(\boldsymbol{\omega}^\top \mathbf{u}) - h(\boldsymbol{\omega}^\top \mathbf{u} + \boldsymbol{\omega}^\top \mathbf{r}_u) \right| \\ &\leq \mathbb{E}_{\boldsymbol{\omega}} |h(\boldsymbol{\omega}^\top \mathbf{u}) - h(\boldsymbol{\omega}^\top \mathbf{u} + \boldsymbol{\omega}^\top \mathbf{r}_u)| \leq \mathbb{E}_{\boldsymbol{\omega}} L_h \cdot |\boldsymbol{\omega}^\top \mathbf{r}_u| \leq L_h C_\Lambda \|\mathbf{r}_u\|_2 \leq 2\eta C_\Lambda \cdot \min(L_f^\mu, L_g^\mu). \end{aligned}$$

Putting everything back together, under \mathcal{E}_1 , \mathcal{E}_2 , and \mathcal{E}_3 , for any $\mathbf{x}', \mathbf{y}' \in \Sigma$:

$$|\widehat{\kappa}_{f,g}(\mathbf{x}', \mathbf{y}') - \kappa_{f,g}(\mathbf{x}', \mathbf{y}')| \leq \epsilon_1 + \epsilon_2 + \epsilon_3 + \eta C_\Lambda \left(L_f^\mu + L_g^\mu + 2 \cdot \min(L_f^\mu, L_g^\mu) \right). \quad (22)$$

It remains to bound the failure probability for each event. For event \mathcal{E}_1 , we have to bound the probability

$$\mathbb{P}[\bar{\mathcal{E}}_1] = \mathbb{P} \left[\exists \mathbf{x} \in \Sigma_\eta \text{ s.t. } \frac{1}{m} \sum_{j=1}^m H_j^f(\eta; \mathbf{x}) \geq L_f^\mu \eta C_\Lambda + \epsilon_1 \right].$$

We first focus on one single center $\mathbf{x} \in \Sigma_\eta$. Each associated $H_j^f(\eta; \mathbf{x}) = \sup_{\mathbf{r} \in \eta\mathbb{B}_2^d} |f(\boldsymbol{\omega}_j^\top \mathbf{x} + \boldsymbol{\omega}_j^\top \mathbf{r} + \xi_j) - f(\boldsymbol{\omega}_j^\top \mathbf{x} + \xi_j)|$ is a random variable identically and independently distributed (where the randomness is due to the draw of $\boldsymbol{\omega}_j$ and ξ_j). The expectation $\mathbb{E} H_j^f$ of those variables is bounded by (we use the mean Lipschitz smoothness (19) with $t = \xi_j$, $r = \boldsymbol{\omega}_j^\top \mathbf{r}$ and $\delta = |\boldsymbol{\omega}_j^\top \mathbf{r}|$):

$$\mathbb{E} H_j^f = \mathbb{E}_{\boldsymbol{\omega}_j} \mathbb{E}_{\xi_j} \sup_{\mathbf{r} \in \eta\mathbb{B}_2^d} |f(\boldsymbol{\omega}_j^\top \mathbf{x} + \boldsymbol{\omega}_j^\top \mathbf{r} + \xi_j) - f(\boldsymbol{\omega}_j^\top \mathbf{x} + \xi_j)| \leq L_f^\mu \mathbb{E}_{\boldsymbol{\omega}_j} |\boldsymbol{\omega}_j^\top \mathbf{r}| = L_f^\mu \eta C_\Lambda.$$

Now we describe how the sum $\frac{1}{m} \sum_{j=1}^m H_j^f(\eta; \mathbf{x})$ concentrates around its mean $\mathbb{E} H_j^f$ with Hoeffding's inequality (note that $0 \leq H_j^f(\eta; \mathbf{x}) \leq 2\|f\|_\infty \leq 2$), and use $\mathbb{E} H_j^f \leq L_f^\mu \eta C_\Lambda$ to get a probabilistic bound

$$\mathbb{P} \left[\frac{1}{m} \sum_j H_j^f(\eta; \mathbf{x}) \geq L_f^\mu \eta C_\Lambda + \epsilon_1 \right] \leq \mathbb{P} \left[\frac{1}{m} \sum_j H_j^f(\eta; \mathbf{x}) - \mathbb{E} H_j^f \geq \epsilon_1 \right] \leq \exp \left(-\frac{m\epsilon_1^2}{2} \right).$$

We take a union bound of this result over the $\mathcal{C}_\eta(\Sigma) = |\Sigma_\eta|$ centers $\mathbf{x} \in \Sigma_\eta$ to obtain

$$\mathbb{P}[\bar{\mathcal{E}}_1] \leq \mathcal{C}_\eta(\Sigma) \exp(-m\epsilon_1^2/2) = \exp(\mathcal{H}_\eta(\Sigma) - m\epsilon_1^2/2).$$

Moreover, if $m \geq 4\mathcal{H}_\eta(\Sigma)\epsilon_1^{-2}$, then we get $\mathbb{P}[\bar{\mathcal{E}}_1] \leq e^{-m\epsilon_1^2/4}$. An identical development for \mathcal{E}_2 yields $\mathbb{P}[\bar{\mathcal{E}}_2] \leq e^{-m\epsilon_2^2/4}$ if $m \geq 4\mathcal{H}_\eta(\Sigma)\epsilon_2^{-2}$. For \mathcal{E}_3 , an union bound of Prop. 4 on all pairs in $\Sigma_\eta \times \Sigma_\eta$ gives

$$\mathbb{P}[\bar{\mathcal{E}}_3] \leq 2(\mathcal{C}_\eta(\Sigma))^2 e^{-m\epsilon_3^2/2} \leq \mathcal{C}_\eta(\Sigma)^2 e^{-m\epsilon_3^2/2} = e^{2\mathcal{H}_\eta(\Sigma) - m\epsilon_3^2/2}.$$

Moreover, if $m \geq 8\mathcal{H}_\eta(\Sigma)\epsilon_3^{-2}$, we get $\mathbb{P}[\bar{\mathcal{E}}_3] \leq e^{-m\epsilon_3^2/4}$. By union bound, and provided

$$m \geq 4\mathcal{H}_\eta(\Sigma) \cdot \max(\epsilon_1^{-2}, \epsilon_2^{-2}, 2\epsilon_3^{-2}),$$

the probability of failure of the deterministic bound above is lower than

$$\mathbb{P}[\bar{\mathcal{E}}_1 \cup \bar{\mathcal{E}}_2 \cup \bar{\mathcal{E}}_3] \leq \mathbb{P}[\bar{\mathcal{E}}_1] + \mathbb{P}[\bar{\mathcal{E}}_2] + \mathbb{P}[\bar{\mathcal{E}}_3] = e^{-m\epsilon_1^2/4} + e^{-m\epsilon_2^2/4} + e^{-m\epsilon_3^2/4}.$$

Finally, the desired result (less generic but more meaningful) is found by imposing equal contributions $\epsilon/4$ by each error term in (22), *i.e.*, $\epsilon_1 = \epsilon_2 = \epsilon_3 = \epsilon/4$ and $\eta = \epsilon/(4C_\Lambda[L_f^\mu + L_g^\mu + 2\min(L_f^\mu, L_g^\mu)])$. \square

Prop. 5 shows that we can control (e.g., by increasing m) the kernel approximation error uniformly, provided we control the smoothness of the “initial” kernel $\kappa^\Delta = \mathcal{F}^{-1}\Lambda$ (through C_Λ) and the mean smoothness of the maps f and g . Improvements are possible, for example, by more carefully setting the values of $\{\epsilon_1, \epsilon_2, \epsilon_3\}$ and η . If Σ is *structured* (e.g., if it consists of sparse vectors or low-rank matrices) and Λ is Gaussian, the value of C_Λ in $\mathbb{E}_\omega |\omega^\top \mathbf{r}| \leq \eta C_\Lambda$ for $\mathbf{r} \in (\Sigma - \Sigma) \cap \eta \mathbb{B}_2^d$ (which controls the bounds on δ_1, δ_2 and δ_4) can be related to the Gaussian mean width of $\Sigma - \Sigma$ [38].

Example 3. Consider once again the case of a Gaussian kernel with unit bandwidth (i.e., $C_\Lambda = 1$), with a signal space Σ made of bounded signals (inside the unit Euclidean ball \mathbb{B}_2^d) lying in a union of S subspaces of \mathbb{R}^d with dimension s . In this case, according to the entropy of this signal model (see Ex. 1), the kernel approximation error $|\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f,g}(\mathbf{x}, \mathbf{y})|$ is uniformly bounded over Σ , with high probability, provided that the number of features satisfies $m \geq C \epsilon^{-2} [s \log(\frac{4}{\epsilon}(L_f^\mu + L_g^\mu + 2 \min(L_f^\mu, L_g^\mu))) + \log S]$. For instance, for bounded s -sparse signals, we need $m \geq C s \cdot \epsilon^{-2} \log(\frac{4ed}{c\epsilon}(L_f^\mu + L_g^\mu + 2 \min(L_f^\mu, L_g^\mu)))$.

We conclude this section by showing that Prop. 5 allows characterizing the proximity of two approximated kernels $\widehat{\kappa}_{f,g}$ and $\widehat{\kappa}_{f',g}$ (given three functions $f, f', g \in \text{PF}$) when they are related by identical expectations $\mathbb{E}\widehat{\kappa}_{f,g} = \mathbb{E}\widehat{\kappa}_{f',g} = \kappa_0$. While this result could be achieved by a simple use of the triangular inequality—from $|\widehat{\kappa}_{f,g} - \widehat{\kappa}_{f',g}| \leq |\widehat{\kappa}_{f,g} - \kappa_0| + |\widehat{\kappa}_{f',g} - \kappa_0|$ and using the same proposition to bound the last two terms—the following corollary provides a more direct bound, possibly tighter.

Corollary 1 (Proximity of approximated RPF kernels). *Given $\epsilon > 0$, a compact set Σ , two 2π -periodic functions f, f' such that their difference $f - f' \in \text{PF}$, as well as a third periodic function $g \in \text{PF}$, such that there exist finite mean smoothness constants $L_{f-f'}^\mu$ and L_g^μ , and Λ such that $C_\Lambda < \infty$, if $\kappa_{f,g}(\cdot, \cdot) = \kappa_{f',g}(\cdot, \cdot)$ and if the feature dimension is larger than*

$$m \geq 128 \cdot \frac{1}{\epsilon^2} \cdot \mathcal{H}_{\epsilon/c}(\Sigma), \quad (23)$$

with constant $c = 4C_\Lambda(L_{f-f'}^\mu + L_g^\mu + 2 \min(L_{f-f'}^\mu, L_g^\mu))$, then

$$|\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \widehat{\kappa}_{f',g}(\mathbf{x}, \mathbf{y})| \leq \epsilon, \quad \forall \mathbf{x}, \mathbf{y} \in \Sigma, \quad (24)$$

with probability exceeding $1 - 3 \exp(-\frac{m\epsilon^2}{64})$.

Proof. We simply observe that, by linearity of the kernels with respect to their supporting functions, for any $\mathbf{x}, \mathbf{y} \in \Sigma$, $\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \widehat{\kappa}_{f',g}(\mathbf{x}, \mathbf{y}) = \widehat{\kappa}_{\tilde{f},g}(\mathbf{x}, \mathbf{y})$ with $\tilde{f} := f - f'$. The proof then follows by applying Prop. 5 to the RPFs supported by $\tilde{f}, g \in \text{PF}$, with the vanishing kernel $\mathbb{E}\widehat{\kappa}_{\tilde{f},g}(\mathbf{x}, \mathbf{y}) = \kappa_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f',g}(\mathbf{x}, \mathbf{y}) = 0$. \square

In Sec. 7.3, we will use this corollary in combination with Prop. 5 to compare the performance of a machine learning algorithm (the kernel support vector machine, SVM) on a given classification task when learning and inference are using identical approximated kernels (i.e., when the learning is performed using the RFF) or only kernels that are asymptotically equal (when the learning stage uses the expected kernel).

5 Asymmetric geometry-preserving embedding

Our approach can be related to the context of geometry-preserving embedding (or coding) developed in [13]. This allows us to provide another version of one of their central results, [13, Thm 3.2], whose proof is incorrect (as described below). While the alternative result we propose looks slightly different, it fulfills the same goal: a non-asymptotic guarantee for the geometry-preserving capabilities of the embedding (6) with discontinuous f , which holds on infinite signal sets. This section can be seen as a first (theoretical) application of Prop. 5.

