
Fast learning in reproducing kernel Kreĭn spaces via signed measures

Fanghui Liu
ESAT-STADIUS
KU Leuven

Xiaolin Huang
Department of Automation
Institute of Medical Robotics
Shanghai Jiao Tong University

Yingyi Chen
ESAT-STADIUS
KU Leuven

Johan A.K. Suykens
ESAT-STADIUS
KU Leuven

Abstract

In this paper, we attempt to solve a long-lasting open question for non-positive definite (non-PD) kernels in machine learning community: can a given non-PD kernel be decomposed into the difference of two PD kernels (termed as positive decomposition)? We cast this question as a distribution view by introducing the *signed measure*, which transforms positive decomposition to measure decomposition: a series of non-PD kernels can be associated with the linear combination of specific finite Borel measures. In this manner, our distribution-based framework provides a sufficient and necessary condition to answer this open question. Specifically, this solution is also computationally implementable in practice to scale non-PD kernels in large sample cases, which allows us to devise the first random features algorithm to obtain an unbiased estimator. Experimental results on several benchmark datasets verify the effectiveness of our algorithm over the existing methods.

1 Introduction

Devising a pairwise similarity/dissimilarity function plays a significant role in metric learning and kernel learning [1, 2, 3]. However, such function is not always positive definite (PD) in practice. For example, we are often faced with *indefinite* (real, symmetric, but not positive definite) kernels including the hyperbolic tangent kernel [4, 5] and truncated ℓ_1 -distance kernel [6]. Interestingly, some common-used PD kernels, e.g., polynomial kernels, Gaussian kernels, would degenerate to indefinite ones in some cases. An intuitive

example is that a linear combination of PD kernels with negative coefficients [7]. Polynomial kernels using ℓ_2 -normalization (i.e., distributed on the unit sphere) are not always PD [8]. Gaussian kernels with some geodesic distances cannot be ensured positive definite [9, 10]. We refer to a survey [11] for details.

Learning with indefinite similarity/dissimilarity functions is typically modeled in Reproducing Kernel Kreĭn Spaces (RKKS) [12], where the (reproducing) indefinite kernel can be decomposed into the difference of two PD kernels, a.k.a, positive decomposition [13]. A series of work [14, 15, 16, 17, 18, 19] rely on the positive decomposition. It is important to note that, indefinite kernel matrices can be decomposed in the difference of two positive semi-definite matrices by eigenvalue decomposition, but *for a given indefinite kernel, does it admit a positive decomposition?* is a long-lasting open question in machine learning community. In fact, it appears non-trivial how to verify that an indefinite kernel can be associated with RKKS except for some intuitive examples, e.g., a linear combination of PD kernels with negative coefficients. In the past, we usually assume that a (reproducing) indefinite kernel is associated with RKKS in practice while the theoretical gap cannot be ignored. In particular, the used eigenvalue decomposition in indefinite kernel based algorithms [15, 16, 17] often incurs huge computational and space complexities and thus is infeasible to large-scale problems.

To answer the open question, we consider indefinite kernel in a distribution view. Our model is based on the *signed measure*, which generalizes Borel measure to be negative. Accordingly, the positive decomposition can be transformed to measure decomposition and thus we provide a sufficient and necessary condition to answer this question. Our distribution-based framework is simple but effective, which naturally allows us to devise unbiased random features based algorithm to scale indefinite kernel methods in large sample cases. Formally, we make the following contributions:

- In Section 3, by introducing the signed measure, we provide a sufficient and necessary condition

to answer the above open question for indefinite kernels via the measure decomposition technique. Moreover, this condition also guides us how to find a specific positive decomposition in practice, and thus we can devise a unbiased estimator to obtain randomized feature maps. To the best of our knowledge, this is the first work to generate unbiased estimation for non-PD kernel approximation by random features.

- In Section 4, we demonstrate that spherically dot-product kernels including polynomial kernels, arc-cosine kernels, and the popular NTK in two-layer ReLU network on the unit sphere¹, are radial and non-PD associated with RKKS. Then we demonstrate the feasibility of our random feature algorithm on several indefinite kernels admitting the positive decomposition.
- In Section 5, we evaluate various non-PD kernels on several typical benchmark datasets to validate the effectiveness of our algorithm.

2 Preliminaries and Related Works

In this section, we briefly sketch some basic ideas of RKKS [13] and Bochner’s theorem in random features [20], and then introduce related works on indefinite kernel approximation.

2.1 Reproducing Kernel Kreĭn Spaces

Here we briefly review on the Kreĭn spaces and the reproducing kernel Kreĭn space (RKKS). Detailed expositions can be found in book [13]. Most of the readers would be familiar with Hilbert spaces. Kreĭn spaces share some properties of Hilbert spaces but differ in some key aspects which we shall emphasize as follows.

Kreĭn spaces are indefinite inner product spaces endowed with a Hilbertian topology.

Definition 1. (Kreĭn space [13]) *An inner product space is a Kreĭn space \mathcal{H}_K if there exist two Hilbert spaces \mathcal{H}_+ and \mathcal{H}_- such that*

- i) $\forall f \in \mathcal{H}_K$, it can be decomposed into $f = f_+ \oplus f_-$, where $f_+ \in \mathcal{H}_+$ and $f_- \in \mathcal{H}_-$, respectively.*
- ii) $\forall f, g \in \mathcal{H}_K$, $\langle f, g \rangle_{\mathcal{H}_K} = \langle f_+, g_+ \rangle_{\mathcal{H}_+} - \langle f_-, g_- \rangle_{\mathcal{H}_-}$.*

The Kreĭn space \mathcal{H}_K can be decomposed into a direct sum $\mathcal{H}_K = \mathcal{H}_+ \oplus \mathcal{H}_-$. Besides, the inner product on \mathcal{H}_K is non-degenerate, i.e., for $f \in \mathcal{H}_K$, if $\langle f, g \rangle_{\mathcal{H}_K} = 0$ for any $g \in \mathcal{H}_K$, we have $f = 0$. From the definition, the decomposition $\mathcal{H}_K = \mathcal{H}_+ \oplus \mathcal{H}_-$ is not necessarily unique. For a fixed decomposition, the inner product $\langle f, g \rangle_{\mathcal{H}_K}$

¹We use the ℓ_2 normalization scheme to ensure the data on the unit sphere as suggested by [8], which is different from directly using spherically i.i.d data.

is given accordingly [15, 16]. The key difference from Hilbert spaces is that the inner products might be negative for Kreĭn spaces, i.e., there exists $f \in \mathcal{H}_K$ such that $\langle f, f \rangle_{\mathcal{H}_K} < 0$. If \mathcal{H}_+ and \mathcal{H}_- are two RKHSs, the Kreĭn space \mathcal{H}_K is an RKKS associated with a unique indefinite reproducing kernel k such that the reproducing property holds, i.e., $\forall f \in \mathcal{H}_K$, $f(x) = \langle f, k(x, \cdot) \rangle_{\mathcal{H}_K}$.

Proposition 1. (positive decomposition [13]) *Let $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a real-valued kernel function. Then there exists an associated RKKS identified with a reproducing kernel k if and only if k admits a positive decomposition $k = k_+ - k_-$, where k_+ and k_- are two positive definite kernels.*

From the definition, this decomposition is not necessarily unique. As mentioned before, not every indefinite kernel function admits a representation as a difference between two positive definite kernels.

2.2 Bochner’s theorem and random features

A positive definite function corresponds to a nonnegative and finite Borel measure, i.e., a probability distribution, via Fourier transform by the following theorem.

