

Data-driven Random Fourier Features using Stein Effect

Wei-Cheng Chang

LTI, CMU

wchang2@cs.cmu.edu

Chun-Liang Li

MLD, CMU

chunlial@cs.cmu.edu

Yiming Yang

LTI, CMU

yiming@cs.cmu.edu

Barnabás Póczos

MLD, CMU

bapoczos@cs.cmu.edu

Abstract

Large-scale kernel approximation is an important problem in machine learning research. Approaches using random Fourier features have become increasingly popular, where kernel approximation is treated as empirical mean estimation via Monte Carlo (MC) or Quasi-Monte Carlo (QMC) integration. A limitation of the current approaches is that all the features receive an equal weight summing to 1. In this paper, we propose a novel shrinkage estimator from “Stein effect”, which provides a data-driven weighting strategy for random features and enjoys theoretical justifications in terms of lowering the empirical risk. We further present an efficient randomized algorithm for large-scale applications of the proposed method. Our empirical results on six benchmark data sets demonstrate the advantageous performance of this approach over representative baselines in both kernel approximation and supervised learning tasks.

1 Introduction

Kernel methods offer a comprehensive collection of non-parametric techniques for modeling a broad range of problems in machine learning. However, standard kernel machines such as kernel SVM or kernel ridge regression suffer from slow training and prediction which limits their use in real-world large-scale applications. Consider, for example, the training of linear regression over n labeled data points $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ where $\mathbf{x}_i \in \mathbb{R}^d$ with $n \gg d$ and y_i is a binary label. The time complexity of standard least square fit is $O(nd^2)$ and the memory demand is $O(nd)$. Its kernel-based counterpart requires solving a linear system with an n -by- n kernel matrix, which takes $O(n^3 + n^2d)$ time and $O(n^2)$ memory usage as typical. Such complexity is far from desirable in big-data applications.

In recent years, significant efforts have been devoted into low-rank kernel approximation for enhancing the scalability of kernel methods. Among existing approaches, random features [Rahimi and Recht, 2007; Rahimi and Recht, 2009; Le *et al.*, 2013; Bach, 2015] have drawn considerable attention and yielded state-of-the-art results in classification accuracy on benchmark data sets [Huang *et al.*, 2014; Dai *et al.*, 2014;

Li and Póczos, 2016]. Specifically, inspired from Bochner’s theorem [Rudin, 2011], random Fourier features have been studied for evaluating the expectation of shift-invariant kernels (i.e., $k(\mathbf{x}, \mathbf{x}') = g(\mathbf{x} - \mathbf{x}')$ for some function g). Rahimi and Recht [2007] proposed to use Monte-Carlo methods (MC) to estimate the expectation; Yang *et al.* [2014] leveraged the low-discrepancy properties of Quasi-Monte Carlo (QMC) sequences to reduce integration errors.

Both the MC and QMC methods rely on the (implicit) assumption that all the M random features are equally important, and hence assign a uniform weight $\frac{1}{M}$ to the features in kernel estimation. Such a treatment, however, is arguably sub-optimal for minimizing the expected risk in kernel approximation. Avron *et al.* [2016b] presented a weighting strategy for minimizing a loose error bound which depends on the maximum norm of data points. On the other hand, Bayesian Quadrature (BQ) is a powerful framework that solves the integral approximation by $\mathbb{E}[f(x)] \approx \sum_m \beta_m f(x_m)$ where $f(x_m)$ is feature function and β_m is with non-uniform weights. BQ leverages Gaussian Process (GP) to model the prior of function $f(\cdot)$, and the resulted estimator is Bayes optimal [Huszár and Duvenaud, 2012]. Thus, BQ could be considered a natural choice for kernel approximation. Nonetheless, a well-known limitation (or potential weakness) is that the performance of Bayesian-based models heavily relies on extensive hyper-parameter tuning, e.g., on the condition of the covariance matrix in the Gaussian prior, which is not suitable for large real-world applications.