5.1 Geometry-preserving embedding: the initial approach

In [13] the authors study when a mapping $\varphi : \Sigma \rightarrow \mathbb{C}^m$ (such as \mathbf{z}_f defined in (6) for $f \in \text{PF}$) defines an embedding of Σ into \mathbb{C}^m approximately preserving the proximity of vectors in Σ . This proximity is measured by the (local) preservation a distance associated with a $\ell_{\#}$ -norm $\|\cdot\|_{\#}$ (e.g., the ℓ_1 or the ℓ_2 -norm). Adapting their setting to our conventions¹, given some $\epsilon, \delta > 0$, and an invertible function, or *distance map*, $\gamma : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, they study the conditions ensuring that φ is a $(\gamma, \delta, \epsilon)$ -embedding of Σ (endowed with the $\ell_{\#}$ -norm) into \mathbb{C}^m ; or mathematically, such that φ respects

$$(1 - \delta) \gamma(\|\mathbf{x} - \mathbf{y}\|_{\#}) - \epsilon \leq \|\varphi(\mathbf{x}) - \varphi(\mathbf{y})\|^2 \leq (1 + \delta) \gamma(\|\mathbf{x} - \mathbf{y}\|_{\#}) + \epsilon, \quad (25)$$

for all $\mathbf{x}, \mathbf{y} \in \Sigma$.

In (25), γ maps distances in Σ to (squared) distances in \mathbb{C}^m , and δ and ϵ quantify the multiplicative and the additive error, respectively, of the embedding associated with the map γ . For instance, if φ is linear with $\varphi(\mathbf{x}) = \mathbf{A}\mathbf{x}$, there exists many random constructions of the $m \times d$ matrix \mathbf{A} with appropriate scaling (e.g., random Gaussian matrix or random partial Fourier matrix [1]) for which (25) holds with high probability with $\ell_{\#} \equiv \ell_2$, $\epsilon = 0$, and $\gamma(t) = t^2$ for $\Sigma = \Sigma_k$ and $m = O(\delta^{-2} k \log(n/k))$. Similarly, in the context of one-bit compressive sensing where $\varphi(\mathbf{x}) = (cm)^{-1/2} \text{sign}(\mathbf{A}\mathbf{x})$ (for some suitable $c > 0$), $\|\varphi(\mathbf{x}) - \varphi(\mathbf{y})\|^2$ represents the (scaled) Hamming distance between the two binary vectors $\varphi(\mathbf{x})$ and $\varphi(\mathbf{y})$, and (25) is verified with high probability over $\Sigma_k \cap \mathbb{B}^d$ with $m = O(\epsilon^{-2} k \log(n/k))$, $\ell_{\#} \equiv \ell_2$, $\delta = 0$, and $\gamma(t) = t$ [39]. The work [13] extends this analysis to general nonlinear feature maps $\varphi(\cdot) = \mathbf{z}_f(\cdot)$ for some periodic function f (such as the universal quantizer q). In such a context, the authors show that (25) holds with a map γ that often displays two regimes: a linear regime for small distances in Σ ($\mathbf{x} \approx \mathbf{y}$), and a saturation regime where γ quickly flattens after a certain distance.

As explained in [13, Sec. 4.5], this approach is connected to the approximation of a kernel $\kappa : \Sigma \times \Sigma \rightarrow \mathbb{R}_+$ from the inner product of the images of two vectors, namely, for which $\langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_f(\mathbf{y}) \rangle \approx \kappa(\mathbf{x}, \mathbf{y})$. Assuming $\|\mathbf{z}_f\| = 1$ for simplicity, which is the case for complex exponential and universal quantization features, we find $\|\mathbf{z}_f(\mathbf{x}) - \mathbf{z}_f(\mathbf{y})\|^2 = 2(1 - \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_f(\mathbf{y}) \rangle)$. Therefore, if \mathbf{z}_f is a $(\gamma, 0, 2\epsilon)$ -embedding of Σ into \mathbb{C}^m , then

$$\kappa(\mathbf{x}, \mathbf{y}) - \epsilon \leq \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_f(\mathbf{y}) \rangle \leq \kappa(\mathbf{x}, \mathbf{y}) + \epsilon, \quad (26)$$

for all $\mathbf{x}, \mathbf{y} \in \Sigma$, provided we define the kernel

$$\kappa(\mathbf{x}, \mathbf{y}) := 1 - \frac{1}{2} \gamma(\|\mathbf{x} - \mathbf{y}\|_{\#}). \quad (27)$$

The smoothness of κ is thus directly connected to the one of γ ; for instance, if γ is Lipschitz continuous with constant $L_{\gamma} \geq 0$ of \mathbb{R}_+ , then, from the invertibility of γ over \mathbb{R}^+ , κ is Lipschitz continuous with constant $L_{\gamma}/2$ with respect to any of its argument. Note that, from Lemma 1 and Lemma 2, we also know that if f is mean smooth with constant L_f^{μ} , then $\kappa(\mathbf{x}, \mathbf{y}) = \kappa_{f,f}^{\Delta}(\mathbf{x} - \mathbf{y})$ is Lipschitz continuous with constant $L_{\kappa} \leq C_{\Lambda} L_f^{\mu}$ with respect to any of its argument (as proved from the bound on δ_4 in the proof of Prop. 5). This shows that, despite their different origin, the smoothness of γ (in the approach [13]) and the one of f (in ours) control the one of κ .

Compared to our approach, [13] imposes the periodic function f to be ‘‘Lipschitz continuous by part’’ (rather than being mean smooth), as defined hereafter in a setting adapted to our needs.

Definition 4 (*T-part Lipschitz continuity* [13, Def. 2.1]). *A function $f : \Sigma \rightarrow \mathbb{C}$ is T-part Lipschitz continuous over $\mathcal{S} \subset \Sigma$ with constant $\bar{L}_f \geq 0$, if there exists a finite partition $\{\mathcal{S}_t\}_{t=1}^T$ of \mathcal{S} into T disjoint sets (i.e., $\bigcup_{t=1}^T \mathcal{S}_t = \mathcal{S}$) such that*

$$\forall t \in [T], \forall \mathbf{x}, \mathbf{y} \in \mathcal{S}_t, \quad |f(\mathbf{x}) - f(\mathbf{y})| \leq \bar{L}_f \cdot \|\mathbf{x} - \mathbf{y}\|_{\#}. \quad (28)$$

¹Hereafter, departing from the general approach of [13], we always consider the simplified case where \mathbb{C}^m is equipped with the Euclidean distance, with a squaring of the corresponding distance in (25).

Moreover, f is exactly T -part Lipschitz continuous over \mathcal{S} with constant $\bar{L}_f \geq 0$, which we write $f \in \overline{\text{Lip}}(\mathcal{S}, T, \bar{L}_f)$, if it is T -part Lipschitz continuous with that constant but is not $(T - 1)$ -part Lipschitz continuous with the same constant.

Note that (28) both generalizes (18) to functions from $\Sigma \subset \mathbb{R}^d$ to \mathbb{C} , and localizes (18) on \mathcal{S} .

Following the convention of our paper, the authors of [13] then prove the following result. We simplify it to the case $\delta = 0$ and where each component of \mathbf{z}_f has at most T parts of continuity (despite its randomness).

Theorem 1 (Adapted from [13, Thm 3.2]). *Given $0 < \epsilon < 1$, an L_γ -Lipschitz continuous distance map $\gamma : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, and a signal set Σ with finite covering number $\mathcal{C}_\eta(\Sigma)$ for any radius $\eta > 0$, let us assume that, for any fixed pair of vectors $\mathbf{x}, \mathbf{y} \in \Sigma$, the mapping \mathbf{z}_f defined in (6) satisfies the embedding relation (25) (for $\delta = 0$) with probability exceeding $1 - Ce^{-cm\epsilon^2}$.*

Let us suppose that there exists a constant $\bar{L}_f \geq 0$ such that, for any $\mathbf{x} \in \Sigma$, integer $t \geq 1$, radius $\eta > 0$, and given $\mathcal{S}_\mathbf{x}(\eta) := \{\mathbf{u} \in \Sigma : \|\mathbf{u} - \mathbf{x}\|_\# \leq \eta\}$ (a neighborhood of \mathbf{x} of radius η),

$$\mathbb{P}[(\mathbf{z}_f(\cdot))_k \in \overline{\text{Lip}}(\mathcal{S}_\mathbf{x}(\eta), T, \bar{L}_f)] \leq p_t(\eta), \quad (29)$$

with p_t independent of \mathbf{x} , and $p_t(\eta) = 0$ if $t > T$ for some integer $T \geq 0$.

In this context, defining $\rho(\eta, T) := \sum_{t=2}^T p_t(\eta) \log t$ and $\nu := \frac{1}{4}(L_\gamma + \bar{L}_f)^{-1}$, provided that $\epsilon^2 \geq C\rho(2\nu\epsilon^2, T)$ and

$$m \geq C\epsilon^{-2}(\mathcal{H}_{\nu\epsilon^2}(\Sigma) + \log T), \quad (30)$$

the mapping \mathbf{z}_f is a $(\gamma, 0, 2\epsilon)$ -embedding with probability exceeding $1 - Ce^{-c\epsilon^2 m}$, i.e., \mathbf{z}_f respects

$$\gamma(\|\mathbf{x} - \mathbf{y}\|_\#) - 2\epsilon \leq \|\mathbf{z}_f(\mathbf{x}) - \mathbf{z}_f(\mathbf{y})\|^2 \leq \gamma(\|\mathbf{x} - \mathbf{y}\|_\#) + 2\epsilon, \quad \forall \mathbf{x}, \mathbf{y} \in \Sigma.$$

The statement of this theorem is an easy adaptation of [13, Thm 3.2] where we set $\delta = 0$, $w(\epsilon, \delta) = c\epsilon^2$, $c_0 = c\epsilon$, $T_{\max} = T$ (so that $P_F = 0$), and $\alpha = \epsilon^2 \leq \epsilon \leq 1$.

Note that the first assumption of this theorem (regarding the fact that (25) holds with high probability for any fixed pair of signals) is proven in a separate result, namely [13, Thm 4.1]. This theorem is similar to our Prop. 4 (up to an easy extension of this proposition to a finite set of pairs by union bound). Since the flaw developed below is independent of that separate result, we abstract this specific assumption away in this work. As explained in [13, App. E], the conditions of this theorem can thus be met for instance in the case where f is the universal quantizer. One can then show that $T = 2$, and defining $p_2(\eta) := D\eta$, with $D > 0$ function of d and Λ , is appropriate for the bound (29). Therefore, $\rho(2\nu\epsilon^2, T) \leq 2D\nu\epsilon^2 \leq \epsilon^2/C$ for an appropriate $C > 0$.

The statement of this theorem bears similarities with our Prop. 5 in the case where $f = g$; in essence, keeping in mind the equivalence (26), up to a smaller covering radius scaling as $\epsilon^2 < \epsilon < 1$ in (30), the constraint (30) is similar to (20) if we consider that the T -part Lipschitz continuity of f replaces its mean smoothness.

However, the proof of Theorem 1 in [13, App. B] is incorrect. Let us see why by sketching their arguments in our system of notations and using $\ell_\# = \ell_2$ for the sake of simplicity. Given $\eta > 0$, $\mathbf{x} \in \Sigma$, and $t \in [T]$, the authors first (implicitly) note that if the random variable $Z(\mathbf{x})$ counts the number of components of $\mathbf{z}_f(\cdot)$ that are exactly t -part Lipschitz over $\mathcal{S}_\mathbf{x}(\frac{\eta}{2})$ with a given constant \bar{L}_f , then $\mathbb{E}Z \leq mp_t$. Therefore, given $c_0 > 0$ and invoking Hoeffding's inequality, they can upper bound the probability that $Z(\mathbf{x}) \geq mp_t(1 + c_0) \geq \mathbb{E}Z(\mathbf{x}) + mp_t c_0$ with

$$\mathbb{P}[Z(\mathbf{x}) \geq mp_t(1 + c_0)] \leq \exp(-2c_0^2 m).$$

From this bound (using a union bound over all $t \in [T]$), they then determine that $\mathcal{S}_{\mathbf{x}}(\frac{\eta}{2})$ is partitioned in at most $S := \exp((1 + c_0) \rho(\eta, T)m)$ cells with probability greater than $1 - Te^{-2c_0^2 m}$. Each cell of this partition of $\mathcal{S}_{\mathbf{x}}(\frac{\eta}{2})$ has thus a diameter of at most η . Moreover, by definition, \mathbf{z}_f is guaranteed to be Lipschitz continuous with constant \bar{L}_f over every such cell.