Theorem 1 (Bochner’s Theorem [21]). *Let $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a bounded continuous function satisfying the stationary property, i.e., $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}')$. Then, k is positive definite if and only if it is the (conjugate) Fourier transform of a nonnegative and finite Borel measure μ (rescale it to a probability measure by setting $k(\mathbf{0}) = 1$)*

$$k(\mathbf{x} - \mathbf{x}') = \int_{\mathbb{R}^d} e^{i\boldsymbol{\omega}^\top(\mathbf{x} - \mathbf{x}')} \mu(d\boldsymbol{\omega}) = \mathbb{E}_{\boldsymbol{\omega} \sim \mu} [e^{i\boldsymbol{\omega}^\top(\mathbf{x} - \mathbf{x}')}].$$

Typically, the kernel in practical uses is real-valued and thus the imaginary part can be discarded, i.e., $k(\mathbf{x} - \mathbf{x}') = \mathbb{E}_{\boldsymbol{\omega} \sim \mu} \cos[\boldsymbol{\omega}^\top(\mathbf{x} - \mathbf{x}')]$. Accordingly, we can use the Monte Carlo method to sample a series of random features $\{\boldsymbol{\omega}_i\}_{i=1}^s$ from the distribution μ to approximate the PD kernel function k , a.k.a. Random Fourier features (RFF) [20]. It brings promising performance and solid theoretical guarantees on scaling up kernel methods in classification [22], nonlinear component analysis [23, 24], and neural tangent kernel (NTK) [25]. Improvements on RFF mainly focus on variance reduction by advanced sampling methods, e.g., quasi-Monte Carlo sampling [26], Monte Carlo sampling with orthogonal constraints [27, 28, 29], leverage-score sampling [30, 31], and quadrature based methods [32, 33, 34], see a survey [35] for details.

2.3 Signed measure

Let $\mu : \mathcal{A} \rightarrow [0, +\infty]$ be a measure on a set Ω satisfying $\mu(\emptyset) = 0$ and σ -additivity (i.e., countably additive). We

call μ a finite measure if $\mu(\Omega) < +\infty$. Specifically, μ is a probability measure if $\mu(\Omega) = 1$, and the triple $(\Omega, \mathcal{A}, \mu)$ is referred as the corresponding probability space. Here we consider the signed measure, a generalized version of a measure allowing for negative values.

Definition 2. (*Signed measure [36]*) Let Ω be some set, \mathcal{A} be a σ -algebra of subsets on Ω . A signed measure is a function $\mu : \mathcal{A} \rightarrow [-\infty, +\infty)$ or $(-\infty, +\infty]$ satisfying σ -additivity.

Based on this definition, the following theorem shows that any signed measure can be represented by the difference of two nonnegative measures.

Theorem 2. (*Jordan decomposition [36]*) Let μ be a signed measure defined on the σ -algebra \mathcal{A} as given in Definition 2. There exists two (nonnegative) measures μ_+ and μ_- (one of them is a finite measure) such that $\mu = \mu_+ - \mu_-$.

The total mass of μ on \mathcal{A} is defined as $\|\mu\| = \|\mu_+\| + \|\mu_-\|$. Note that this decomposition is not unique.

2.4 Related works

Learning with indefinite kernels in RKKS can be solved by eigenvalue transformation [37, 38], stabilization [12, 15], and minimization [16]. However, these methods need eigenvalue decomposition and cannot be directly applied to large-scale problems.

To scale indefinite kernel matrices in large sample problems, Nyström approximation works in a data-dependent way, and is a good choice to seek a low-rank representation to approximate indefinite kernel matrices, e.g., [17, 39, 40]. Besides, Liu et al. [41] decompose (a subset of) kernel matrix into two PD kernel matrices, and then learn their respective randomized feature maps by infinite Gaussian mixtures. However, this approach in fact focuses on approximating kernel matrices rather than kernel functions. If we consider indefinite kernel approximation by random features in a data-independent way, Pennington et al. [8] find that the polynomial kernel using ℓ_2 -normalized data is not PD, and then use (positive) mixtures of Gaussian distributions, associated with a PD kernel, to approximate it. This is in essence using a PD kernel to approximate an indefinite one. Till now, approximating non-PD kernels by random features cannot ensure unbiased and has not yet been fully investigated. In this paper, our work provides an unbiased estimator without extra parameters, so as to achieve both simplicity and effectiveness.

Besides, our algorithm can be also applied to dot-product kernels with ℓ_2 -normalized data, e.g., polynomial kernels on the unit sphere. Recent works for polynomial kernel approximation include Maclaurin expansion [42], the tensor sketch technique [43, 44], and oblivious subspace embedding [45, 46].

3 Model

In this section, by introducing the concept of signed measures [36], we attempt to answer the *open question* and then devise the sampling strategy for random features. For notational simplicity, we denote $z := \|\mathbf{z}\|_2 = \|\mathbf{x} - \mathbf{x}'\|_2$ and $\omega := \|\boldsymbol{\omega}\|_2$. Moreover, a function $k(\mathbf{z})$ is called *radial* if $k(\mathbf{z}) = k(\|\mathbf{z}\|_2)$. To notify, the considered stationary kernels in this paper are all *radial*, and their Fourier transforms are also *radial*, i.e., $\mu(\omega) = \mu(\boldsymbol{\omega})$, refer to [8, 47].

3.1 Answer to the open question in RKKS

As mentioned before, not every indefinite kernel admits a representation as a difference between two positive definite kernels. In fact we do not know how to verify that an indefinite kernel can be associated with RKKS except for some intuitive examples, e.g., a linear combination of PD kernels with negative coefficients. By virtue of measure decomposition of the signed measure in Theorem 2, we provide a sufficient and necessary condition in the following theorem to answer the question in RKKS: *for a given indefinite kernel, does it admit a positive decomposition?*

Theorem 3. Assume that an indefinite kernel is stationary, i.e., $k(\mathbf{x}, \mathbf{x}') = k(\mathbf{x} - \mathbf{x}')$. Denote its (generalized) Fourier transform as the measure μ , then we have the following results:

- (i) **Existence:** k admits the positive decomposition, i.e., $k = k_+ - k_-$, if and only if the total mass of the measure μ is finite, i.e., $\|\mu\| < \infty$. Here k_+ and k_- are two reproducing kernels associated with two reproducing kernel Hilbert spaces (RKHS) \mathcal{H}_+ and \mathcal{H}_- , respectively.
- (ii) **Representation:** If $\|\mu\| < \infty$, we choose μ_+ and μ_- such that $\mu = \mu_+ - \mu_-$, then the associated RKHSs \mathcal{H}_\pm are given by

$$\mathcal{H}_\pm = \left\{ f : \|f\|_{\mathcal{H}_\pm}^2 = \int_{\mathbb{R}^d} \frac{|F(\boldsymbol{\omega})|^2}{\mu_\pm(\boldsymbol{\omega})} d\boldsymbol{\omega} < \infty \right\},$$

where $F(\boldsymbol{\omega})$ is the Fourier transform of f .