To address the fundamental limitations of the aforementioned work, we propose a novel approach to data-driven feature weighting in the approximation of shift-invariant kernels, which motivated by the by Stein Effect in the statistical literature (Section 4), and solve it using an efficient stochastic algorithm with a convex optimization objective. We also present a natural extension of BQ to the applications of kernel approximation (Section 3). The adapted BQ together with standard MC and QMC methods serve as representative baselines in our empirical evaluations on six benchmark data sets. The empirical results (Section 5) show that the proposed Stein-Effect Shrinkage (SES) estimator consistently outperforms the baseline methods in terms of lowering the kernel approximation errors, and was competitive to or better than the best performer among the baseline methods in most task-oriented evaluations (on supervised learning of regression and classifi-

cation tasks).

2 Preliminaries

Let us briefly outline the related background and key concepts in random-feature based kernel approximation.

2.1 Reproducing Kernel Hilbert Space

At the heart of kernel methods lies the idea that inner products in high-dimensional feature spaces can be computed in an implicit form via a kernel function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, which is defined on an input data domain $\mathcal{X} \subset \mathbb{R}^d$ such that

$$k(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathcal{H}}.$$

where $\phi : \mathcal{X} \rightarrow \mathcal{H}$ is a feature map that associates kernel k with an embedding of the input space into a high-dimensional Hilbert space \mathcal{H} . A key to kernel methods is that as long as kernel algorithms have access to k , one need not to represent $\phi(\mathbf{x})$ explicitly. Most often, that means although $\phi(\mathbf{x})$ can be high-dimensional or even infinite-dimensional, their inner products, can be evaluated by k . This idea is known as the kernel trick.

The bottleneck of kernel methods lies in both training and prediction phase. The former requires to compute and store the kernel matrix $K \in \mathbb{R}^{n \times n}$ for n data points, while the latter need to evaluate the decision function $f(\mathbf{x})$ via kernel trick $f(\mathbf{x}) = \sum_{i=1}^N \alpha_i k(\mathbf{x}_i, \mathbf{x})$, which sums over the nonzero support vectors α_i . Unfortunately, Steinwart and Christmann [2008] show that the number of nonzero support vectors growth linearly in the size of training data n . This posts a great challenge for designing scalable algorithms for kernel methods in large-scale applications. Some traditional way to address this difficulty is by making trade-offs between time and space. [Kivinen *et al.*, 2004; Chang and Lin, 2011].

2.2 Random Features (RF)

Rahimi and Recht [2007] propose a low-dimensional feature map $z(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}^{2M}$ for the kernel function k

$$k(\mathbf{x}, \mathbf{x}') \approx \langle z(\mathbf{x}), z(\mathbf{x}') \rangle \quad (1)$$

under the assumption that k fall in the family of shift-invariant kernel. The starting point is a celebrated result that characterizes the class of positive definite functions:

Theorem 1 (Bochner’s theorem [Rudin, 2011]). *A continuous, real valued, symmetric and shift-invariant function k on \mathbb{R}^d is a positive definite kernel if and only if there is a positive finite measure $\mathbb{P}(\mathbf{w})$ such that*

$$k(\mathbf{x} - \mathbf{x}') = \int_{\mathbb{R}^d} 2 [\cos(\mathbf{w}^T \mathbf{x}) \cos(\mathbf{w}^T \mathbf{x}') + \sin(\mathbf{w}^T \mathbf{x}) \sin(\mathbf{w}^T \mathbf{x}')] d\mathbb{P}(\mathbf{w}). \quad (2)$$

The most well-known kernel that belongs to the shift-invariant family is Gaussian kernel $k(\mathbf{x}, \mathbf{x}') = \exp(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\sigma^2})$, the associated density $\mathbb{P}(\mathbf{w})$ is again Gaussian, $N(\mathbf{0}, \sigma^{-2}\mathbf{I}_d)$.