The author then consider the possibility to pick one point per such cell, called cell center, and to gather them in a finite set of at most S elements. One can then repeat this construction for all vectors \mathbf{x} of a $\frac{\eta}{2}$ -covering $\Sigma_{\eta/2}$ of Σ , and collect, for each such vector, all cell centers of its related neighborhood into a global set \mathcal{G} of centers of at most $S \times \mathcal{C}_{\frac{\eta}{2}}(\Sigma)$ elements. By definition, \mathcal{G} is thus a $\eta/2$ -covering of Σ with the additional property that \mathbf{z}_f is \bar{L}_f -Lipschitz continuous over each cell.

The authors then leverage this local continuity as follows. Since, by hypothesis, the mapping \mathbf{z}_f defined in (6) satisfies the embedding relation (25) (for $\delta = 0$) with probability exceeding $1 - Ce^{-cm\epsilon^2}$ over any *fixed* pair of vectors $\mathbf{x}, \mathbf{y} \in \Sigma$, they first expand this property over all pairs of vectors taken in $\mathcal{G} \times \mathcal{G}$. This is ensured with probability exceeding $1 - CS^2 \mathcal{C}_{\eta/2}^2(\Sigma) e^{-cm\epsilon^2}$, by union bound, since $|\mathcal{G} \times \mathcal{G}| \leq S^2 \mathcal{C}_{\eta/2}^2(\Sigma)$. Next, they extend this property to all $\mathbf{x}, \mathbf{y} \in \Sigma$ by continuity, exploiting the (local) Lipschitz continuity of f over each cell.

The flaw, which happens in the first step above, is analogous to how we cannot show the wrong statement $\mathbb{P}[\|\mathbf{g}\|^2 < 0] = 1/2$ for a Gaussian vector $\mathbf{g} \sim \mathcal{N}^d(0, 1)$ by assigning another vector \mathbf{u} to \mathbf{g} in the correct equality $\mathbb{P}[\langle \mathbf{u}, \mathbf{g} \rangle < 0] = 1/2$, valid for \mathbf{u} fixed. Indeed, the vectors of \mathcal{G} , the collection of all cell centers, are built from the random mapping \mathbf{z}_f — each center must be taken in a cell whose frontiers are controlled by the discontinuities of the components of \mathbf{z}_f . These vectors are thus dependent of both $\mathbf{\Omega} \sim \Lambda^m$ and the dither $\boldsymbol{\xi} \sim \mathcal{U}^m([0, 2\pi))$, through their dependence in \mathbf{z}_f . Therefore, one cannot ensure that the probability that (25) holds (with $\delta = 0$) on two cell centers $\mathbf{x} = \mathbf{x}(\mathbf{\Omega}, \boldsymbol{\xi}), \mathbf{y} = \mathbf{y}(\mathbf{\Omega}, \boldsymbol{\xi})$ exceeds $1 - Ce^{-cm\epsilon^2}$, since that probability is itself taken over $\mathbf{\Omega}, \boldsymbol{\xi}$. This flaw breaks the proof of [13, Thm 3.2].

5.2 An alternative geometry-preserving embedding

One can use Prop. 5 to get a variant of Thm 1 relying on the equivalence between (25) and (26). This variant achieves the same high-level goal (*i.e.*, a non-asymptotic guarantee on the approximation error achieved by the embedding \mathbf{z}_f that holds infinite signal sets even for discontinuous f), but the assumptions it relies on differ in two aspects. First, the smoothness of the distance map γ is not anymore characterized by its Lipschitz smoothness directly, but by the constant C_Λ , defined by the sampling scheme Λ driving the random projections $\mathbf{\Omega}$. Second, we use the mean Lipschitz property instead of the T -part Lipschitz property as notion of “generalized smoothness” for the map f . It is not clear if this change is fundamentally necessary to be able to prove a variant of Thm 1, but we leave an investigation of this issue for future work (the universal quantization q satisfies both properties anyway).

In fact, the following corollary shows that one can define novel *asymmetric embeddings* from Σ into \mathbb{C}^m ; we can map two vectors of Σ with different random feature mappings \mathbf{z}_f and \mathbf{z}_g achieved with distinct periodic functions f and g , respectively, and still show that, under certain conditions on f, g , and the frequency distribution Λ , $\|\mathbf{z}_f(\mathbf{x}) - \mathbf{z}_g(\mathbf{y})\|^2$ approximates a distortion of the distance between any $\mathbf{x}, \mathbf{y} \in \Sigma$ provided m is large compared to the complexity of Σ . Then, setting $f = g$ provides a specific embedding of Σ into \mathbb{C}^m , in the sense described by [13].

Corollary 2 (Asymmetric geometry-preserving embedding). *Let Σ be a compact set with finite covering number, $f, g \in \text{PF}$ be two real 2π -periodic functions and finite mean smoothness constants $L_f^\mu > 0$ and $L_g^\mu > 0$, respectively. We assume that the frequency distribution Λ is such that $C_\Lambda < \infty$, and there exists a real, one-dimensional p.d. kernel $\kappa_0^\Delta : \mathbb{R}_+ \rightarrow [0, 1]$ such that $[\mathcal{F}^{-1}\Lambda](\mathbf{u}) = \kappa_0^\Delta(\|\mathbf{u}\|_\#)$ for some norm $\|\cdot\|_\#$.*

For all error level $\epsilon > 0$, provided the feature dimension is larger than

$$m \geq 128 \cdot \frac{1}{\epsilon^2} \cdot \mathcal{H}_{\epsilon/c}(\Sigma), \quad (31)$$

with constant $c = 4C_\Lambda(L_f^\mu + L_g^\mu + 2 \min(L_f^\mu, L_g^\mu))$, we have, with probability exceeding $1 - 9 \exp(-\frac{m\epsilon^2}{64})$,

$$\gamma_{f,g}(\|\mathbf{x} - \mathbf{y}\|_\#) - 4\epsilon \leq \|\mathbf{z}_f(\mathbf{x}) - \mathbf{z}_g(\mathbf{y})\|^2 \leq \gamma_{f,g}(\|\mathbf{x} - \mathbf{y}\|_\#) + 4\epsilon, \quad \forall \mathbf{x}, \mathbf{y} \in \Sigma, \quad (32)$$

according to the distance map $\gamma_{f,g}$ defined by

$$\gamma_{f,g} : s \in \mathbb{R}_+ \rightarrow \gamma_{f,g}(s) = \|f\|^2 + \|g\|^2 - 2 \sum_{k \in \mathbb{Z}} F_k G_k^* \kappa_0^\Delta(|k|s) \in \mathbb{R}_+.$$

Therefore, if $f = g$ and if $\gamma_{f,f}$ is invertible, \mathbf{z}_f is a $(\gamma_{f,f}, 0, 4\epsilon)$ -embedding of Σ (equipped with the norm $\|\cdot\|_\#$) into \mathbb{C}^m , with $\gamma_{f,f}(0) = 0$ and $\gamma_{f,f}(s) \in [0, 2]$.

Proof. Under the hypothesis of this corollary and remembering that $\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle$, Prop. 5 tells us that the event

$$|\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) - \kappa_{f,g}(\mathbf{x}, \mathbf{y})| \leq \epsilon, \quad \forall \mathbf{x}, \mathbf{y} \in \Sigma,$$

holds with probability exceeding $1 - 3 \exp(-\frac{m\epsilon^2}{64})$. Similarly, with the same probability, $|\widehat{\kappa}_{f,f}(\mathbf{x}, \mathbf{y}) - \kappa_{f,f}(\mathbf{x}, \mathbf{y})| \leq \epsilon$ and $|\widehat{\kappa}_{g,g}(\mathbf{x}, \mathbf{y}) - \kappa_{g,g}(\mathbf{x}, \mathbf{y})| \leq \epsilon$, for all $\mathbf{x}, \mathbf{y} \in \Sigma$. Therefore, by union bound, these three events jointly hold with probability larger than $1 - 9 \exp(-\frac{m\epsilon^2}{64})$.

Conditionally to this combined occurrence, since $\kappa_{f,f}(\mathbf{x}, \mathbf{x}) = \kappa_{f,f}^\Delta(\mathbf{0}) = \|f\|^2$ and $\kappa_{g,g}(\mathbf{y}, \mathbf{y}) = \kappa_{g,g}^\Delta(\mathbf{0}) = \|g\|^2$, we find

$$|\langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_f(\mathbf{x}) \rangle - \|f\|^2| \leq \epsilon, \quad |\langle \mathbf{z}_g(\mathbf{y}), \mathbf{z}_g(\mathbf{y}) \rangle - \|g\|^2| \leq \epsilon,$$

and

$$\|\mathbf{z}_f(\mathbf{x}) - \mathbf{z}_g(\mathbf{y})\|^2 \leq \|f\|^2 + \|g\|^2 - 2\kappa_{f,g}(\mathbf{x}, \mathbf{y}) + 4\epsilon = \|f\|^2 + \|g\|^2 - 2\kappa_{f,g}^\Delta(\mathbf{x} - \mathbf{y}) + 4\epsilon.$$

Moreover, from Prop. 3, since $\kappa_{f,g}^\Delta(\mathbf{u}) = \sum_{k \in \mathbb{Z}} F_k G_k^* \kappa^\Delta(k\mathbf{u})$ with $\kappa^\Delta(\mathbf{u}) = (\mathcal{F}^{-1}\Lambda)(\mathbf{u}) = \kappa_0^\Delta(\|\mathbf{u}\|_\#)$ and $\mathbf{u} \in \mathbb{R}^d$, the definition of $\gamma_{f,g}$ provides

$$\gamma_{f,g}(\|\mathbf{u}\|_\#) = \|f\|^2 + \|g\|^2 - 2\kappa_{f,g}^\Delta(\mathbf{u}),$$

which proves the upper bound of (32), the lower bound being established similarly.

Since $f, g, \kappa_0^\Delta \in \mathbb{R}$, $\sum_k F_k G_k^* \beta_k \in \mathbb{R}$ for any real coefficients β_k , we show easily that $\gamma_{f,g} \in \mathbb{R}$ with $\gamma_{f,g}(0) = \|f\|^2 + \|g\|^2 - 2\langle f, g \rangle \geq \|f\|^2 + \|g\|^2 - 2\|f\| \|g\| \geq 0$. Moreover, if $f = g$, we get $\gamma_{f,f}(0) = 0$ since $\kappa_0^\Delta(0) = 0$, and $\gamma_{f,f}(s) \in [0, 2]$ since $0 \leq \kappa_0^\Delta(s) \leq 1$ for all $s \geq 0$ and $\sum_{k \in \mathbb{Z}} |F_k|^2 \kappa_0^\Delta(|k|s) \leq \|f\|^2$. \square

In this corollary, the existence of a norm $\|\cdot\|_\#$ controlling the behavior of $\mathcal{F}^{-1}\Lambda$ is ensured, for instance, if Λ is a centered Gaussian distribution, in which case the $\ell_\#$ -norm is the ℓ_2 -norm. If Λ is the Cartesian product of d Cauchy distributions in \mathbb{R}^d (with zero location parameter and scale parameter $\tau > 0$), *i.e.*,

$$\Lambda(\boldsymbol{\omega}) = \frac{1}{\pi^d \tau^d} \prod_{k=1}^d \frac{\tau^2}{\omega_k^2 + \tau^2}, \quad (33)$$

then $\mathcal{F}^{-1}\Lambda$ amounts to the Laplace distribution and $\|\cdot\|_\# = \|\cdot\|_1$ [13, Sec. 4.2.2.]. Moreover, if Λ is set to any α -stable distribution with $\alpha \geq 1$, *i.e.*, a distribution with characteristic function $(\mathcal{F}^{-1}\Lambda)(\mathbf{x}) \propto \exp(-c\|\mathbf{x}\|_\alpha^\alpha)$ with the Gaussian and the Cauchy distributions as special cases, we can reach an (asymmetric) embedding associated with the norm $\|\cdot\|_\alpha$ [40].