Proof. The proof can be found in Appendix A. \square

Remark: We provide an explicit sufficient and necessary condition to link the Jordan decomposition of signed measures to positive decomposition in RKKS. (i) Functions that can be written as a difference between two positive functions have been studied and characterized intrinsically in the field of harmonic analysis [48, 49, 50]. They partly answered this question either restricted in the one-dimensional case (i.e., $d = 1$) or assuming the indefinite kernel $k(\mathbf{x}, \mathbf{x}')$ to be jointly analytic of \mathbf{x} and \mathbf{x}' in a neighborhood of the origin. However, the used univariate condition could not be directly applied to the machine

learning society and the smoothness requirement on kernels excludes some non-differentiable kernels, e.g., arc-cosine kernels. Instead, Theorem 3 provides an access via Fourier transform to verify whether a (reproducing) indefinite kernel belongs to RKKS or not. The measure decomposition is much easier to be founded than positive decomposition in RKKS that cannot be verified in practice.

(ii) Theorem 3 can be further improved to cover some non-squared-integrable kernel functions, e.g., conditionally positive definite kernels [51], of which the standard Fourier transform does not exist. In this case, Theorem 3 needs to consider the Fourier transform in Schwartz space [52]. For example, conditionally positive kernels correspond to a positive Borel measure μ on $\mathbb{R}^d \setminus \{\mathbf{0}\}$ with an analytic function in Schwartz space, refer to [53, Theorem 2.3].

3.2 Randomized feature map

The condition in Theorem 3 serves as a guidance for us to find a specific positive decomposition in practice. Hence we are ready to develop our random feature algorithm for (real-valued) non-PD kernels. One intuitive implementation way is choosing $\mu_+ := \max\{\mu, 0\}$ and $\mu_- := \min\{0, \mu\}$ such that $\mu = \mu_+ - \mu_-$. Then the stationary indefinite kernel k can be expressed by Eq. (1) via $k = k_+ - k_-$ with two positive constants c_1, c_2 . The decomposed two Borel measures $\tilde{\mu}_+ := \mu_+ / \|\mu_+\|$, $\tilde{\mu}_- := \mu_- / \|\mu_-\|$ are associated with two (normalized) PD kernels k_+ and k_- , respectively. Accordingly, the stationary indefinite kernel k can be approximated by $k \approx \tilde{k} := \tilde{k}_+ - \tilde{k}_-$. That is, $k(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\omega, \nu} \langle \Phi(\mathbf{x}), \Phi(\mathbf{x}') \rangle \approx \tilde{k}(\mathbf{x}, \mathbf{x}') := \frac{1}{s} \sum_{i=1}^s \langle \varphi_i(\mathbf{x}), \varphi_i(\mathbf{x}') \rangle$ where $\Phi(\mathbf{x})$ is the explicit feature mapping $\Phi(\mathbf{x}) = [\varphi_1(\mathbf{x}), \dots, \varphi_s(\mathbf{x})]^\top$ with $\varphi_i(\mathbf{x})$

$$\begin{aligned} \varphi_i(\mathbf{x}) = & \left[\sqrt{c_1 \|\mu_+\|} \cos(\omega_i^\top \mathbf{x}), \sqrt{c_1 \|\mu_+\|} \sin(\omega_i^\top \mathbf{x}), \right. \\ & \left. i \sqrt{c_2 \|\mu_-\|} \cos(\nu_i^\top \mathbf{x}), i \sqrt{c_2 \|\mu_-\|} \sin(\nu_i^\top \mathbf{x}) \right], \end{aligned} \quad (2)$$

where i is the imaginary unit. Then random features are obtained by sampling $\{\omega_i\}_{i=1}^s \sim \mu_+ / \|\mu_+\|$ and $\{\nu_i\}_{i=1}^s \sim \mu_- / \|\mu_-\|$. The employed sampling method can be Monte Carlo sampling, orthogonal Monte Carlo sampling [27, 54], or leverage-score based sampling [30, 55]. The real and imaginary part in $\varphi_i(\mathbf{x})$ correspond to k_+ and k_- , and thus our estimation is unbiased. It is important to note that, we need the imaginary unit in the feature mapping due to the difference operation², and then the approximated kernel is still real-valued.

The complete random features process is summarized in Algorithm 1. For a given stationary indefinite kernel,

²A simple example is that $a - b = \langle (\sqrt{a}, i\sqrt{b}), (\sqrt{a}, i\sqrt{b}) \rangle$ for two nonnegative real numbers a, b .

Algorithm 1: Random features for various indefinite kernels via generalized measures.

Input: A kernel function $k(\mathbf{x}, \mathbf{x}') = k(z)$ with $z := \|\mathbf{x} - \mathbf{x}'\|_2$ and the number of random features s .

Output: Random feature map

$\Phi(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^{4s}$ such that

$$k(\mathbf{x}, \mathbf{x}') \approx \frac{1}{s} \sum_{i=1}^s \langle \varphi_i(\mathbf{x}), \varphi_i(\mathbf{x}') \rangle.$$

1. Obtain the measure $\mu(\cdot)$ of the kernel k via (generalized) Fourier transform ;
 2. Given μ , let $\mu := \mu_+ - \mu_-$ be the Jordan decomposition with two nonnegative measures μ_\pm and compute the total mass $\|\mu\| = \|\mu_+\| + \|\mu_-\|$;
 3. Sample $\{\omega_i\}_{i=1}^s \sim \mu_+ / \|\mu_+\|$ and $\{\nu_i\}_{i=1}^s \sim \mu_- / \|\mu_-\|$;
 4. Output the explicit feature mapping $\Phi(\mathbf{x})$ with $\varphi_i(\mathbf{x})$ given in Eq. (2).
-

its μ, μ_\pm can be pre-computed, which is independent of the training data. In this way, our algorithm achieves the same complexity with the standard RFF by $\mathcal{O}(ns^2)$ time and $\mathcal{O}(ns)$ memory. Besides, the formulation in Eq. (1), as well as Algorithm 1, is general enough to cover various PD and non-PD kernels. Stationary PD kernels admit Eq. (1) by choosing $c_1 = 1$ and $c_2 = 0$ where we have $\mu = \mu_+$ associated with $\|\mu\| = 1$, i.e., a Borel measure. Hence, the Bochner's theorem can be regarded as a special case of the considered integration representation (1) in this paper.

The approximation performance in our method for indefinite kernels still achieves theoretical guarantees with those of PD kernels by the following proposition. The result can be easily derived from [56, 57], of which the proof refers to Appendix B for completeness.

Proposition 2. *Let k be a stationary indefinite kernel in RKKS with two Borel measures $\tilde{\mu}_\pm$ defined in Eq. (1), we have the following results:*

(i) **Approximation:** *Let \mathcal{S}_R be the compact ball by $\mathcal{S}_R = \{\Delta | \Delta \in \mathbb{R}^d, \|\Delta\|_2 \leq R\}$, then given the approximated kernel \tilde{k} obtained by our algorithm via Monte Carlo sampling, for any $\epsilon > 0$*

$$\Pr \left[\sup_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}_R} |k(\mathbf{x}, \mathbf{x}') - \tilde{k}(\mathbf{x}, \mathbf{x}')| \geq \epsilon \right] \leq 66 \left(\frac{2\sigma R}{\epsilon} \right)^2 e^{-\frac{s\epsilon^2}{32(d+2)}},$$

where $\sigma^2 := \mathbb{E}_{\omega \sim \tilde{\mu}_+} [\omega^\top \omega] + \mathbb{E}_{\omega \sim \tilde{\mu}_-} [\omega^\top \omega] < \infty$.