Rahimi and Recht [2007] approximate the integral representation of the kernel (2) by Monte Carlo method as follows,

$$k(\mathbf{x} - \mathbf{x}') \approx \frac{1}{M} \sum_{m=1}^M h_{\mathbf{x}, \mathbf{x}'}(\mathbf{w}_m) = \langle z(\mathbf{x}), z(\mathbf{x}') \rangle, \quad (3)$$

where

$$\begin{aligned} h_{\mathbf{x}, \mathbf{x}'}(\mathbf{w}) &:= 2 [\cos(\mathbf{w}^T \mathbf{x}) \cos(\mathbf{w}^T \mathbf{x}') + \sin(\mathbf{w}^T \mathbf{x}) \sin(\mathbf{w}^T \mathbf{x}')] \\ z(\mathbf{x}) &:= \frac{1}{\sqrt{M}} [\phi_{\mathbf{w}_1}(\mathbf{x}), \dots, \phi_{\mathbf{w}_M}(\mathbf{x})] \\ \phi_{\mathbf{w}}(\mathbf{x}) &:= \sqrt{2} [\cos(\mathbf{w}^T \mathbf{x}), \sin(\mathbf{w}^T \mathbf{x})] \end{aligned} \quad (4)$$

and $\mathbf{w}_1, \dots, \mathbf{w}_M$ are i.i.d. samples from $\mathbb{P}(\mathbf{w})$.

After obtaining the randomized feature map z , training data could be transformed into $\{(z(\mathbf{x}_i), y_i)\}_{i=1}^n$. As long as M is sufficiently smaller than n , this leads to more scalable solutions, e.g., for regression we get back to $O(nM^2)$ training and $O(Md)$ prediction time, with $O(nM)$ memory requirements. We can also apply random features to unsupervised learning problems, such as kernel clustering [Chitta *et al.*, 2012].

2.3 Quasi-Monte Carlo Technique

Instead of using plain MC approximation, Yang *et al.* [2014] propose to use the low-discrepancy properties of Quasi-Monte Carlo (QMC) sequences to reduce the integration error in approximations of the form (3). Due to space limitation, we restrict our discussion to the background that is necessary for understanding subsequent sections. We refer interested readers to the comprehensive reviews [Caflish, 1998; Dick *et al.*, 2013] for more detailed exposition.

The QMC method is generally applicable to integrals over a unit cube. So the procedure is to first generate a discrepancy sequence $\mathbf{t}_1, \dots, \mathbf{t}_M \in [0, 1]^d$, and transform it into a sequence $\mathbf{w}_1, \dots, \mathbf{w}_M$ in \mathbb{R}^d . To convert the integral presentation of kernel (2) to an integral over the unit cube, a simple change of variables suffices. For $\mathbf{t} \in \mathbb{R}^d$, define $\Phi^{-1}(\mathbf{t}) = [\Phi_1^{-1}(t_1), \dots, \Phi_d^{-1}(t_d)] \in \mathbb{R}^d$, where we assume that the density function in (2) can be written as $p(\mathbf{w}) = \prod_{j=1}^M p_j(\mathbf{w}_j)$ and ϕ_j is the cumulative distribution function (CDF) of $p_j, j = 1, \dots, d$. By setting $\mathbf{w} = \Phi^{-1}(\mathbf{t})$, the integral (2) is equivalent to

$$k(\mathbf{x} - \mathbf{x}') = \int_{\mathbb{R}^d} h_{\mathbf{x}, \mathbf{x}'}(\mathbf{w}) p(\mathbf{w}) d\mathbf{w} = \int_{[0, 1]^d} h_{\mathbf{x}, \mathbf{x}'}(\Phi^{-1}(\mathbf{t})) d\mathbf{t}.$$

3 Bayesian Quadrature for Random Features

Random features constructed by either MC [Rahimi and Recht, 2007] or QMC [Yang *et al.*, 2014] approaches employ equal weights $1/M$ on integrand functions as shown in (3). An important question is how to construct different weights on the integrand functions to have a better kernel approximation.