Regarding the distance map $\gamma_{f,g}$, we observe that it does not necessarily vanish at the origin, when $\mathbf{x} = \mathbf{y}$ in (32). As soon as $f \neq g$, a bias exists since

$$\gamma_{f,g}(0) = \|f\|^2 + \|g\|^2 - 2 \sum_{k \in \mathbb{Z}} F_k G_k^* = \|f\|^2 + \|g\|^2 - 2\langle f, g \rangle = \|f - g\|^2, \quad (34)$$

using $\kappa_0^\Delta(0) = \int_{\mathbb{R}^d} \Lambda(\boldsymbol{\omega}) d\boldsymbol{\omega} = 1$. For instance, if $f = q$ (with q the universal quantizer defined in (8)), and $g(\cdot) = \cos(\cdot)$, $\gamma_{q,\cos}(s) = \frac{3}{2} - 2\Re(F_1) \kappa_0^\Delta(s) = \frac{3}{2} - \frac{4}{\pi} \kappa_0^\Delta(s)$ since $\|f\|^2 = 1$, $\|g\|^2 = 1/2$, $2G_k = \delta_{k,1} + \delta_{k,-1}$, and $F_1 = F_{-1} = \frac{2}{\pi}$ from (8). Therefore, if Λ is a Gaussian distribution with unit standard deviation,

$$\|\mathbf{z}_f(\mathbf{x}) - \mathbf{z}_g(\mathbf{y})\|^2 \approx \gamma_{q,\cos}(\|\mathbf{x} - \mathbf{y}\|) = \frac{3}{2} - \frac{4}{\pi} \exp(-\frac{1}{2}\|\mathbf{x} - \mathbf{y}\|^2).$$

Compared to the case $f(\cdot) = g(\cdot) = \cos(\cdot)$ where

$$\gamma_{\cos,\cos}(\|\mathbf{x} - \mathbf{y}\|) = 1 - \exp(-\frac{1}{2}\|\mathbf{x} - \mathbf{y}\|^2),$$

and $\gamma_{\cos,\cos}(0) = 0$, we thus observe a systematic bias $\gamma_{q,\cos}(0) = \frac{3}{2} - \frac{4}{\pi} \approx 0.2268$ at the origin.

This non-vanishing bias² in the case $f \neq g$ is not a drawback per se, since $\gamma_{f,g}$ can still be invertible. For $\gamma_{q,\cos}$ and a Gaussian Λ with unit variance, we find

$$\gamma_{q,\cos}^{-1}(s') = \left(-2 \ln\left(\frac{3\pi}{8} - \frac{\pi}{4}s'\right)\right)^{1/2}, \text{ with } s' \in \left[\frac{3}{2} - \frac{4}{\pi}, \frac{3}{2}\right].$$

This shows that, if ϵ is small enough, we can still reliably infer the distance between \mathbf{x} and \mathbf{y} from $\|\mathbf{z}_f(\mathbf{x}) - \mathbf{z}_g(\mathbf{y})\|$ provided that $\mathbf{x} \approx \mathbf{y}$. Indeed, estimating $\gamma_{q,\cos}^{-1}(\|\mathbf{z}_f(\mathbf{x}) - \mathbf{z}_g(\mathbf{y})\|^2) \approx \|\mathbf{x} - \mathbf{y}\|$ leads to a first order error [13] proportional to

$$\left(\frac{d}{ds} \gamma_{q,\cos}(s)\right)^{-1} \epsilon = \frac{4}{\pi s} \epsilon \exp\left(\frac{s^2}{2}\right),$$

for $s = \|\mathbf{x} - \mathbf{y}\|$. As expected from the local nature of the embedding, this error quickly explodes when s is large.

Remark. When $f = g$, $\gamma(s) = 2\|f\|^2 - 2\sum_k |F_k|^2 \kappa_0^\Delta(|k|s)$ is invertible iff $\sum_k |F_k|^2 \kappa_0^\Delta(|k|s)$ is invertible. This occurs, for instance, if the one-dimensional kernel κ_0^Δ is differentiable with $\frac{d}{ds} \kappa_0^\Delta(s) < 0$ for all $s > 0$, which is the case of any symmetric α -stable distribution Λ for which $(\mathcal{F}^{-1}\Lambda)(\mathbf{x}) \propto \exp(-c\|\mathbf{x}\|_\alpha)$. In this case, we easily verify that $\frac{d}{ds} \gamma(s) > 0$ for $s > 0$, and γ is monotonically increasing when s increases, starting from $\gamma(0) = 0$. This ensures the injectivity of γ .

6 Semi-quantized random Fourier features

In this section, we explore one practical application of our general results from Sec. 4 by instantiating them on the semi-quantized scenario motivated in the Introduction (see Fig. 1). More precisely, we consider the asymmetric RPF setting $\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{z}_f(\mathbf{x}), \mathbf{z}_g(\mathbf{y}) \rangle$ in the particular case where: (i) one of the signals \mathbf{x} is available through its one-bit universal features $\mathbf{z}_q(\mathbf{x}) \in \{-\frac{1}{\sqrt{m}}, +\frac{1}{\sqrt{m}}\}^m$ (that is, the first periodic map f is the square wave q , alternating between ± 1 with period 2π , see Fig. 2); (ii) the other signal \mathbf{y} is available through its classical (full-precision) random Fourier features $\mathbf{z}_{\cos}(\mathbf{y})$ (that is, the second map g is a cosine).

Concretely, we start by highlighting a striking general result: when the classical RFF (*i.e.*, for which $g(\cdot) = \exp(i\cdot)$ or $\cos(\cdot)$) are combined with *any* mean smooth function $f \in \text{PF}$, then the asymmetric inner product $\widehat{\kappa}_{f,g}(\mathbf{x}, \mathbf{y})$ exactly recovers the initial kernel κ that would be approached by symmetric usual RFF $\widehat{\kappa}_{g,g}(\mathbf{x}, \mathbf{y})$. Then, to combine this fact with the binary square wave $f = q$, we prove that q is mean smooth (Def. 3). This finally allows us to obtain a probabilistic uniform bound on the kernel approximation error for the semi-quantized scenario pair, demonstrating in the process how to deal with the scaling issues that appear in such schemes by using the normalization (13).

²This bias is here demonstrated when the feature space \mathbb{C}^m is equipped with the squared ℓ_2 -distance; the question of its existence for other metrics, such as the ℓ_1 -distance, remains open.

Expected kernel with a single-frequency nonlinearity: Let us begin by noting an interesting consequence of Prop. 3. From RPF $\mathbf{z}_f(\mathbf{x})$ captured on \mathbf{x} with *any* nonlinearity $f \in \text{PF}$ whose fundamental period is exactly 2π , one can recover in expectation, for a given vector \mathbf{y} , the evaluation the shift-invariant kernel $\kappa(\mathbf{x}, \mathbf{y}) = \kappa^\Delta(\mathbf{x} - \mathbf{y})$ associated with the sampling of the projections $\omega_j \sim \Lambda = \mathcal{F}\kappa^\Delta$.

Indeed, using Prop. 3 in the complex field, and setting $g(\cdot) = \exp(i\cdot)$ for the RPF of \mathbf{y} —which in this case is the RFF (Def. 1)—ensures that $\kappa_{f, \exp(i\cdot)}(\mathbf{x}, \mathbf{y}) = F_1 \kappa(\mathbf{x}, \mathbf{y})$. Intuitively, the dithering averages out all the high-frequency components in f , leaving only its fundamental frequency. When dealing with real-valued quantities $\kappa, f \in \mathbb{R}$, we can use the real RFF (where $g(\cdot) = \Re \exp(i\cdot) = \cos(\cdot)$) instead, and using the normalized kernel (13) with (11) gives

$$\dot{\kappa}_{f, \cos}(\mathbf{x}, \mathbf{y}) = \frac{1}{\Re F_1} \kappa_{f, \cos}(\mathbf{x}, \mathbf{y}) = \frac{1}{\Re F_1} \mathbb{E} \langle f(\mathbf{\Omega}^\top \mathbf{x} + \boldsymbol{\xi}), \cos(\mathbf{\Omega}^\top \mathbf{y} + \boldsymbol{\xi}) \rangle = \kappa(\mathbf{x}, \mathbf{y}), \quad (35)$$

since $\langle f, g \rangle = \sum_k F_k G_k^* = \Re F_1$. We thus recover, through $\dot{\kappa}_{f, \cos}$ the initial kernel κ , thanks to a *rescaling* by $1/\langle f, g \rangle = 1/\Re\{F_1\}$ which must be taken into account for a fair comparison.

Remark. *In theory, we can thus recover, from $\mathbf{z}_f(\mathbf{x})$, the kernel κ at many different scales by “probing” it with $\mathbf{z}_{\cos(k\cdot)}(\mathbf{y})$ for any scale k such that $F_k \neq 0$. However, we observe in practice that the kernel approximation error quickly increases with k . This can be understood in the light of Prop. 5, since one easily show that $L_{\cos(k\cdot)} = |k|$ and³ $\frac{2}{\pi}|k| \leq L_{\cos(k\cdot)}^\mu \leq |k|$. Moreover, if $\|f\|^2 = \sum_k |F_k|^2$ is bounded, each rescaling factor $(\Re\{F_k\})^{-1}$ grows as k increases; for instance, $(\Re\{F_k\})^{-1} \propto |k|$ for $f = q$.*

The asymmetric scheme in (35) is interesting because it allows the same level of control over the approximated kernel as the usual RFF (which is an improvement compared to the scale mixture of RFF kernels imposed by Prop. 2) while still enjoying the freedom to use any type of features $\mathbf{z}_f(\mathbf{x})$ for one of the signals being compared—a particularly appealing choice being $f = q$, the one-bit universal quantization. However, in order to use Prop. 5 to obtain uniform error bounds, we still need to prove the mean Lipschitz smoothness of this (discontinuous) map.

Mean Lipschitz smoothness of universal quantization: We now show that the one-bit universal quantization function q (*i.e.*, the square wave) has the mean Lipschitz property—although it is discontinuous. The same strategy could be used to prove the mean Lipschitz smoothness of any function in PF with a finite number of discontinuities per period.

Proposition 6. *The one-bit universal quantization function q , defined in (8), has the mean Lipschitz smoothness property (Def. 3) with constant*

$$L_q^\mu = \frac{4}{\pi} \|q\|_\infty = \frac{4}{\pi}. \quad (36)$$

Proof. By definition of the mean smoothness property, we must find L_q^μ such that

$$\frac{1}{2\pi} \int_0^{2\pi} \max_{r \in [-\delta, \delta]} \{|q(t+r) - q(t)|\} dt \leq L_q^\mu \cdot \delta.$$

We start by characterizing the integrand $I_\delta(t) := \max_{|r| \leq \delta} \{|q(t+r) - q(t)|\}$. Since $q(t)$ is constant (in particular, $q(t) = \pm \|q\|_\infty$) everywhere except on discontinuities at $t \in \frac{\pi}{2} + \pi\mathbb{Z}$ where its height changes by an absolute step of $2\|q\|_\infty$, we have for $k \in \mathbb{Z}$ that (see Fig. 2)

$$I_\delta(t) = \begin{cases} 0 & \frac{\pi}{2} + k\pi + \delta < t < \frac{\pi}{2} + (k+1)\pi - \delta \\ 2\|q\|_\infty & \frac{\pi}{2} + k\pi - \delta \leq t \leq \frac{\pi}{2} + k\pi + \delta. \end{cases}$$

Integrating this over one period gives $\int_0^{2\pi} I_\delta(t) dt = \min(4\delta, 2\pi) \cdot 2\|q\|_\infty \leq 8\|q\|_\infty \cdot \delta$, *i.e.*, $L_q^\mu = \frac{4}{\pi} \|q\|_\infty$. \square

³First, $L_{\cos(k\cdot)}^\mu \leq L_{\exp(ik\cdot)}^\mu = |k|$ since $|\cos(\alpha) - \cos(\beta)| \leq |e^{i\alpha} - e^{i\beta}|$ for all $\alpha, \beta \in \mathbb{R}$. Second, from $|\cos(k(t+r)) - \cos(kt)| = 2|\sin(k(t + \frac{r}{2})) \sin \frac{kr}{2}|$ for all $t, r \in \mathbb{R}$, we get, by fixing $r = \delta$ in (19), $L_f^\mu \geq \sup_{\delta > 0} \frac{2}{\delta} |\sin \frac{k\delta}{2}| \mathbb{E}_{t \sim \mathcal{U}([0, 2\pi])} |\sin(k(t + \frac{\delta}{2}))| = \frac{2}{\pi} |k|$.