(ii) **Variance reduction:** *If we consider orthogonal Monte Carlo (OMC) sampling [27, 54] in our algorithm, it admits $\text{MSE}[\tilde{k}^{\text{OMC}}(\mathbf{x}, \mathbf{x}')] \leq \text{MSE}[\tilde{k}^{\text{MC}}(\mathbf{x}, \mathbf{x}')] for sufficiently large d , where the mean-squared error (MSE) is defined as $\mathbb{E}[\tilde{k}(\mathbf{x}, \mathbf{x}')] = \mathbb{E}[k(\mathbf{x}, \mathbf{x}') - \tilde{k}(\mathbf{x}, \mathbf{x}')]$.$*

$$\begin{aligned}
 k(\mathbf{x} - \mathbf{x}') &= c_1 \int_{\mathbb{R}^d} e^{i\boldsymbol{\omega}^\top \mathbf{z}} \mu_+(\mathrm{d}\boldsymbol{\omega}) - c_2 \int_{\mathbb{R}^d} e^{i\boldsymbol{\nu}^\top \mathbf{z}} \mu_-(\mathrm{d}\boldsymbol{\nu}) = c_1 \|\mu_+\| \mathbb{E}_{\boldsymbol{\omega} \sim \tilde{\mu}_+} [\cos(\boldsymbol{\omega}^\top \mathbf{z})] - c_2 \|\mu_-\| \mathbb{E}_{\boldsymbol{\nu} \sim \tilde{\mu}_-} [\cos(\boldsymbol{\nu}^\top \mathbf{z})] \\
 &:= k_+(\mathbf{x} - \mathbf{x}') - k_-(\mathbf{x} - \mathbf{x}') \approx \tilde{k}_+(\mathbf{x} - \mathbf{x}') - \tilde{k}_-(\mathbf{x} - \mathbf{x}') \\
 &= \frac{1}{s} \sum_{i=1}^s \langle \operatorname{Re}[\varphi_i(\mathbf{x})], \operatorname{Re}[\varphi_i(\mathbf{x}')] \rangle - \frac{1}{s} \sum_{i=1}^s \langle \operatorname{Im}[\varphi_i(\mathbf{x})], \operatorname{Im}[\varphi_i(\mathbf{x}')] \rangle.
 \end{aligned} \tag{1}$$

4 Examples

In this section, we investigate a series of indefinite kernels for a better understanding of our random features algorithm. We begin with an intuitive example, the linear combination of PD kernels with negative constraints. Then we discuss several dot-product kernels using ℓ_2 normalization data, including the polynomial kernel [8], the arc-cosine kernel [58], and the NTK kernel in a two-layer ReLU network [59].

A linear combination of positive definite kernels with negative coefficients: Kernels in this class admit the formulation $k = \sum_{i=1}^t a_i k_i$, where $\{k_i\}_{i=1}^t$ is the set of PD kernels, and $a_i \in \mathbb{R}$. This is a typical example of indefinite kernels in RKKS, which admits positive decomposition such that $k = k_+ - k_-$ with two PD kernels k_\pm . Theorem 3 guides us to find μ_\pm based on the sign of a_i . Hence we explicitly decompose an indefinite kernel in this class into the difference of two PD kernels, i.e., $k = k_+ - k_- := \sum_{i=1}^t \max(0, a_i) k_i - \sum_{i=1}^t \max(0, -a_i) k_i$. Then the corresponding nonnegative measures μ_\pm can be subsequently obtained due to the additivity of Fourier transform. We take the Delta-Gaussian kernel [16] $k(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|_2^2 / 2\tau_1^2) - \exp(-\|\mathbf{x} - \mathbf{x}'\|_2^2 / 2\tau_2^2)$ as an example. This kernel admits $c_1 = c_2 = 1$ and $\|\mu_+\| = \|\mu_-\| = 1$ in Eq. (1), and its random feature mapping is given by Eq. (2) with $\{\boldsymbol{\omega}_i\}_{i=1}^s \sim \mathcal{N}(0, \tau_1^{-2} \mathbf{I}_d)$ and $\{\boldsymbol{\nu}_i\}_{i=1}^s \sim \mathcal{N}(0, \tau_2^{-2} \mathbf{I}_d)$.

After providing the above simple warming-up example, we now discuss dot-product kernels on the unit sphere, and demonstrate the feasibility of our algorithm.

Polynomial kernels on the sphere: Pennington et al. [8] point out that a polynomial kernel on the unit sphere by ℓ_2 normalization is of $k(\mathbf{x}, \mathbf{x}') = \left(1 - \frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{a^2}\right)^p$ for $a \geq 2$ and $p \geq 1$ and $z := \|\mathbf{x} - \mathbf{x}'\|_2 \in [0, 2]$. This kernel is indefinite since its Fourier transform is not a nonnegative measure in [8]

$$\mu(\omega) = \sum_{i=0}^p \frac{p!}{(p-i)!} \left(1 - \frac{4}{a^2}\right)^{p-i} \left(\frac{2}{a^2}\right)^i \left(\frac{2}{\omega}\right)^{\frac{d}{2}+i} J_{\frac{d}{2}+i}(2\omega), \tag{3}$$

which results from the oscillatory behavior of the Bessel function of the first kind $J_{d/2+i}(2\omega)$. We demonstrate $\|\mu\| < \infty$ (see in Appendix C), which makes the integration our random features algorithm feasible by

decomposing $\mu = \mu_+ - \mu_-$ with $\mu_+ = \max\{0, \mu\}$ and $\mu_- = \min\{0, \mu\}$. Then random feature map for this kernel can be also given by Eq. (2) with $\{\boldsymbol{\omega}_i\}_{i=1}^s \sim \mu_+ / \|\mu_+\|$ and $\{\boldsymbol{\nu}_i\}_{i=1}^s \sim \mu_- / \|\mu_-\|$. Therefore, Algorithm 1 is suitable for this kernel. Note that the (scaled) measure μ_\pm is not a typical probability distribution, but the radial property of the Fourier transform allows us to conduct rejection sampling in one dimension to sample from this ‘‘complex’’ distribution, which does not incur too much computational cost. We experimentally evaluate this with other sampling schemes in Section 5. Compared to [8] using a positive sum of Gaussians to approximate $\mu(\omega)$, where parameters in Gaussians need to be optimized beforehand, our algorithm achieves both simplicity and effectiveness by having (i) an unbiased estimator, (ii) incurring no extra parameters. Figure 1 shows the superiority of our method to SRF on approximating the spherical polynomial kernel $k(z)$. It can be found that, our method is unbiased to achieve lower mean squared error since SRF directly overlooks the negative part of the signed measure μ .

Next we consider the NTK of two-layer ReLU networks on the unit sphere³ [59]. Since this kernel in fact consists of zero/first-order arc-cosine kernels [58], we combine them together for discussion.