Given a fixed QMC sequence, Avron *et al.* [2016b] solve the weights by optimizing a derived error bound based on the observed data. There are two potential weakness of [Avron *et al.*, 2016b]. First, Avron *et al.* [2016b] does not fully utilize the data information because the optimization objective only depends on the upper bound information of $|\mathbf{x}_i|$ instead of

the distribution of \mathbf{x}_i . Second, the error bound used in Avron *et al.* [2016b] is potentially loose. Due to technical reasons, Avron *et al.* [2016b] relaxes $h_{\mathbf{x},\mathbf{x}'}(\mathbf{w})$ to $h_{\mathbf{x},\mathbf{x}'}(\mathbf{w})\text{sinc}(T\mathbf{w})$ when they derive the error bound to optimize. It is not studied whether the relaxation results in a proper upper bound.

On the other hand, selecting weights by such way is closely connected to the Bayesian Quadrature (BQ), originally proposed by Ghahramani and Rasmussen [2002] and theoretically guaranteed under Bayes assumptions. BQ is a Bayesian approach that puts prior knowledge on functions to solve the integral estimation problem. We first introduce the work of Ghahramani and Rasmussen [2002], and then discuss how to apply BQ to kernel approximations.

Standard BQ considers a single integration problem with form

$$\bar{f} = \int f(\mathbf{w})d\mathbb{P}(\mathbf{w}). \quad (5)$$

To utilize the information of f , BQ puts a Gaussian Process (GP) prior on f with mean 0 and covariance matrix K_{GP} where

$$K_{GP}(\mathbf{w}, \mathbf{w}') = \text{Cov}(f(\mathbf{w}), f(\mathbf{w}')) = \exp\left(-\frac{\|\mathbf{w} - \mathbf{w}'\|_2^2}{2\sigma_{GP}^2}\right).$$

After conditioning on sample $\mathbf{w}_1, \dots, \mathbf{w}_M$ from \mathbb{P} , we obtain a closed-form GP posterior over f and use the posterior to get the optimal Bayes estimator \hat{f}^1 as

$$\begin{aligned} \hat{f} &= \mathbb{E}_{GP} \left[\int f(\mathbf{w})d\mathbb{P}(\mathbf{w}) \right] \\ &= \boldsymbol{\gamma}^\top K_{GP}^{-1} f(W) = \sum_{m=1}^M \beta_{BQ}^{(m)} f(\mathbf{w}_m), \end{aligned}$$

where $\gamma_m = \int K_{GP}(\mathbf{w}, \mathbf{w}_m)d\mathbb{P}(\mathbf{w})$, $\beta_{BQ}^m = (K_{GP}^{-1}\boldsymbol{\gamma})_m$, and $f(W) = [f(\mathbf{w}_1), \dots, f(\mathbf{w}_M)]$. Clearly, the resulting estimator is simply a new weighted combination of the empirical observations $f(\mathbf{w}_m)$ by $\beta_{BQ}^{(m)}$ derived from the Bayes perspective.

In kernel approximation, we could treat it as a series of the single integration problem. Therefore, we could directly apply BQ on it. Here we discuss the technical issues of applying BQ on kernel approximation. First, we need to evaluate $\boldsymbol{\gamma}$. For Gaussian kernel, we derive the closed-form expression

$$\gamma_m = \int K_{GP}(\mathbf{w}, \mathbf{w}_m)\mathbb{P}(\mathbf{w}) = \exp\left(\frac{-\|\mathbf{w}_m - \mathbf{0}\|_2^2}{\sigma_{GP}^2 + \sigma^{-2}}\right)$$

where γ_m is a random variable evaluated at \mathbf{w}_m .