Combining quantized and cosine features: We are interested in approximating a specific kernel $\kappa(\mathbf{x}, \mathbf{y})$ by the asymmetric features product $\langle \mathbf{z}_q(\mathbf{x}), \mathbf{z}_{\cos}(\mathbf{y}) \rangle$. This product gives on average $\kappa_{q,\cos} = \frac{2}{\pi}\kappa$ (recall from (8) that $Q_k = \frac{2}{k\pi}(-1)^{(k-1)/2}$ for k odd and 0 otherwise), and the re-scaled approximation $\tilde{\kappa}_{q,\cos}$ defined in (13) in this case is given by

$$\tilde{\kappa}_{q,\cos}(\mathbf{x}, \mathbf{y}) := \frac{\pi}{2} \cdot \hat{\kappa}_{q,\cos}(\mathbf{x}, \mathbf{y}) = \frac{\pi}{2} \langle \mathbf{z}_q(\mathbf{x}), \mathbf{z}_{\cos}(\mathbf{y}) \rangle \approx \kappa(\mathbf{x}, \mathbf{y}). \quad (37)$$

We bound the error of approximating the kernel over an infinite compact set Σ thanks to Prop. 5.

Corollary 3 (Uniform kernel approximation error from quantized-complex asymmetric features). *Given $\epsilon > 0$, a compact set Σ , and the frequency distribution Λ such that $C_\Lambda < \infty$, provided that*

$$m \geq 32\pi^2 \cdot \frac{1}{\epsilon^2} \cdot \mathcal{H}_{\epsilon/((8+6\pi)C_\Lambda)}(\Sigma), \quad (38)$$

the following kernel approximation bound holds uniformly:

$$|\tilde{\kappa}_{q,\cos}(\mathbf{x}, \mathbf{y}) - \kappa(\mathbf{x}, \mathbf{y})| \leq \epsilon, \quad \forall \mathbf{x}, \mathbf{y} \in \Sigma, \quad (39)$$

with probability exceeding $1 - 3 \exp(-\frac{m\epsilon^2}{16\pi^2})$

Proof. Apply Prop. 5 with $f = q$, $g = \cos$, using that $L_q^\mu = \frac{4}{\pi}$ (from Prop. 6) and $L_{\cos}^\mu = 1$: for any given $\epsilon' > 0$, if $m \geq 128 \cdot \frac{1}{\epsilon'^2} \cdot \mathcal{H}_{\epsilon'/c}(\Sigma)$ with $c = (12 + \frac{16}{\pi})C_\Lambda$,

$$\mathbb{P}[\exists \mathbf{x}, \mathbf{y} \in \Sigma : |\hat{\kappa}_{q,\cos}(\mathbf{x}, \mathbf{y}) - \frac{2}{\pi}\kappa(\mathbf{x}, \mathbf{y})| > \epsilon'] \leq 3 \exp\left(-\frac{m\epsilon'^2}{64}\right).$$

To take into account the scaling of the kernel, set $\epsilon = \frac{\epsilon'}{\Re Q_1} = \frac{\pi}{2}\epsilon'$. □

Example 4. *Consider a final time our example of a union of S s -dimensional subspaces (see Ex. 1) combined with the Gaussian kernel with unit bandwidth (and $C_\Lambda = 1$). In this case, the kernel approximation error $|\tilde{\kappa}_{q,\cos}(\mathbf{x}, \mathbf{y}) - \kappa(\mathbf{x}, \mathbf{y})|$ is uniformly bounded over Σ , with high probability, provided that the number of features satisfies $m \geq C\epsilon^{-2}(s \log(\frac{8+6\pi}{\epsilon}) + \log S)$, which reduces to $m \geq C\epsilon^{-2}s \log(\frac{(8+6\pi)\epsilon d}{s\epsilon})$ for bounded s -sparse signals.*

Corollary 3 provides a theoretical guarantee justifying the semi-quantized scheme presented in the Introduction. In the next section, we further validate this approach from numerical simulations.

7 Experiments

In all our experiments, we are interested in approximating a kernel $\kappa(\mathbf{x}, \mathbf{y})$, associated with the RFF sampled with Λ^m , by the inner product of random periodic features. We focus on (combinations of) the two types of features discussed in the previous section: the “real” random Fourier features $\mathbf{z}_{\cos}(\mathbf{x}) = \frac{1}{\sqrt{m}} \cos(\mathbf{\Omega}^\top \mathbf{x} + \boldsymbol{\xi}) \in \mathbb{R}^m$, and the universal features $\mathbf{z}_q(\mathbf{x}) = \frac{1}{\sqrt{m}} q(\mathbf{\Omega}^\top \mathbf{x} + \boldsymbol{\xi}) \in \{-\frac{1}{\sqrt{m}}, +\frac{1}{\sqrt{m}}\}^m$, where m is the number of features (or dimension), and where we generate $\mathbf{\Omega} \sim \Lambda^m$ and $\boldsymbol{\xi} \sim \mathcal{U}^m([0, 2\pi))$. Recalling the rescaling (13) for fair comparisons of the approximated kernels with κ , we thus consider three possible combinations: the classical (real) random Fourier features (with $\|f\|^2 = \|\cos(\cdot)\|^2 = \frac{1}{2}$),

$$\tilde{\kappa}_{\cos,\cos}(\mathbf{x}, \mathbf{y}) = 2 \langle \mathbf{z}_{\cos}(\mathbf{x}), \mathbf{z}_{\cos}(\mathbf{y}) \rangle \approx \kappa(\mathbf{x}, \mathbf{y}),$$

our asymmetric “semi-quantized” scheme (with $\langle f, g \rangle = 2/\pi$),

$$\tilde{\kappa}_{q,\cos}(\mathbf{x}, \mathbf{y}) = \frac{\pi}{2} \langle \mathbf{z}_q(\mathbf{x}), \mathbf{z}_{\cos}(\mathbf{y}) \rangle \approx \kappa(\mathbf{x}, \mathbf{y}),$$

and the fully quantized inner product from [7] (with $\|q\|^2 = \sum_k |Q_k|^2 = 1$),

$$\tilde{\kappa}_{q,q}(\mathbf{x}, \mathbf{y}) = \langle \mathbf{z}_q(\mathbf{x}), \mathbf{z}_q(\mathbf{y}) \rangle \approx \kappa_{q,q}(\mathbf{x}, \mathbf{y}) = \sum_{k \in \mathbb{Z}} |Q_k|^2 \kappa(k\mathbf{x}, k\mathbf{y}) \neq \kappa(\mathbf{x}, \mathbf{y}).$$

7.1 Qualitative analysis of the expected kernel

As a first experiment, we visually demonstrate that our asymmetric product $\tilde{\kappa}_{q,\text{cos}}$ indeed approaches a target kernel κ . As target, we use the Gaussian kernel $\kappa(\mathbf{x}, \mathbf{y}) = \exp(-\frac{\|\mathbf{x}-\mathbf{y}\|_2^2}{2\sigma^2})$ (for which Λ is the Gaussian distribution $\mathcal{N}(\mathbf{0}, \sigma^{-2}\mathbf{I}_d)$), as well as the Laplace kernel $\kappa(\mathbf{x}, \mathbf{y}) = \exp(-\frac{\|\mathbf{x}-\mathbf{y}\|_1}{\tau})$ (where Λ is the Cauchy distribution (33)), both in dimension $d = 5$.

We evaluate the three inner products $\tilde{\kappa}_{\text{cos,cos}}$, $\tilde{\kappa}_{q,\text{cos}}$ and $\tilde{\kappa}_{q,q}$ on $n = 2000$ pairs of vectors $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, that are generated as follows. We first sample $\mathbf{x}_i \in \mathbb{R}^5$ according to a standard normal distribution, then pick $\mathbf{y}_i = \mathbf{x}_i + \lambda_i \mathbf{u}_i$, where \mathbf{u}_i is a randomly chosen unit vector (*i.e.*, normalized such that $\|\mathbf{u}_i\|_p = 1$ with $p = 2$ for the Gaussian kernel and $p = 1$ for the Laplace one), and $\lambda_i = \frac{(i-1)\lambda_{\max}}{(n-1)}$ is a controlled distance which is incremented for each pair, linearly increasing from 0 to $\lambda_{\max} = 5$. This ensures that we test the kernel approximations uniformly in the desired range of distances $\|\mathbf{x} - \mathbf{y}\|_p$.

We then generate one realization of $\mathbf{\Omega}, \mathbf{\xi}$ (with $m = 200$ for the Gaussian kernel, and $m = 2000$ for the Laplace kernel, values which were arbitrarily chosen to get pleasing visualizations), which we use to compute the real RFF $\{(\mathbf{z}_{\text{cos}}(\mathbf{x}_i), \mathbf{z}_{\text{cos}}(\mathbf{y}_i))\}_{i=1}^n$ as well as the universal quantization features $\{(\mathbf{z}_q(\mathbf{x}_i), \mathbf{z}_q(\mathbf{y}_i))\}_{i=1}^n$, from which we get n evaluations of the classical RFF inner product $\tilde{\kappa}_{\text{cos,cos}}(\mathbf{x}_i, \mathbf{y}_i) = 2\langle \mathbf{z}_{\text{cos}}(\mathbf{x}_i), \mathbf{z}_{\text{cos}}(\mathbf{y}_i) \rangle$, the asymmetric product $\tilde{\kappa}_{q,\text{cos}}(\mathbf{x}_i, \mathbf{y}_i) = \frac{\pi}{2}\langle \mathbf{z}_q(\mathbf{x}_i), \mathbf{z}_{\text{cos}}(\mathbf{y}_i) \rangle$, and the fully quantized product $\tilde{\kappa}_{q,q}(\mathbf{x}_i, \mathbf{y}_i) = \langle \mathbf{z}_q(\mathbf{x}_i), \mathbf{z}_q(\mathbf{y}_i) \rangle$.

Those evaluations are shown as black dots in Fig. 3 for the Gaussian and Laplace kernels in the top and bottom rows, respectively. As predicted by the theory, both the RFF product $\tilde{\kappa}_{\text{cos,cos}}$ and our semi-quantized product $\tilde{\kappa}_{q,\text{cos}}$ concentrate around the target kernel κ (in red). As expected from [13], in the fully quantized case the product $\tilde{\kappa}_{q,q}$ rather concentrates around a different “distorted” kernel, $\kappa_{q,q}$. Note that we increased the feature space dimension m tenfold for the Cauchy kernel, which reduced the variance of the approximation. However, it is difficult to notice a substantial difference of approximation quality between the plain RFF $\tilde{\kappa}_{\text{cos,cos}}$ and the semi-quantized asymmetric scheme $\tilde{\kappa}_{q,\text{cos}}$. We thus perform a more quantitative exploration of the error $|\tilde{\kappa}_{q,\text{cos}} - \kappa|$ in the next experiment.

7.2 Quantitative analysis of the approximation error

To perform a more quantitative analysis of the kernel approximation from Cor. 3, we perform another set of experiments that highlight the evolution of the worst-case error (associated with the hybrid estimation (37)),

$$\epsilon_{q,\text{cos}}(\Sigma) := \sup_{\mathbf{x}, \mathbf{y} \in \Sigma} |\tilde{\kappa}_{q,\text{cos}}(\mathbf{x}, \mathbf{y}) - \kappa(\mathbf{x}, \mathbf{y})|,$$

as a function of m . In this synthetic experiment, we work with a finite set of signals $\Sigma = \{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$ obtained from a Gaussian distribution $\mathbf{x}_i \sim_{\text{i.i.d.}} \mathcal{N}(\mathbf{0}, \tilde{\sigma}^2 \mathbf{I}_d)$ in with $\tilde{\sigma} = 10$. We target a Gaussian kernel κ of bandwidth $\sigma = 0.25$, and evaluate the absolute approximation error $\epsilon_{q,\text{cos}}(\Sigma)$ over all vector pairs of Σ . We record the largest error encountered this way, and repeat this process for several feature dimensions m .