NTK of Two-layer ReLU networks on the unit sphere: Bietti and Mairal [59] consider a two-layer ReLU network of the form $f(\mathbf{x}; \boldsymbol{\theta}) = \sqrt{2s} \sum_{j=1}^s \sum_{i=1}^s a_j \max\{\boldsymbol{\omega}_j^\top \mathbf{x}, 0\}$, with the parameter $\boldsymbol{\theta} = (\boldsymbol{\omega}_1^\top, \dots, \boldsymbol{\omega}_s^\top, a_1, \dots, a_s)$ initialized according to $\mathcal{N}(0, 1)$. By formulating ReLU as $\max\{\boldsymbol{\omega}_j^\top \mathbf{x}, 0\} = (\boldsymbol{\omega}_j^\top \mathbf{x})_+$, we have the following formulation corresponding to NTK [59, 60]

$$\begin{aligned}
 k(\mathbf{x}, \mathbf{x}') &= 2\mathbb{E}_{\boldsymbol{\omega} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [(\boldsymbol{\omega}^\top \mathbf{x})_+ (\boldsymbol{\omega}^\top \mathbf{x}')_+] \\
 &\quad + 2(\mathbf{x}^\top \mathbf{x}') \mathbb{E}_{\boldsymbol{\omega} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [1\{\boldsymbol{\omega}^\top \mathbf{x} \geq 0\} 1\{\boldsymbol{\omega}^\top \mathbf{x}' \geq 0\}],
 \end{aligned} \tag{4}$$

which can be further represented by $k(\mathbf{x}, \mathbf{x}') = \|\mathbf{x}\| \|\mathbf{x}'\| \cdot \kappa(\langle \mathbf{x}, \mathbf{x}' \rangle / (\|\mathbf{x}\| \|\mathbf{x}'\|))$ with $\kappa(u) := u\kappa_0(u) + \kappa_1(u)$. Here, $\kappa_0(u) = 1 - \frac{1}{\pi} \arccos(u)$ corresponds to the zero-order arc-cosine kernel and $\kappa_1(u) = \frac{1}{\pi}(u(\pi - \arccos(u)) + \sqrt{1 - u^2})$ is the first-order arc-cosine ker-

³This setting is actually different from the considered ℓ_2 -normalization case in this paper that cannot ensure the data are i.i.d on the unit sphere.

nel [58]. Furthermore, such NTK kernel is proved to be stationary but indefinite by the following theorem.

Theorem 4. *For any $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$, by ℓ_2 normalization, the NTK kernel of a two layer ReLU network of the form $f(\mathbf{x}; \boldsymbol{\theta}) = \sqrt{2s} \sum_{j=1}^s \sum_{j=1}^s a_j \max\{\boldsymbol{\omega}_j^\top \mathbf{x}, 0\}$ is stationary, that is,*

$$k(\mathbf{x}, \mathbf{x}') = \frac{2 - z^2}{\pi} \arccos\left(\frac{1}{2}z^2 - 1\right) + \frac{z}{2\pi} \sqrt{4 - z^2},$$

where $z := \|\mathbf{x} - \mathbf{x}'\|_2 \in [0, 2]$. However, the function $k(z), z \in [0, 2]$ is not positive definite.⁴

Proof. The proof can be found in Appendix D. \square

Since the above NTK on the unit sphere can be formulated as $k(\mathbf{x}, \mathbf{x}') = \langle \mathbf{x}, \mathbf{x}' \rangle \kappa_0(\mathbf{x}, \mathbf{x}') + \kappa_1(\mathbf{x}, \mathbf{x}')$ associated with arc-cosine kernels, we have the direct corollary for arc-cosine kernels.

Corollary 4.1. *The zero/first order arc-cosine kernel is not positive definite if the data are ℓ_2 -normalized, and its measure μ is given in Appendix E.*

Remark: These spherical dot-product kernels including polynomial kernels, arc-cosine kernels, and NTK are indefinite by ℓ_2 normalization, which extends the classical insight on spherical dot-product kernels via spherical harmonics [61]. Besides, our findings motivate us to scrutinize functional spaces, the approximation performance, and generalization properties of over-parameterized networks in deep learning theory if considering ℓ_2 -normalization data, which in return expands the usage scope of indefinite kernels.

In Appendix E, we compute the measure μ of arc-cosine kernels, which is quite complex as it involves with the sum of infinite series with Bessel functions. When taking finite series (e.g., one term) as an approximation, we demonstrate $\|\mu\| < \infty$. In this case, Algorithm 1 is accordingly suitable for arc-cosine kernels and NTK on the unit sphere. If we take more terms in finite series, the calculation appears non-trivial. And specifically, there exists a gap between the original μ and its approximation by finite series, so we do not include these two kernels in our experiments.

5 Experiments

We evaluate the proposed method on four representative benchmark datasets including *letter*⁵, *ijcnn1*⁶, *covtype*⁶, and *cod-RNA*⁶, see in Table 1. The datasets are normalized to $[0, 1]^d$ by an ℓ_2 -norm scaling scheme

⁴The behavior of $k(z)$ with $z > 2$ is undefined. Following [8], we set $k(z) = 0$ for $z > 2$.

⁵<https://archive.ics.uci.edu/ml/datasets.html>.

⁶<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

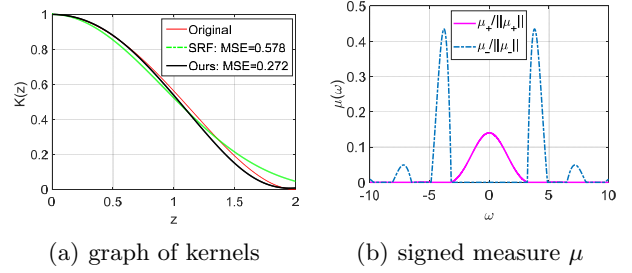


Figure 1: Approximation of the spherical polynomial kernel with $a = p = 2$.

Table 1: Benchmark datasets.

Datasets	d	#training	#test
<i>letter</i>	16	12,000	6,000
<i>ijcnn1</i>	22	49,990	91,701
<i>covtype</i>	54	290,506	290,506
<i>cod-RNA</i>	8	59,535	157,413

and have been given with training/test partition except for *covtype*. Hence, we randomly split the *covtype* dataset into the training and test sets by half. In our experiment, the used indefinite kernels are the spherical polynomial kernel $k(\mathbf{x}, \mathbf{x}') = (1 - \|\mathbf{x} - \mathbf{x}'\|_2^2/a^2)^p$ with $a = 2$, $p = 2$ in [8], and the Delta-Gaussian kernel $k(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2/2\tau_1^2) - \exp(-\|\mathbf{x} - \mathbf{x}'\|^2/2\tau_2^2)$ with $\tau_1 = 1$ and $\tau_2 = 10$ in [16]. The compared algorithms include SRF (Spherical Random Features) [8], DIGMM (Double-Infinite Gaussian Mixtures Model) [41], and Nyström with leverage score [17]. Moreover, we also include Random Maclaurin (RM) [42], Tensor Sketch (TS) [43], and Tensorized Random Projection (TRP) [44] for polynomial kernel approximation. Note that the related error bars and standard deviations are obtained by running the experiments for 10 times. All experiments are implemented in MATLAB and carried out on a PC with Intel® i7-8700K CPU (3.70 GHz) and 64 GB RAM. The source code of our implementation can be found in <http://www.lfhsgre.org>.