3.1 Data-driven by Kernel Hyperparameters

The most import issue of applying BQ on kernel approximation is that it models the information of $h_{\mathbf{x},\mathbf{x}'}$ by the kernel K_{GP} . In practice, we usually use Gaussian kernel for K_{GP} , then the feature information of $(\mathbf{x}, \mathbf{x}')$ is modeling by the kernel bandwidth. As suggested by Rasmussen [2006], one could tune the best hyper-parameter either by maximizing the marginal likelihood or cross-validation based on the metrics

¹For the square loss.

in interested. Note that the existing BQ approximates the integral representation of a kernel function value, evaluated at a single pair of data $(\mathbf{x}, \mathbf{x}')$. However, it is unfeasible to tune individual σ_{GP} for n^2 pairs of data. Therefore, we tune a global hyper-parameter σ_{GP} on a subset of pair data points.

4 Stein Effect Shrinkage (SES) Estimator

In practice, it is well known that the performance of Gaussian Process based models heavily depend on hyper-parameter σ_{GP} , i.e. the bandwidth of kernel function k_{GP} . Therefore, instead of using Bayesian way to obtain the optimal Bayes estimator from BQ, we start from the risk minimization perspective by considering the *Stein effect* to derive the *shrinkage estimator* for random features.

The standard estimator of Monte-Carlo methods is the unbiased empirical average using equal weights $\frac{1}{M}$ for M samples which sum to 1. However, *Stein effect* (or James-Stein effect) suggests a family of biased estimators could result in a lower risk than the standard empirical average. Here we extend the analysis of non-parametric Stein effect to the kernel approximation with random features by considering the risk function $\mathcal{R}(k, k') = \mathbb{E}_{\mathbf{x}, \mathbf{x}', \mathbf{w}}(k(\mathbf{x} - \mathbf{x}') - k'(\mathbf{x} - \mathbf{x}'; \mathbf{w}))^2$.

Theorem 2. (*Stein effect on kernel approximation*) Let $\hat{k}(\mathbf{x} - \mathbf{x}', \mathbf{w}) = \frac{1}{M} \sum_{m=1}^M h_{\mathbf{x}, \mathbf{x}'}(\mathbf{w}_m)$. For any estimator μ , which is independent of \mathbf{x} and \mathbf{w} , there exists $0 \leq \alpha < 1$ such that $\tilde{k}(\cdot) = \alpha\mu + (1 - \alpha)\hat{k}(\cdot)$ is an estimator with lower risk $\mathcal{R}(k, \tilde{k}) < \mathcal{R}(k, \hat{k})$.

Proof. The proof closely follows Muandet *et al.* [2014] for a different problem under the non-parametric setting². By bias-variance decomposition, we have

$$\begin{aligned} \mathcal{R}(k, \hat{k}) &= \mathbb{E}_{\mathbf{x}, \mathbf{x}'} [\text{Var}_{\mathbf{w}}(\hat{k}(\mathbf{x} - \mathbf{x}', \mathbf{w}))] \\ \mathcal{R}(k, \tilde{k}) &= (1 - \alpha)^2 \mathbb{E}_{\mathbf{x}, \mathbf{x}'} [\text{Var}_{\mathbf{w}}(\hat{k}(\mathbf{x} - \mathbf{x}', \mathbf{w}))] \\ &\quad + \alpha^2 \mathbb{E}_{\mathbf{x}, \mathbf{x}'} [(\mu - k(\mathbf{x} - \mathbf{x}'))^2]. \end{aligned}$$

By elementary calculation, we have $\mathcal{R}(k, \tilde{k}) < \mathcal{R}(k, \hat{k})$ when

$$0 \leq \alpha \leq \frac{2\mathbb{E}_{\mathbf{x}, \mathbf{x}'} [\text{Var}_{\mathbf{w}}(\hat{k}(\mathbf{x} - \mathbf{x}', \mathbf{w}))]}{\mathbb{E}_{\mathbf{x}, \mathbf{x}'} [\text{Var}_{\mathbf{w}}(\hat{k}(\mathbf{x} - \mathbf{x}', \mathbf{w}))] + \mathbb{E}_{\mathbf{x}, \mathbf{x}'} [(\mu - k(\mathbf{x} - \mathbf{x}'))^2]}$$

and the optimal value happened at

$$\alpha^* = \frac{\mathbb{E}_{\mathbf{x}, \mathbf{x}'} [\text{Var}_{\mathbf{w}}(\hat{k}(\mathbf{x} - \mathbf{x}', \mathbf{w}))]}{\mathbb{E}_{\mathbf{x}, \mathbf{x}'} [\text{Var}_{\mathbf{w}}(\hat{k}(\mathbf{x} - \mathbf{x}', \mathbf{w}))] + \mathbb{E}_{\mathbf{x}, \mathbf{x}'} [(\mu - k(\mathbf{x} - \mathbf{x}'))^2]} < 1$$