First, we let $n = |\Sigma|$, the number of signals, vary between 10 and 500 (by sampling 21 equally-spaced values for $\log_{10}(n)$), and generate a new dataset in dimension $d = 32$ each time. For each value of m (varying uniformly between 100 and 1300), we repeat 50 independent draws of $\mathbf{\Omega}$ and $\mathbf{\xi}$ and report Fig. 4a the number of times that $\epsilon_{q,\text{cos}}(\Sigma) \leq \bar{\epsilon}$ for a fixed threshold $\bar{\epsilon} = 0.15$ (*i.e.*, we report the empirical “success rate” of the embedding). As expected, the feature space dimension m needed to succeed (highlighted in red for 50% success rate) scales as $O(\log n)$. We also show in dashed yellow the same transition for the worst-case error $\epsilon_{\text{cos,cos}}(\Sigma)$ (evaluating $|\tilde{\kappa}_{\text{cos,cos}} - \kappa|$ over all vector pairs of Σ) committed by the plain RFF, which shows the price to pay for quantization. Roughly speaking, the same success rate is achieved for $\tilde{\kappa}_{q,\text{cos}}(\mathbf{x}_i, \mathbf{x}_j)$ as for $\tilde{\kappa}_{\text{cos,cos}}(\mathbf{x}_i, \mathbf{x}_j)$ provided we take $\sim 33\%$ more random features, which still corresponds to a bitrate reduction for the features of the first signal \mathbf{x}_i . Finally, for the sake of




Figure 3: Comparison between the target kernel $\kappa(\mathbf{x}, \mathbf{y})$ (red curves), and the approximations (the black scatter plots each evaluated over $n = 200$ pairs $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ using, for (a,d), the plain random Fourier features $\tilde{\kappa}_{\text{cos,cos}}(\mathbf{x}_i, \mathbf{y}_i) = 2\langle \mathbf{z}_{\text{cos}}(\mathbf{x}_i), \mathbf{z}_{\text{cos}}(\mathbf{y}_i) \rangle$, for (b,e), our asymmetric cosine-quantized pair $\tilde{\kappa}_{q,\text{cos}}(\mathbf{x}_i, \mathbf{y}_i) = \frac{\pi}{2}\langle \mathbf{z}_q(\mathbf{x}_i), \mathbf{z}_{\text{cos}}(\mathbf{y}_i) \rangle$, and for (c,f), only quantized features $\tilde{\kappa}_{q,q}(\mathbf{x}_i, \mathbf{y}_i) = \langle \mathbf{z}_q(\mathbf{x}_i), \mathbf{z}_q(\mathbf{y}_i) \rangle$. In the last case, the “distorted” expected kernel $\mathbb{E} \tilde{\kappa}_{q,q} = \kappa_{q,q}$ is shown in blue. For (a-c; top row), the comparison is made for the target Gaussian kernel $\kappa(\mathbf{x}, \mathbf{y}) = \exp(-\frac{\|\mathbf{x}-\mathbf{y}\|_2^2}{2\sigma^2})$ with scale $\sigma = 1.5$, and the approximated kernels use $m = 200$ random features evaluated. For (d-f; bottom row), the target kernel is the Laplace kernel $\kappa(\mathbf{x}, \mathbf{y}) = \exp(-\frac{\|\mathbf{x}-\mathbf{y}\|_1}{\tau})$ with scale $\tau = 1.5$, and its different approximations are set with $m = 2000$.

comparison we also show in blue the same success rate but when measuring the proximity error between the two approximations (semi-quantized and usual RFF), *i.e.*, $|\tilde{\kappa}_{q,\text{cos}} - \kappa_{\text{cos,cos}}|$, which relates to our bound in Cor. 1.

Second, we fix one single dataset Σ of $n = 200$ signals in \mathbb{R}^5 , but record the precise value of the worst-case error $\epsilon := \epsilon_{q,\text{cos}}(\Sigma)$ for each of the 50 draw of Ω, ξ at different values of m (this time varying along a logarithmic scale). We display the various errors ϵ obtained as box-plots in Fig. 4b. As can be seen by comparison with the $-1/2$ slope in red, the error is controlled with high probability (discarding the outliers from the box-plots) provided that $m = O(\epsilon^{-2})$, as expected from Prop. 5.

7.3 Application: semi-quantized support vector machines

As a last experiment, we demonstrate how the asymmetric features can be used in practice, for the particular case of Support Vector Machine classification [8, 9], where the goal is to assign a class label $y' \in \mathbb{Z}$ to new query vectors $\mathbf{x}' \in \Sigma$ from labeled training data $\mathcal{T} := \{(\mathbf{x}_i, y_i)\}_{i=1}^n$. In the binary classification case (labels $y_i \in \{\pm 1\}$), given a kernel $\kappa^\sharp : \Sigma \times \Sigma \rightarrow \mathbb{R}$, the learned SVM classifier θ predicts the class of an incoming vector \mathbf{x}' as

$$\theta(\mathbf{x}') = \text{sign} \left(\sum_{i \in \mathcal{S}^*} \alpha_i y_i \kappa^\sharp(\mathbf{x}', \mathbf{x}_i) + b \right), \quad (40)$$

where $\mathcal{S}^* \subset [n]$ is the index set of support vectors \mathbf{x}_i , $0 < \alpha_i \leq R$ are the related weights, and b is a bias (or “intercept”) term. The quantities $\{\mathcal{S}^*, \alpha_i, b\}$ are the parameters to be learned during the training

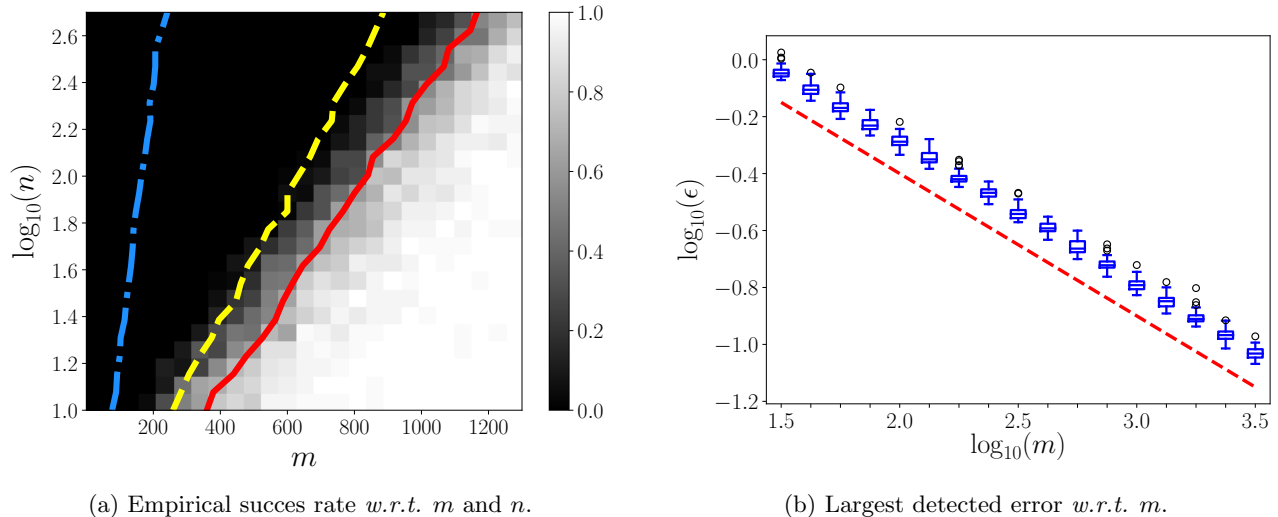


Figure 4: (a) Empirical success rate (from 100% success in white to 0% in black) of the kernel approximation (defining success as $\epsilon_{q,\text{cos}}(\Sigma) < \bar{\epsilon} = 0.15$), as a function of m (log scale) for varying dataset size $n = |\Sigma|$. The transition to 50% success or more is highlighted in solid red; the same curve is shown for the success rate of the classical RFF (when $\epsilon_{\text{cos},\text{cos}}(\Sigma) < \bar{\epsilon}$) in dashed yellow. The blue line represents the success rate related to the proximity between those two kernel approximations, *i.e.*, when $\sup_{\mathbf{x}, \mathbf{y} \in \Sigma} |\kappa_{q,\text{cos}}(\mathbf{x}, \mathbf{y}) - \kappa_{\text{cos},\text{cos}}(\mathbf{x}, \mathbf{y})| < \bar{\epsilon}$. (b) Largest kernel approximation error $\epsilon := \epsilon_{q,\text{cos}}(\Sigma)$ as a function of m for 50 draws of Ω and ξ (the blue box-plots). The dashed red line shows the slope $\log_{10}(\epsilon) \sim -\frac{1}{2} \log_{10}(m)$, for reference.

stage, while the kernel κ^\sharp and regularization strength $R > 0$ (where a smaller R corresponds to more regularization) are hyper-parameters to be set beforehand. In the multi-class case (where $y_i \in [N]$ for N classes), we use the “one-versus-rest⁴” strategy where one binary classifier is trained to recognize each class.

In this experiment, given a simple classification task described below, we propose to train the SVM with a given kernel κ_L^\sharp , and to test the classification of new samples with another kernel κ_T^\sharp that approximates κ_L^\sharp , hence assessing how the classifier θ is impacted by this modification. We consider two options. In the first we train a kernel SVM on the raw data \mathcal{T} with a “true” kernel $\kappa_L^\sharp = \kappa$ and use the approximated kernels provided by random periodic features (setting κ_T^\sharp to the kernels $\tilde{\kappa}_{\text{cos},\text{cos}}$, $\tilde{\kappa}_{q,\text{cos}}$ and $\tilde{\kappa}_{q,q}$ defined from (13)) *only at the inference stage*. In this mode, which is the viewpoint we adopted in most of this work (*e.g.*, in Prop. 5), we thus interpret the RPF inner products as a means to approximate as well as possible the given kernel κ .

In a second case, we directly train a linear SVM on the cosine random Fourier features of the training set $\mathcal{T}' := \{(\mathbf{z}_{\text{cos}}(\mathbf{x}_i), y_i)\}_{i=1}^n$, which amounts to using $\kappa_L^\sharp = \tilde{\kappa}_{\text{cos},\text{cos}}$ in (40) as the reference kernel during training. At the testing stage, we still set κ_T^\sharp to $\tilde{\kappa}_{\text{cos},\text{cos}}$, $\tilde{\kappa}_{q,\text{cos}}$ and $\tilde{\kappa}_{q,q}$. In this scenario, the random periodic features are rather (implicitly) used to define a specific kernel $\tilde{\kappa}_{\text{cos},\text{cos}}$ that generalizes as well as possible without caring about the approximation $\tilde{\kappa}_{\text{cos},\text{cos}} \approx \kappa$; this view is more faithful to recent research on the generalization capabilities of learning from RFF [14, 41–43]. Our other RPF products used at the test ($\tilde{\kappa}_{q,\text{cos}}$ and $\tilde{\kappa}_{q,q}$) are then to be understood as approximations to $\kappa_L^\sharp = \tilde{\kappa}_{\text{cos},\text{cos}}$ rather than to the original κ , as explained in Cor. 1, and as measured by the blue curve in Fig. 4a.

Synthetic data

Specifically, for both contexts, we generate a synthetic dataset of 10 000 samples in \mathbb{R}^2 by generating a mixture of 4 Gaussians for each of the $N = 5$ different classes, separated into $n = 8000$ training

⁴also known as “one-versus-all”.

Table 1: Scenarios and color coding for Fig. 5.