Kernel approximation: The relative error $\|\mathbf{K} - \widetilde{\mathbf{K}}\|_F/\|\mathbf{K}\|_F$ is chosen to measure the approximation quality where \mathbf{K} and $\widetilde{\mathbf{K}}$ denote the exact kernel matrix on 1,000 random selected samples and its approximated kernel matrix, respectively. Figure 2 shows the approximation error under two indefinite kernels as a function of #random features s . Our method achieves lower approximation error than the other algorithms across such two kernels on these datasets in most cases. A clear look at the case of Delta-Gaussian kernel approximation will find that our approach significantly improves the approximation quality compared to random features

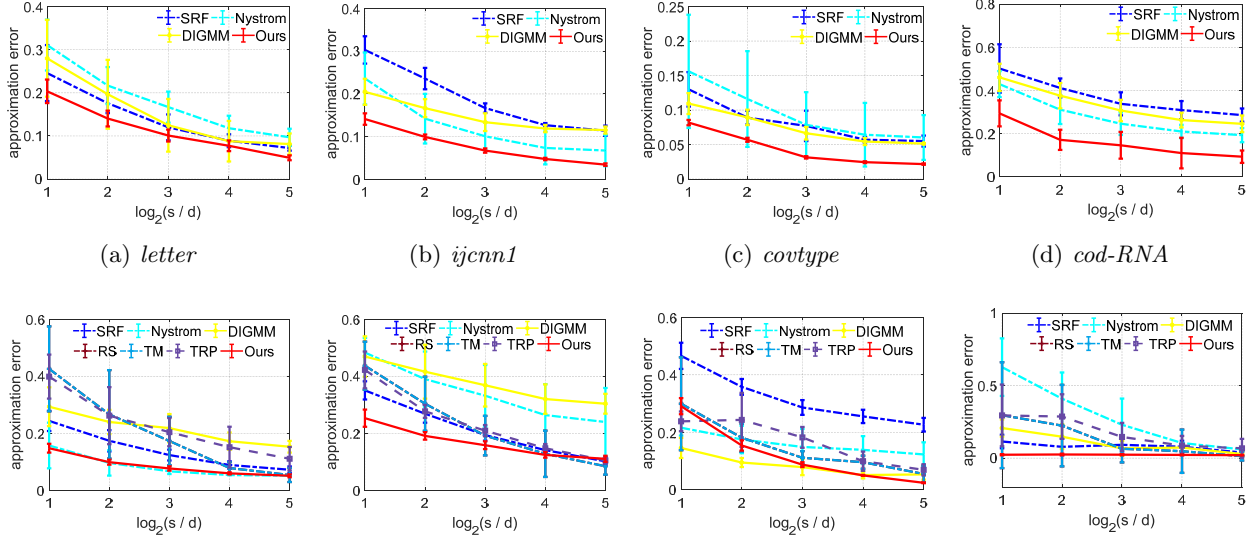


Figure 2: Comparisons of various algorithms for approximation error across the Delta-Gaussian kernel (top) and the spherical polynomial kernel (bottom) on four datasets.

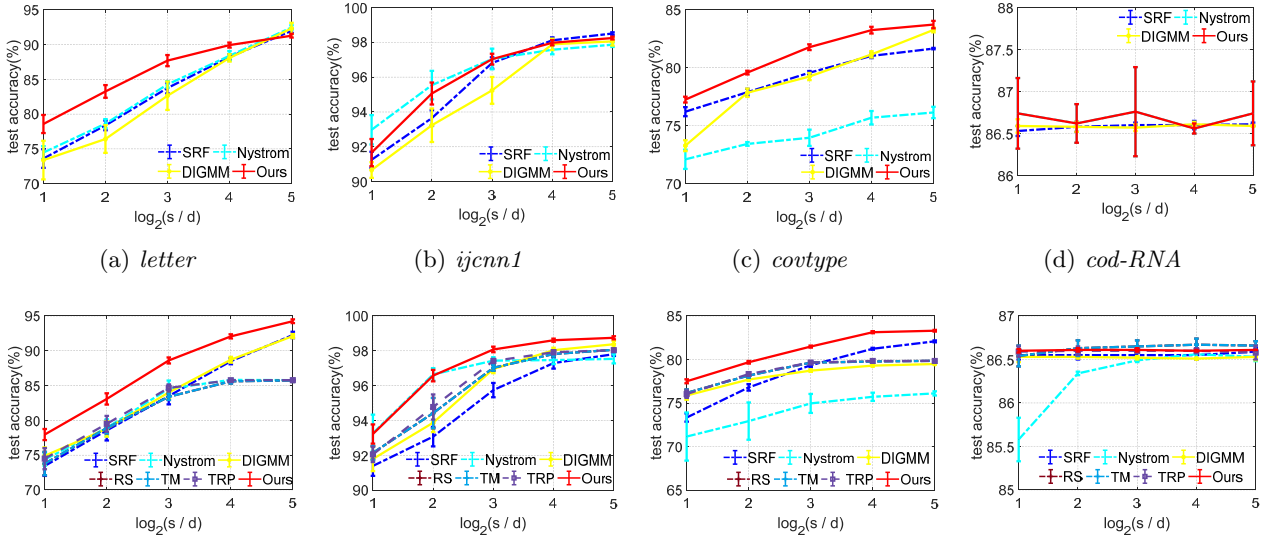


Figure 3: Comparisons of various algorithms for classification accuracy with LibLinear across the Delta-Gaussian kernel (top) and the spherical polynomial kernel (down) on four datasets.

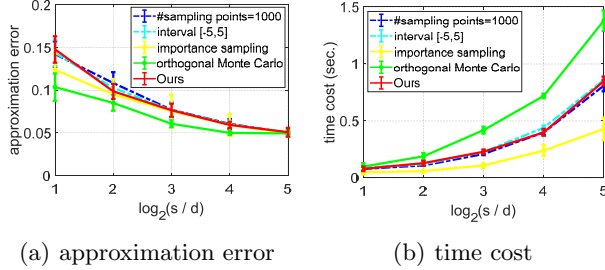
based algorithms: SRF and DIGMM. There always exists a gap in SRF that uses a PD kernel to approximate an indefinite one since the negative part is overlooked. DIGMM only focuses on approximating a subset of the kernel matrix. Different from these two, our method directly approximates the indefinite kernel function by an unbiased estimator, which incurs no extra loss for kernel approximation. Besides, when compared with several representative algorithms for polynomial kernels, e.g., RM, TS, and TRP, our method still performs well, which extends the application of our model.

Classification with linear SVM: We train a linear classifier: LibLinear [62] with the obtained randomized feature map. The balanced parameter in linear SVM is tuned by five-fold cross validation on a grid of points: $C = [0.01, 0.1, 1, 10, 100]$. The test accuracy of various algorithms are shown in Figure 3. As we expected, higher-dimensional randomized feature map outputs higher classification accuracy except the *cod-RNA* dataset. On this dataset, all of algorithms achieve the similar classification accuracy under various s . Apart from this dataset, our method performs best in most cases.

Table 2: Time cost (sec.) for generating feature matrices of various algorithms.

Datasets	s	Spherical polynomial kernel							Delta-Gaussian kernel			
		SRF ¹	Nystrom	DIGMM	RM	TS	TRP	Ours	SRF ¹	Nystrom	DIGMM	Ours
<i>letter</i>	$2d$	17.7 +0.02	0.12	0.21	0.04	0.07	0.01	0.08	17.2 +0.03	0.14	0.32	0.07
	$8d$	0.09	0.24	0.41	0.12	0.10	0.21	0.23	0.10	0.39	1.19	0.23
	$32d$	0.33	1.27	2.69	0.38	0.31	0.89	0.85	0.30	1.69	4.56	0.92
<i>ijcnn1</i>	$2d$	10.3 +0.24	0.83	0.42	0.33	0.53	0.73	0.70	20.3 +0.23	1.23	0.61	0.41
	$8d$	0.89	2.78	1.08	1.20	0.90	2.74	1.87	0.86	4.38	1.96	1.44
	$32d$	3.30	16.44	8.36	4.50	2.64	10.50	7.31	3.42	22.64	6.86	5.67

¹ On each dataset, SRF obtains parameters in GMM by an off-line grid search scheme in advance, of which this extra time cost is reported in **bold**.



(a) approximation error


(b) time cost

Figure 4: Comparisons of various sampling schemes across the spherical polynomial kernel on *letter*.