□

Corollary 2.1. (*Shrinkage estimator*) If $k(\mathbf{x} - \mathbf{x}') > 0$, there is a shrinkage estimator $\tilde{k}(\cdot) = (1 - \alpha^*)\hat{k}(\cdot)$ that has lower risk than the empirical mean estimator $\hat{k}(\cdot)$ by setting $\mu = 0$.

The standard shrinkage estimator derived from Stein effect uses the equal weights $\frac{1-\alpha}{M}$ to minimize the variance; however, a non-uniform weights usually results in better performance in practice [Huszár and Duvenaud, 2012]. Therefore,

²The original Stein effect has a Gaussian distribution assumption.

based on the analysis of Theorem 2, we proposed the estimator $\tilde{k}(\mathbf{x} - \mathbf{x}', \mathbf{w}) = \sum_{m=1}^M \beta_m h(\mathbf{w}_m)$ obtained by minimizing the empirical risk with the constraint $\sum_{m=1}^M \beta_m^2 \leq c$ to shrink the weights, where c is a small constant. With Lagrangian multiplier, we end up the following risk minimization formulation which is *data-driven* by directly taking $k(\mathbf{x}, \mathbf{x}')$ into account,

$$\min_{\beta \in \mathbb{R}^M} \sum_{(\mathbf{x}, \mathbf{x}') \sim D} \left[k(\mathbf{x}, \mathbf{x}') - [z(\mathbf{x}) \circ z(\mathbf{x}')]^T \beta \right]^2 + \lambda_\beta \|\beta\|^2, \quad (6)$$

where \circ denotes the Hadamard product, and λ is the regularization coefficient to shrink the weights.

The objective function (6) can be further simplified as a least squares regression (LSR) formulation,

$$\min_{\beta \in \mathbb{R}^M} \|\mathbf{t} - Z\beta\|^2 + \lambda_\beta \beta^T \beta, \quad (7)$$

where $\mathbf{t} \in \mathbb{R}^{n^2}$, $Z \in \mathbb{R}^{n^2 \times M}$, such that for all $i, j \in P = \{1, \dots, n\}$, there exist a corresponding pair mapping function $\xi : \xi_P(i, j) \rightarrow k$ satisfying $t_k := \phi(\mathbf{x}_i)^T \phi(\mathbf{x}_j)$ and $Z(k, :) := \mathbf{z}(\mathbf{x}_i) \circ \mathbf{z}(\mathbf{x}_j)$.

We approximately solve optimization problem (7) via the matrix sketching techniques [Avron *et al.*, 2013; Avron *et al.*, 2016a; Woodruff and others, 2014]:

$$\min_{\beta \in \mathbb{R}^M} \|\mathbf{S}\mathbf{t} - \mathbf{S}Z\beta\|^2 + \lambda_\beta \beta^T \beta, \quad (8)$$

where $\mathbf{S} \in \mathbb{R}^{r \times n^2}$ is the randomized sketching matrix. Solving the sketched LSR problem (8) now costs only $O(rd^2 + T_s)$ where T_s is the time cost of sketching. From the optimization perspective, suppose β^* be the optimal solution of LSR problem (7), and $\tilde{\beta}$ be the solution of sketched LSR problem (8). [Wang *et al.*, 2017] prove that if $r = O(M/\epsilon + \text{poly}(M))$, the objective function value of (8) at $\tilde{\beta}$ is at most ϵ worse than the value of (7) at β^* . In practice, we consider an efficient $T_s = O(r)$ sampling-based sketching matrix where \mathbf{S} is a diagonal matrix with $S_{i,i} = 1/p_i$ if row i was include in the sample, and $S_{i,i} = 0$ otherwise. That being said, we form a sub-matrix of Z by including each row of Z in the sub-matrix independently with probability p_i . [Woodruff and others, 2014] suggests to set p_i proportional to $\|Z_i\|$, the norm of i th row of Z , to have a meaningful error bound.