Color	RPF for		Kernel at the test: $\kappa_{\mathbb{T}}^{\sharp}(\mathbf{x}', \mathbf{x}_i)$
	Query vector \mathbf{x}'	Support vectors $\{\mathbf{x}_i\}_{i \in \mathcal{S}^*}$	
black	$\mathbf{z}_{\cos}(\mathbf{x}')$	$\mathbf{z}_{\cos}(\mathbf{x}_i)$	$\tilde{\kappa}_{\cos, \cos}(\mathbf{x}', \mathbf{x}_i) = 2 \langle \mathbf{z}_{\cos}(\mathbf{x}'), \mathbf{z}_{\cos}(\mathbf{x}_i) \rangle \approx \kappa(\mathbf{x}', \mathbf{x}_i)$
green	$\mathbf{z}_q(\mathbf{x}')$	$\mathbf{z}_{\cos}(\mathbf{x}_i)$	$\tilde{\kappa}_{q, \cos}(\mathbf{x}', \mathbf{x}_i) = \frac{\pi}{2} \langle \mathbf{z}_q(\mathbf{x}'), \mathbf{z}_{\cos}(\mathbf{x}_i) \rangle \approx \kappa(\mathbf{x}', \mathbf{x}_i)$
blue	$\mathbf{z}_{\cos}(\mathbf{x}')$	$\mathbf{z}_q(\mathbf{x}_i)$	$\tilde{\kappa}_{\cos, q}(\mathbf{x}', \mathbf{x}_i) = \frac{\pi}{2} \langle \mathbf{z}_{\cos}(\mathbf{x}'), \mathbf{z}_q(\mathbf{x}_i) \rangle \approx \kappa(\mathbf{x}', \mathbf{x}_i)$
red	$\mathbf{z}_q(\mathbf{x}')$	$\mathbf{z}_q(\mathbf{x}_i)$	$\tilde{\kappa}_{q, q}(\mathbf{x}', \mathbf{x}_i) = \langle \mathbf{z}_q(\mathbf{x}'), \mathbf{z}_q(\mathbf{x}_i) \rangle$

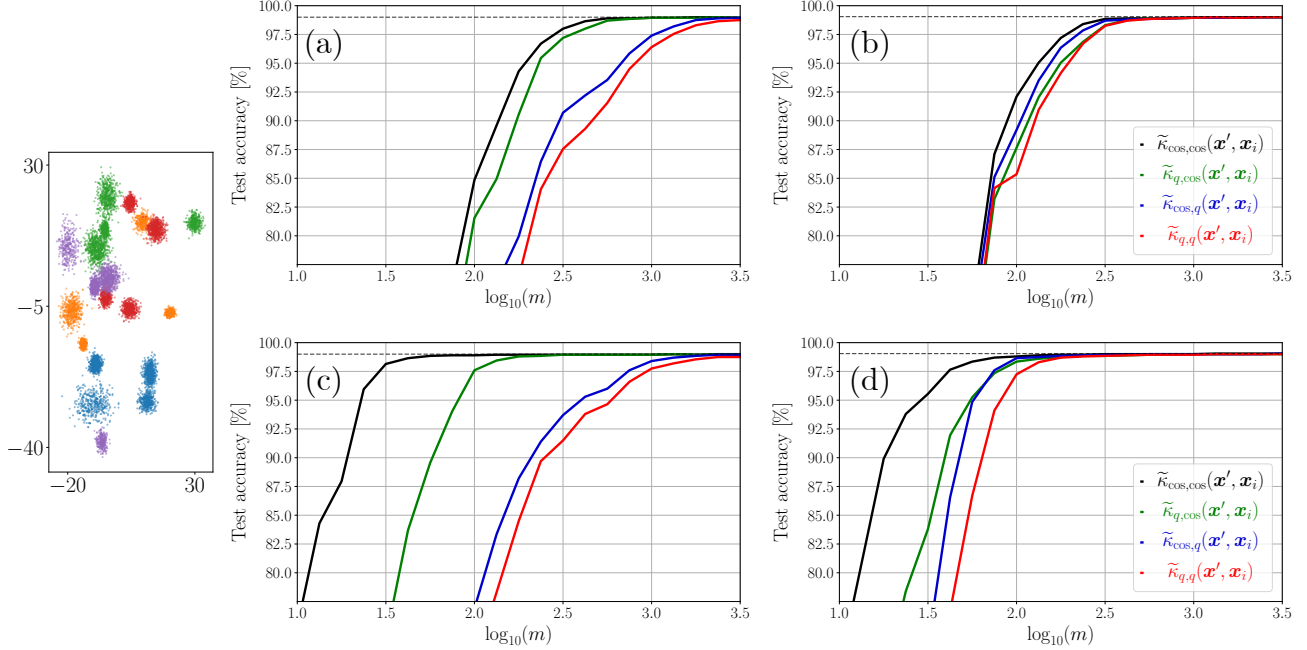


Figure 5: Left: Considered dataset (each arbitrary color corresponds to one of the 5 classes). Right: Test accuracy of an SVM classifier learned with a Gaussian kernel on the raw training set \mathcal{T} (top (a-b)) or on their cosine RFF (bottom (c-d)), for regularization parameters $R = 5$ (weak regularization; left (a-c)) and $R = 0.25$ (strong regularization; right (b-d)), and evaluated with several RPF combinations as a function of their dimension m . Classification performance is measured on the test set according to the scenarios and curve colors described in Table 1. The curves are the median out of 25 independent draws of Ω and ξ . The dashed line indicates the test accuracy of the exact SVM classifier (no random features are used).

and 2000 testing samples (see Fig. 5, left). Regarding the true kernel, we set it to a Gaussian kernel $\kappa(\mathbf{x}, \mathbf{x}') = \kappa^\Delta(\mathbf{x} - \mathbf{x}') = \exp(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2\sigma^2})$ with bandwidth $\sigma = 2$, and fixed the regularization R either to 5.0 (mild regularization) or 0.25 (strong regularization). For various feature space dimensions m , we generate the projections $\Omega \sim \Lambda^m$ (with $\Lambda = \mathcal{F}\kappa^\Delta$, see Sec. 2) and dithering $\xi \sim \mathcal{U}^m([0, 2\pi))$, and thus train (with Scikit-learn [44]) both an SVM classifier from the raw data with the Gaussian kernel, and another linear SVM from the associated cosine random Fourier features. We then evaluate these classifiers (in “inference mode”) on the separate test set, using the different random features inner products and report the median accuracy (out of 25 draws) as a function of m in Fig. 5, right. In the four plots of this figure, we use a specific color coding of the RPF for both the incoming query vector \mathbf{x}' and the learned SVM support vectors $\{\mathbf{x}_i\}_{i \in \mathcal{S}^*}$, as summarized in Table 1 for convenience. Note that, according to this table, the green and the blue curves are associated with the scenarios represented in the Introduction in Fig. 1a and Fig. 1b, respectively, and the red curves relate to the (symmetric) quantized approach of [3].

When approximating the “exact” SVM classifier (where $\kappa_{\mathbb{L}}^{\sharp} = \kappa$; top row in Fig. 5), a substantial number of features is required to reach the same performances ($m \approx 300$ to obtain accuracy $> 97.5\%$). As could be intuitively expected, the drop of accuracy is larger when more quantization of the features

is being performed (the price to pay is particularly high when the support vectors are quantized, in blue and red). This difference is probably related to the fact that the SVM decision function is relatively sensitive to the position of the support vectors (in the feature space), because they directly lie on the decision boundary. Notice also the importance of regularization: although changing R does not incur a noticeable change on the exact SVM accuracy (the horizontal dashed lines), it appears in Fig. 5b that stronger regularization improves the RPF-based classifiers.

When the SVM is directly trained on the random Fourier features (*i.e.*, $\kappa_{\mathbf{L}}^{\sharp} = \tilde{\kappa}_{\text{cos,cos}}$; bottom row in Fig. 5) the accuracy of $\tilde{\kappa}_{\text{cos,cos}}(\mathbf{x}', \mathbf{x}_i)$ (black) is strongly boosted (reaching 97.5% accuracy or higher with less than $m = 50$ features). The role of regularization is here exacerbated: reducing R hurts the performances of the plain RFF classifier, but is necessary to maintain good accuracy with the semi-quantized schemes for a reasonable feature dimension m . It appears that proper regularization is needed to account for the quantization noise.

Finally, note that in all cases, the fully quantized scheme (in red) is not much worse than our semi-quantized solution with dataset quantization (the kernel mismatch shown Fig. 3c is apparently not too harmful in this case), but still suffers from strictly more classification errors. Therefore, it can be replaced by one of the asymmetric schemes for a negligible cost, it should be done. However, because we focus on approximating an imposed kernel without considerations for the underlying machine learning model, we did not compare to the fully-quantized case where the quantized features are already used during the training [3], *i.e.*, where the distorted kernel $\kappa_{q,q} = \mathbb{E}\tilde{\kappa}_{q,q}$ is directly embraced to train the SVM rather than being used as an approximation for κ .

Real data: remote classification of hyperspectral pixels

In a last experiment, to prove the concept of using semi-quantized features in a concrete and practical setting, we consider the problem of hyperspectral pixel classification, *i.e.*, determine the class of a spatial pixel given its electromagnetic spectral response across d wavelengths. Kernel SVMs have been a quite popular solution to this challenge: our approach in particular is inspired by [45, 46] as well as [47] for the use of RFF, but many more references can be found in the extensive review [48].

To have a concrete and quantifiable measure of the computational gains allowed by the quantization of RFF, we focus on the quantized query context (illustrated Fig. 1a). More precisely, we consider the scenario of an aircraft (or satellite) equipped with a hyperspectral sensor that must send its readings \mathbf{x}' for remote classification of the pixels it observes. We assume this task is entrusted to a kernel SVM, involving a weighted sum of $\kappa_{\mathbf{T}}^{\sharp}(\mathbf{x}', \mathbf{x}_i)$ terms. Since the communication link between the satellite and the remote server is presumably costly, it is important that the number of bits used to encode this query, noted b , is as small as possible.

We compare three strategies. First, the baseline strategy is to send the “raw” measurements $\mathbf{x}' \in \mathbb{R}^d$, which requires $b = Bd$ bits, where B designates the bit-depth of full-precision readings (we consider $B = 64$ bits in our experiments). In this strategy, the kernel in the learning and testing stages is the “true” Gaussian kernel, *i.e.*, $\kappa_{\mathbf{L}}^{\sharp} = \kappa_{\mathbf{T}}^{\sharp} = \kappa$. Second, the usual RFF strategy is to send the full-precision RFF $\mathbf{z}_{\text{cos}}(\mathbf{x}') \in \mathbb{R}^m$, which requires $b = Bm$ bits. Following the observations from the previous experiments, we both learn and test on the RFF kernel, *i.e.*, $\kappa_{\mathbf{L}}^{\sharp} = \kappa_{\mathbf{T}}^{\sharp} = \tilde{\kappa}_{\text{cos,cos}}$. Finally, the quantized RFF query strategy is to send the quantized RFF $\mathbf{z}_q(\mathbf{x}') \in \{-\frac{1}{\sqrt{m}}, +\frac{1}{\sqrt{m}}\}^m$, which takes up $b = m$ bits⁵. Although the kernel at test time is now the hybrid product $\kappa_{\mathbf{T}}^{\sharp} = \tilde{\kappa}_{q,\text{cos}}$, we still learn using the usual RFF kernel $\kappa_{\mathbf{L}}^{\sharp} = \tilde{\kappa}_{\text{cos,cos}}$. Note that to simplify the comparison, we thus consider only a naive encoding of the query, neglecting the use of *e.g.*, entropy coding strategies.

⁵Strictly speaking, we would actually need to transmit $\frac{1}{2}(\sqrt{m}\mathbf{z}_q(\mathbf{x}') + 1) \in \{0, 1\}^m$ and to remotely recover $\mathbf{z}_q(\mathbf{x}')$ from this binary stream, assuming m is known to the receiver.

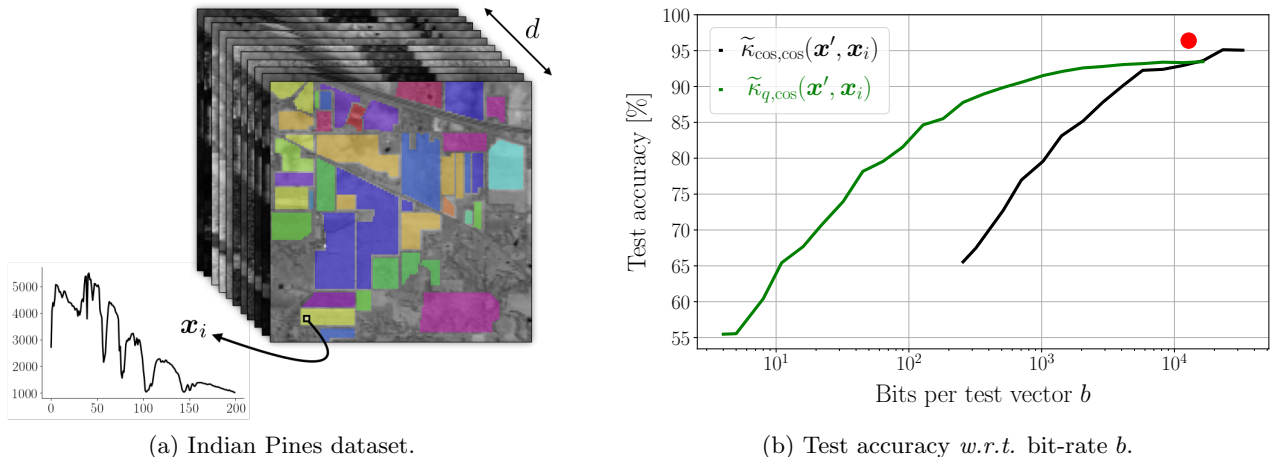


Figure 6: (a) Schematic representation of the Indian Pines hyperspectral volume, containing 10 249 labeled pixels $\mathbf{x}_i \in \mathbb{R}^d$, which feature $d = 200$ wavelengths⁷. (b) Test accuracy as a function of the number of bits b to be transmitted to the remote server, for each of the three considered strategies: the baseline strategy (red dot) sends the raw test vector \mathbf{x}' ; the RFF strategy (solid black) sends the full-precision RFF $\mathbf{z}_{\text{cos}}(\mathbf{x}')$ of varying size m ; and the quantized RFF strategy (solid green) sends the one-bit equivalent $\mathbf{z}_q(\mathbf{x}')$. The curves are the median out of 30 independent draws of Ω and ξ .