Computational time: Table 2 reports the time cost on generating randomized feature map with various dimensions s on two datasets. Our method achieves the same complexity with the standard RFF with $\mathcal{O}(ns^2)$ time and $\mathcal{O}(ns)$ memory. In practice, as reported by Table 2, our method takes a little more time than SRF to generate randomized feature maps due to the introduced extra imaginary part. Nevertheless, on each dataset, SRF requires extra time to obtain parameters of a sum of Gaussians in advance.

Sampling schemes in spherical polynomial kernel: The measures $\mu_+/\|\mu_+\|$ and $\mu_-/\|\mu_-\|$ associated with the spherical polynomial kernel are not typical distributions, so we conduct rejection sampling to acquire them by generating a set of uniformly 10,000 samples in a range of $[-10, 10]$. Here we compare various sampling schemes in our method for spherical polynomial kernel approximation, including sampling with 1,000 points, sampling in a sub-interval $[-5, 5]$, importance sampling, and orthogonal Monte Carlo. Here the *surrogate* distribution in importance sampling is chosen as the Gaussian distribution. The applied orthogonal Monte Carlo, followed by [27], aims to obtain orthogonal random features.

Figure 4 shows the approximation error and time cost of various sampling schemes. It can be found that, orthogonal random features achieve lower approximation error but require more computational cost, as suggested by [27, 54]. Instead, the applied importance sampling

(a) order $p = 1$ (b) order $p = 3$ Figure 5: Approximation error of various algorithms across the spherical polynomial kernel with different orders on the *letter* dataset.

decreases the time cost with a slight improvement on the approximation performance. If we choose the ridge leverage function in importance sampling, our model works with the leverage score based sampling, refer to [30] for details.

Different orders in spherical polynomial kernel: Apart from the used $p = 2$ in the spherical polynomial kernel in our experiment, we evaluate our model on spherical polynomial kernels with various orders, e.g., $p = 1$ and $p = 3$. Kernel approximation results in Figure 5 show that, under different orders, our method performs better than other algorithms in terms of the approximation error.

6 Conclusion

We answer the open question of indefinite kernels in machine learning community by the introduced measure decomposition technique, which motivates us to develop a general random features algorithm across various kernels that are stationary indefinite kernels. Albeit simple, our algorithm is effective to output unbiased estimates for indefinite kernel approximation. Besides, our findings on the indefiniteness of NTK on the unit sphere (by ℓ_2 normalization) encourages us to better understand the approximation performance, functional spaces, and generalization properties in over-parameterized networks in the future.

Acknowledgements

The research leading to these results has received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation program / ERC Advanced Grant E-DUALITY (787960). This paper reflects only the authors’ views and the Union is not liable for any use that may be made of the contained information. This work was supported in part by Research Council KU Leuven: Optimization frameworks for deep kernel machines C14/18/068; Flemish Government: FWO projects: GOA4917N (Deep Restricted Kernel Machines: Methods and Foundations), PhD/Postdoc grant. This research received funding from the Flemish Government (AI Research Program). This work was supported in part by Ford KU Leuven Research Alliance Project KUL0076 (Stability analysis and performance improvement of deep reinforcement learning algorithms), EU H2020 ICT-48 Network TAILOR (Foundations of Trustworthy AI - Integrating Reasoning, Learning and Optimization), Leuven.AI Institute; and in part by the National Natural Science Foundation of China 61977046. Corresponding author: Fanghui Liu (fanghui.liu@kuleuven.be) and Xiaolin Huang (xiaolin.huang@sjtu.edu.cn).

References

- [1] Brian Kulis. Metric learning: a survey. *Foundations and Trends in Machine Learning*, 5(4), 2013.
- [2] Lalit Jain, Blake Mason, and Robert Nowak. Learning low-dimensional metrics. In *Proceedings of Advances in Neural Information Processing Systems*, pages 4142–4150, 2017.
- [3] Han-Jia Ye, De-Chuan Zhan, and Yuan Jiang. Fast generalization rates for distance metric learning. *Machine Learning*, 108(2):267–295, 2019.
- [4] Alex J. Smola, Zoltan L. Ovari, and Robert C. Williamson. Regularization with dot-product kernels. In *Advances in Neural Information Processing Systems*, pages 308–314, 2001.
- [5] Hyunghoon Cho, Benjamin DeMeo, Jian Peng, and Bonnie Berger. Large-margin classification in hyperbolic space. In *International Conference on Artificial Intelligence and Statistics*, pages 1832–1840. PMLR, 2019.
- [6] Xiaolin Huang, Johan A.K. Suykens, Shuning Wang, Joachim Hornegger, and Andreas Maier. Classification with truncated ℓ_1 distance kernel. *IEEE Transactions on Neural Networks and Learning Systems*, 29(5):2025 – 2030, 2018.
- [7] Cheng Soon Ong, Alexander J. Smola, and Robert C Williamson. Learning the kernel with hyperkernels. *Journal of Machine Learning Research*, 6(Jul):1043–1071, 2005.
- [8] Jeffrey Pennington, Felix Xinnan X. Yu, and Sanjiv Kumar. Spherical random features for polynomial kernels. In *Advances in Neural Information Processing Systems*, pages 1846–1854, 2015.
- [9] Sadeep Jayasumana, Richard Hartley, Mathieu Salzmann, Hongdong Li, and Mehrtaash Harandi. Kernel methods on the Riemannian manifold of symmetric positive definite matrices. In *the IEEE Conference on Computer Vision and Pattern Recognition*, pages 73–80, 2013.
- [10] Aasa Feragen, François Lauze, and Søren Hauberg. Geodesic exponential kernels: when curvature and linearity conflict. In *the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3032–3042, 2015.
- [11] Frank Michael Schleif and Peter Tino. Indefinite proximity learning: a review. *Neural Computation*, 27(10):2039–2096, 2015.
- [12] Cheng Soon Ong, Xavier Mary, and Alexander J. Smola. Learning with non-positive kernels. In *the International Conference on Machine Learning*, pages 81–89, 2004.
- [13] János Bognár. *Indefinite inner product spaces*. Springer, 1974.
- [14] Volker Roth, Julian Laub, Motoaki Kawanabe, and Joachim Buhmann. Optimal cluster preserving embedding of nonmetric proximity data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(12):1540–1551, 2003.
- [15] Gaëlle Loosli, Stéphane Canu, and Soon Ong Cheng. Learning SVM in Krein spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(6):1204–1216, 2016.
- [16] Dino Oglic and Thomas Gärtner. Learning in reproducing kernel Krein spaces. In *the International Conference on Machine Learning*, pages 3859–3867, 2018.
- [17] Dino Oglic and Thomas Gärtner. Scalable learning in reproducing kernel krein spaces. In *International Conference on Machine Learning*, pages 4912–4921, 2019.
- [18] Akash Saha and Balamurugan Palaniappan. Learning with operator-valued kernels in reproducing kernel krein spaces. In *Advances in Neural Information Processing Systems*, pages 1–11, 2020.