Algorithm 1 Data-driven random feature (SES)

- 1: **procedure** SES
- 2: Compute the Fourier transform p of the kernel
- 3: Generate a sequence $\mathbf{w}_1, \dots, \mathbf{w}_M$ by QMC or MC
- 4: Compute un-weighted random feature $z(\mathbf{x})$ by (4)
- 5: Calculate shrinkage weight by (8)
- 6: Compute data-driven random features by

$$z'(\mathbf{x}) = \text{diag}(\sqrt{\beta})z(\mathbf{x})$$

5 Experiments

In this section, we empirically demonstrate the benefits of our proposed method on six data sets listed in Table 1. The

Dataset	Task	n	d
CPU	regression	6554	21
Census	regression	18,186	119
Years	regression	463,715	90
Adult	classification	32,561	123
MNIST	classification	60,000	778
Covtype	classification	581,012	54

Table 1: Data Statistics. n is number of instances and d is dimension of instances.

features in all the six data sets are scaled to zero mean and unit variance in a preprocessing step.

We consider two different tasks: i) the quality of kernel approximation and ii) the applicability in supervised learning. Both the relative kernel approximation errors and the regression/classification errors are reported on the test sets.

5.1 Methods for Comparison

- MC: Monte Carlo approximation with uniform weight.
- QMC: Quasi-Monte Carlo approximation with uniform weight. We adapt the implementation of Yang *et al.* [2014] on Github. ³
- BQ: QMC sequence with Bayesian Quadrature weights. We modify the source code from Huszár and Duvenaud [2012]. ⁴

Gaussian RBF kernel is used through all the experiments where the kernel bandwidth σ is tuned on the set $\{2^{-10}, 2^{-8}, \dots, 2^8, 2^{10}\}$ that is in favor of Monte Carlo methods. For QMC method, we use scrambling and shifting techniques recommended in the QMC literature (See Dick *et al.* [2013] for more details). In BQ framework, we consider the GP prior with the covariance matrix $k_{GP}(\mathbf{w}, \mathbf{w}') = \exp(-\frac{\|\mathbf{w} - \mathbf{w}'\|_2^2}{2\sigma_{GP}^2})$ where σ_{GP} is tuned on the set $\{2^{-8}, 2^{-6}, \dots, 2^6, 2^8\}$ over a subset of sampled pair data $(\mathbf{x}, \mathbf{x}')$. Likewise, regularization coefficient λ_β of in the proposed objective function (6) is tuned on the set $\{2^{-8}, 2^{-6}, \dots, 2^6, 2^8\}$ over the same subset of sampled pair data. For regression task, we solve ridge regression problems where the regularization coefficient λ is tuned by five folds cross-validation. For classification task, we solve L2-loss SVM where the regularization coefficient C is tuned by five folds cross-validation. All experiments code are written in MATLAB and run on a Intel(R) Xeon(R) CPU 2.40GHz Linux server with 32 cores and 190GB memory.

5.2 Quality of Kernel Approximation

In our setting, the most natural metric for comparison is the quality of approximation of the kernel matrix. We use the relative kernel approximation error $\|K - \tilde{K}\|_F / \|K\|_F$ to measure the quality, as shown in Figure 1.