We use the standard Indian Pines dataset [49], a hyperspectral volume which contains 10 249 labeled pixels⁶, measured across $d = 200$ wavelengths⁷, separated into $N = 16$ classes (see Fig. 6a). We first separated 20% of those pixels into testing set, which left $n = 8204$ pixels for training. In order to select the hyper-parameters (kernel bandwidth σ and regularization strength R), following their sensitivity observed in the previous experiment, we performed a separate cross-validation (with 5 folds from the training set) for each of the three individual strategies. We then evaluated the test set accuracy reached by each strategy, while letting m vary for the two RFF-based strategies, and report the results Fig. 6b. In this figure, the baseline strategy is represented by the red dot, and the usual (resp. quantized) RFF strategies are represented by the black (resp. green) solid curves, which are obtained by varying m .

The baseline (red) achieves the best accuracy overall, but at the price of a quite substantial bandwidth usage. When using full-precision random features (black), the accuracy is only slightly reduced, but only if a relatively large number of random features m is used, which does bring substantial bitrate reduction. Indeed, to reduce the bandwidth b by, say, an order of magnitude, the full-precision RFF strategy must sacrifice more than 10% accuracy, which is probably not acceptable in practice. On the other hand, with the quantized query RFF strategy, we are able to achieve this same bitrate reduction by an order of magnitude at the cost of only about 4%, which sounds more reasonable. Overall, (keeping in mind that more involved compression methods could be applied to transmit the raw measurement \mathbf{x}' in our scenario above, and hence still apply the first classification strategy after decompression) the quantized RFF strategy performs better whenever the bitrate b is significantly smaller than the baseline bitrate, hence showing the potential of the approach.

8 Conclusion

We introduced the framework of asymmetric random periodic features, where random projections are passed through two different periodic maps, and whose inner products are used to approximate a kernel.

⁶The size of the full volume is 145×145 which gives 21025 pixels in total, but many of them are unlabeled, which we discard for this experiment.

⁷The initial volume contains 220 wavelengths, but following the workflow commonly adopted with this dataset, we removed the water absorption bands (*i.e.*, the spectral indices [104-108], [150-163], and 220 from the initial dataset).

We provided an expression of this kernel, with a uniform error bound holding on infinite compact signal sets, provided the periodic maps satisfy a property we called the *mean smoothness*. The mean smoothness holds for some discontinuous maps such as the one-bit universal quantization (a square wave). As a first (theoretical) application of those developments, we generalized the local geometry-preserving embeddings from [13], and corrected an error in their main result in the process. For a second more practical application, we studied (theoretically and empirically) semi-quantized kernel approximations, and showed how the impact of quantization can be controlled.

As highlighted by our last experiments, these theoretical guarantees do not necessarily ensure an accurate control over the generalization performances in a machine learning context. Indeed, it seems crucial to incorporate the random periodic features directly into the training stage, and to anticipate that some features might be quantized later (for example, when picking the regularization strength). We leave for future work the question of efficiently incorporating the asymmetric random periodic features strategy during the training stage.

In the two applicative scenarii we proposed in the Introduction (Fig. 1), the ultimate objective is to improve the bitrate-kernel accuracy trade-off. Allowing a finer (but still coarse) quantization of the RFF (*i.e.*, coding each entry on $b > 1$ bits instead of 1 as we proposed, for example with the b -bit universal quantization [13]) is a promising idea to reach a better trade-off. Since the mean smoothness of these finer quantization functions can also be verified, our results would carry over easily to this multi-bit quantization of the RFF, but a detailed analysis of this approach is also left for future work. More generally, the RPF with other periodic maps might also be worth studying. Finally, let us mention that our results could be used to obtain formal guarantees in the context of compressive learning [20, 21], where the dataset sketch (which pools the RPF of each dataset sample; see Sec. 1.1) and the algorithm estimating the dataset distribution (such as CLOMPR [21]) relies on asymmetric periodic functions, *e.g.*, the universal quantizer and the cosine function [22].

References

- [1] Simon Foucart and Holger Rauhut. A mathematical introduction to compressive sensing. *Bull. Am. Math.*, 54:151–165, 2017.
- [2] Ali Rahimi and Benjamin Recht. Random Features for Large-Scale Kernel Machines. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 1177–1184. Curran Associates, Inc., 2008.
- [3] Petros T Boufounos, Laurent Jacques, Felix Krahmer, and Rayan Saab. Quantization and compressive sensing. In *Compressed sensing and its applications*, pages 193–237. Springer, 2015.
- [4] C Sinan Güntürk, Mark Lammers, Alexander M Powell, Rayan Saab, and Ö Yılmaz. Sobolev duals for random frames and $\sigma\delta$ quantization of compressed sensing measurements. *Foundations of Computational mathematics*, 13(1):1–36, 2013.
- [5] Sjoerd Dirksen. Quantized compressed sensing: a survey. In *Compressed Sensing and Its Applications*, pages 67–95. Springer, 2019.
- [6] Petros T Boufounos and Hassan Mansour. Universal embeddings for kernel machine classification. In *2015 International Conference on Sampling Theory and Applications (SampTA)*, pages 307–311. IEEE, 2015.
- [7] Petros T. Boufounos and Shantanu Rane. Efficient Coding of Signal Distances Using Universal Quantized Embeddings. In *2013 Data Compression Conference*, pages 251–260, March 2013.

- [8] Bernhard Schölkopf, Alexander J Smola, Francis Bach, et al. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press, 2002.
- [9] Bernhard E Boser, Isabelle M Guyon, and Vladimir N Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152, 1992.
- [10] Young-Gyu Yoon, Jaewook Kim, Tae-Kwang Jang, and SeongHwan Cho. A time-based bandpass ADC using time-interleaved voltage-controlled oscillators. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 55(11):3571–3581, 2008.
- [11] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2010.
- [12] Maxim Raginsky and Svetlana Lazebnik. Locality-sensitive binary codes from shift-invariant kernels. In *Advances in neural information processing systems*, pages 1509–1517, 2009.
- [13] Petros T. Boufounos, Shantanu Rane, and Hassan Mansour. Representation and coding of signal geometry. *Information and Inference: A Journal of the IMA*, 6(4):349–388, 2017.
- [14] Jian Zhang, Avner May, Tri Dao, and Christopher Ré. Low-Precision Random Fourier Features for Memory-Constrained Kernel Approximation. *arXiv preprint arXiv:1811.00155*, 2018.
- [15] Nick Ryder, Zohar Karnin, and Edo Liberty. Asymmetric random projections. *arXiv preprint arXiv:1906.09489*, 2019.
- [16] Wei Dong, Moses Charikar, and Kai Li. Asymmetric distance estimation with sketches for similarity search in high-dimensional spaces. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 123–130, 2008.
- [17] Albert Gordo, Florent Perronnin, Yunchao Gong, and Svetlana Lazebnik. Asymmetric distances for binary embeddings. *IEEE transactions on pattern analysis and machine intelligence*, 36(1):33–47, 2013.
- [18] Xiaoyun Li and Ping Li. Random projections with asymmetric quantization. In *Advances in Neural Information Processing Systems*, pages 10857–10866, 2019.
- [19] Ping Li. Sign-full random projections. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4205–4212, 2019.
- [20] R. Gribonval, G. Blanchard, N. Keriven, and Y. Traonmilin. Compressive Statistical Learning with Random Feature Moments. *ArXiv e-prints*, June 2017.
- [21] Nicolas Keriven et al. Compressive K-means. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6369–6373. IEEE, 2017.
- [22] Vincent Schellekens and Laurent Jacques. Quantized compressive K-means. *IEEE Signal Processing Letters*, 25(8):1211–1215, 2018.
- [23] AN Kolmogorov and VM Tihomirov. ϵ -entropy and ϵ -capacity of sets in functional space. *Amer. Math. Soc. Transl.(2)*, 17:277–364, 1961.
- [24] Gilles Pisier. *The volume of convex bodies and Banach space geometry*, volume 94. Cambridge University Press, 1999.

- [25] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [26] Ulaş Ayaz, Sjoerd Dirksen, and Holger Rauhut. Uniform recovery of fusion frame structured sparse signals. *Applied and Computational Harmonic Analysis*, 41(2):341–361, 2016.
- [27] Emmanuel J Candes and Yaniv Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359, 2011.
- [28] Armin Eftekhari and Michael B Wakin. New analysis of manifold embeddings and signal recovery from compressive measurements. *Applied and Computational Harmonic Analysis*, 39(1):67–109, 2015.
- [29] Walter Rudin. *Fourier Analysis on Groups*. Interscience Publishers, 1962.
- [30] Dougal J Sutherland and Jeff Schneider. On the error of random Fourier features. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 862–871, 2015.
- [31] Bharath Sriperumbudur and Zoltán Szabó. Optimal rates for random Fourier features. In *Advances in Neural Information Processing Systems*, pages 1144–1152, 2015.
- [32] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Yonina C. Eldar and Gitta Kutyniok, editors, *Compressed Sensing*, pages 210–268. Cambridge University Press, 2012.
- [33] Richard G Baraniuk, Volkan Cevher, Marco F Duarte, and Chinmay Hegde. Model-based compressive sensing. *IEEE Transactions on information theory*, 56(4):1982–2001, 2010.
- [34] Laurent Jacques and Valerio Cambareri. Time for dithering: fast and quantized random embeddings via the restricted isometry property. *Information and Inference: A Journal of the IMA*, 6(4):441–476, 2017.
- [35] Gilles Puy, Mike E Davies, and Rémi Gribonval. Recipes for stable linear embeddings from Hilbert spaces to \mathbb{R}^m . *IEEE Transactions on Information Theory*, 63(4):2171–2187, 2017.
- [36] Otto Szász. Fourier series and mean moduli of continuity. *Transactions of the American Mathematical Society*, 42(3):366–395, 1937.
- [37] Ingemar Wik. Criteria for absolute convergence of Fourier series of functions of bounded variation. *Transactions of the American Mathematical Society*, 163:1–24, 1972.
- [38] Venkat Chandrasekaran, Benjamin Recht, Pablo A Parrilo, and Alan S Willsky. The convex geometry of linear inverse problems. *Foundations of Computational mathematics*, 12(6):805–849, 2012.
- [39] Laurent Jacques, Jason N. Laska, Petros T. Boufounos, and Richard G. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transactions on Information Theory*, 59(4):2082–2102, apr 2013.
- [40] Daniel Otero and Gonzalo R Arce. Generalized restricted isometry property for alpha-stable random projections. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3676–3679. IEEE, 2011.
- [41] Tianbao Yang, Yu-Feng Li, Mehrdad Mahdavi, Rong Jin, and Zhi-Hua Zhou. Nyström method vs random Fourier features: A theoretical and empirical comparison. In *Advances in neural information processing systems*, pages 476–484, 2012.

- [42] Alessandro Rudi and Lorenzo Rosasco. Generalization properties of learning with random features. In *Advances in Neural Information Processing Systems*, pages 3215–3225, 2017.
- [43] Federica Gerace, Bruno Loureiro, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. Generalisation error in learning with random features and the hidden manifold model. *arXiv preprint arXiv:2002.09339*, 2020.
- [44] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [45] JA Gualtieri, Samir R Chettri, RF Crompt, and LF Johnson. Support vector machine classifiers as applied to AVIRIS data. In *Proc. Eighth JPL Airborne Geoscience Workshop*, 1999.
- [46] Farid Melgani and Lorenzo Bruzzone. Classification of hyperspectral remote sensing images with support vector machines. *IEEE Transactions on geoscience and remote sensing*, 42(8):1778–1790, 2004.
- [47] Nikhila Haridas, V Sowmya, and KP Soman. Gurls vs libsvm: Performance comparison of kernel methods for hyperspectral image classification. *Indian Journal of Science and Technology*, 8(24):1, 2015.
- [48] Giorgos Mountrakis, Jungho Im, and Caesar Ogole. Support vector machines in remote sensing: A review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(3):247–259, 2011.
- [49] Marion F. Baumgardner, Larry L. Biehl, and David A. Landgrebe. 220 Band AVIRIS Hyperspectral Image Data Set: June 12, 1992 Indian Pine Test Site 3, Sep 2015.