- [19] Fanghui Liu, Lei Shi, Xiaolin Huang, Jie Yang, and Johan A.K. Suykens. Analysis of regularized least squares in reproducing kernel kreĭn spaces. *Machine Learning*, pages 1–20, 2021.
- [20] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, pages 1177–1184, 2007.
- [21] Salomon Bochner. *Harmonic Analysis and the Theory of Probability*. Courier Corporation, 2005.
- [22] Yitong Sun, Anna Gilbert, and Ambuj Tewari. But how does it work in theory? Linear SVM with random features. In *Advances in Neural Information Processing Systems*, pages 3383–3392, 2018.
- [23] Bo Xie, Yingyu Liang, and Le Song. Scale up non-linear component analysis with doubly stochastic gradients. In *Advances in Neural Information Processing Systems*, pages 2341–2349, 2015.
- [24] David Lopez-Paz, Suvrit Sra, Alex J. Smola, Zoubin Ghahramani, and Bernhard Schölkopf. Randomized nonlinear component analysis. In *the International Conference on Machine Learning*, pages 1359–1367, 2014.
- [25] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- [26] Jiyan Yang, Vikas Sindhwani, Haim Avron, and Michael Mahoney. Quasi-Monte Carlo feature maps for shift-invariant kernels. In *the International Conference on Machine Learning*, pages 485–493, 2014.
- [27] Felix Xinnan Yu, Ananda Theertha Suresh, Krzysztof Choromanski, Daniel Holtmannrue, and Sanjiv Kumar. Orthogonal random features. In *Advances in Neural Information Processing Systems*, pages 1975–1983, 2016.
- [28] Yueming Lyu. Spherical structured feature maps for kernel approximation. In *the 34th International Conference on Machine Learning*, pages 2256–2264. JMLR.org, 2017.
- [29] Krzysztof M. Choromanski, Mark Rowland, and Adrian Weller. The unreasonable effectiveness of structured random orthogonal embeddings. In *Advances in Neural Information Processing Systems*, pages 219–228, 2017.
- [30] Haim Avron, Michael Kapralov, Cameron Musco, Christopher Musco, Ameya Velingker, and Amir Zandieh. Random Fourier features for kernel ridge regression: Approximation bounds and statistical guarantees. In *the 34th International Conference on Machine Learning*, pages 253–262, 2017.
- [31] Zhu Li, Jean-Francois Ton, Dino Oglic, and Dino Sejdinovic. Towards a unified analysis of random Fourier features. In *the 36th International Conference on Machine Learning*, pages 3905–3914, 2019.
- [32] Tri Dao, Christopher M. De Sa, and Christopher Ré. Gaussian quadrature for kernel features. In *Advances in neural information processing systems*, pages 6107–6117, 2017.
- [33] Marina Munkhoeva, Yermek Kapushev, Evgeny Burnaev, and Ivan Oseledets. Quadrature-based features for kernel approximation. In *Advances in Neural Information Processing Systems*, pages 9147–9156, 2018.
- [34] Fanghui Liu, Xiaolin Huang, Yudong Chen, and Johan A.K. Suykens. Towards a unified quadrature framework for large-scale kernel machines. *arXiv preprint arXiv:2011.01668*, 2020.
- [35] Fanghui Liu, Xiaolin Huang, Yudong Chen, and Johan A.K. Suykens. Random features for kernel approximation: A survey in algorithms, theory, and beyond. *arXiv preprint arXiv:2004.11154*, 2020.
- [36] Krishna B Athreya and Soumendra N Lahiri. *Measure theory and probability theory*. Springer Science & Business Media, 2006.
- [37] Yihua Chen, Maya R Gupta, and Benjamin Recht. Learning kernels from indefinite similarities. In *the International Conference on Machine Learning*, pages 145–152. ACM, 2009.
- [38] Yiming Ying, Colin Campbell, and Mark Girolami. Analysis of SVM with indefinite kernels. In *Advances in Neural Information Processing Systems*, pages 2205–2213, 2009.
- [39] Siamak Mehrkanoon, Xiaolin Huang, and Johan A.K. Suykens. Indefinite kernel spectral learning. *Pattern Recognition*, 78:144–153, 2018.
- [40] Frank-Michael Schleif, Andrej Gisbrecht, and Peter Tino. Probabilistic classifiers with low rank indefinite kernels. *arXiv preprint arXiv:1604.02264*, 2016.

- [41] Fanghui Liu, Lei Shi, Xiaolin Huang, Jie Yang, and Johan A.K. Suykens. A double-variational Bayesian framework in random Fourier features for indefinite kernels. *IEEE Transactions on Neural Networks and Learning Systems*, 31(8):2965–2979, 2020.
- [42] Purushottam Kar and Harish Karnick. Random feature maps for dot product kernels. In *the International Conference on Artificial Intelligence and Statistics*, pages 583–591, 2012.
- [43] Ninh Pham and Rasmus Pagh. Fast and scalable polynomial kernels via explicit feature maps. In *ACM International Conference on Knowledge Discovery and Data Mining*, pages 239–247, 2013.
- [44] Michela Meister, Tamas Sarlos, and David Woodruff. Tight dimensionality reduction for sketching low degree polynomial kernels. In *Advances in Neural Information Processing Systems*, pages 9475–9486, 2019.
- [45] Haim Avron, Huy Nguyen, and David Woodruff. Subspace embeddings for the polynomial kernel. In *Advances in neural information processing systems*, pages 2258–2266, 2014.
- [46] David P Woodruff and Amir Zandieh. Near input sparsity time kernel embeddings via adaptive sampling. In *International Conference on Machine Learning*, pages 10324–10333, 2020.
- [47] Aurélie Boisbunon. The class of multivariate spherically symmetric distributions. *Université de Rouen, Technical Report, # 2012-005*, 2012.
- [48] P.H. Maserick. BV-functions, positive-definite functions and moment problems. *Transactions of the American Mathematical Society*, 214:137–152, 1975.
- [49] Daniel Alpay. Some reproducing kernel spaces of continuous functions. *Journal of Mathematical Analysis and Applications*, 160(2):424–433, 1991.
- [50] Daniel Alpay. Some remarks on reproducing kernel Krein spaces. *The Rocky Mountain Journal of Mathematics*, pages 1189–1205, 1991.
- [51] Holger Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- [52] William F Donoghue. *Distributions and Fourier transforms*. Academic Press, 2014.
- [53] Xingping Sun. Conditionally positive definite functions and their application to multivariate interpolations. *Journal of Approximation Theory*, 74(2):159–180, 1993.
- [54] Krzysztof Choromanski, Mark Rowland, Wenyu Chen, and Adrian Weller. Unifying orthogonal Monte Carlo methods. In *International Conference on Machine Learning*, pages 1203–1212, 2019.
- [55] Francis Bach. On the equivalence between kernel quadrature rules and random feature expansions. *Journal of Machine Learning Research*, 18(1):714–751, 2017.
- [56] Dougal J. Sutherland and Jeff Schneider. On the error of random Fourier features. In *the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 862–871, 2015.
- [57] Krzysztof Choromanski, Mark Rowland, Tamás Sarlós, Vikas Sindhwani, Richard Turner, and Adrian Weller. The geometry of random features. In *International Conference on Artificial Intelligence and Statistics*, pages 1–9, 2018.
- [58] Youngmin Cho and Lawrence K Saul. Kernel methods for deep learning. In *Advances in Neural Information Processing Systems*, pages 342–350, 2009.
- [59] Alberto Bietti and Julien Mairal. On the inductive bias of neural tangent kernels. In *Proceedings of Advances in Neural Information Processing Systems*, pages 12873–12884, 2019.
- [60] Lenaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pages 2933–2943, 2019.
- [61] Claus Müller. *Spherical harmonics*, volume 17. Springer, 2006.
- [62] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. LIBLINEAR: a library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.
- [63] Thomas Hotz and Fabian JE Telschow. Representation by integrating reproducing kernels. *arXiv preprint arXiv:1202.4443*, 2012.
- [64] Isaac J Schoenberg. Metric spaces and completely monotone functions. *Annals of Mathematics*, pages 811–841, 1938.