SES outperforms the other methods on cpu, census, adult, mnist, covtype, and is competitive with the best baseline methods over the years data sets. Notice that SES is particularly

³<https://github.com/chocjy/QMC-features>

⁴<http://www.cs.toronto.edu/duvenaud/>

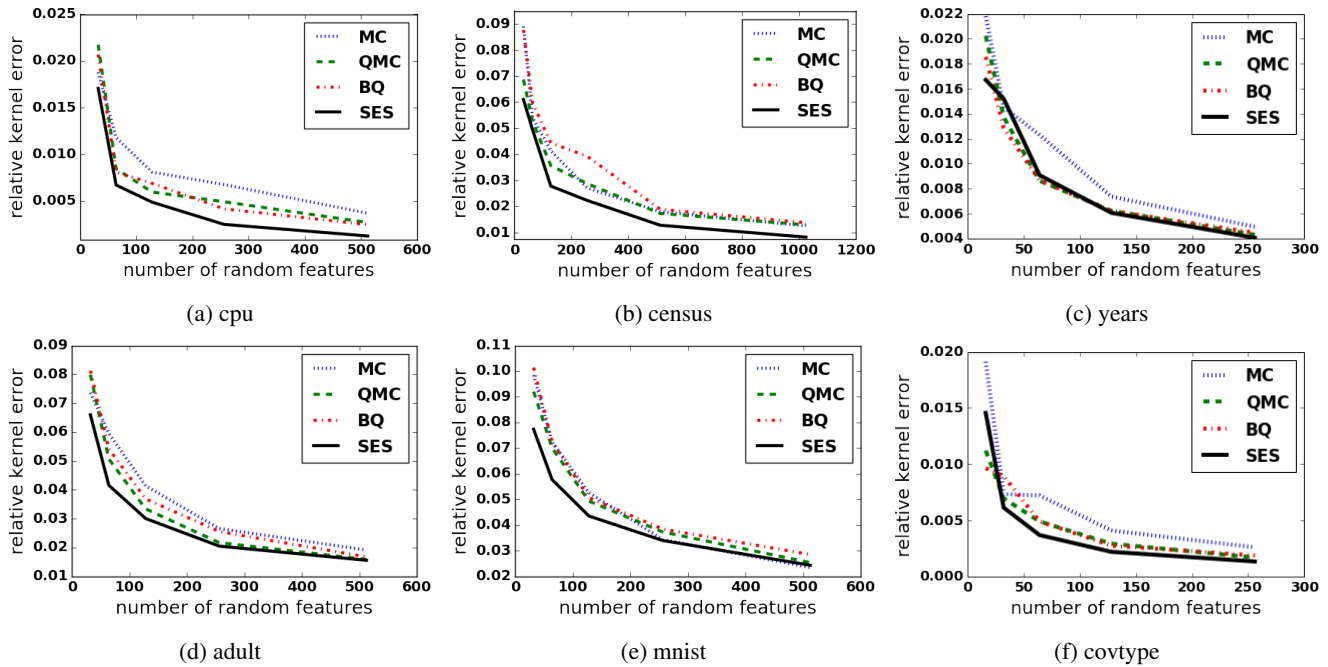


Figure 1: Kernel approximation results. x -axis is number of samples used to approximate the integral and y -axis shows the relative kernel approximation error.

strong when the number of random features is relatively small. This phenomenon confirms our theoretical analysis, that is, the smaller the random features are, the larger the variance would be in general. Our shrinkage estimator with Stein effect enjoys better performance than the empirical mean estimators (MC and QMC) using uniformly weighted features without shrinkage. Although BQ does use weighted random features, its performance is still not as good as SES because tuning the covariance matrix in the Gaussian Process is rather difficult. Moreover, it is even harder to tune a global bandwidth for all integrand functions that approximate the kernel function.

5.3 Supervised Learning

For regression task, we report the relative regression error $\|\mathbf{y} - \tilde{\mathbf{y}}\|_2 / \|\mathbf{y}\|_2$. For classification task, we report the classification error.

We investigate if better kernel approximations reflect better predictions in the supervised learning tasks. The results are shown in Table 2. We can see that among the same data

Dataset	M	MC	QMC	BQ	SES
cpu	512	3.35%	3.29%	3.29%	3.27%
census	256	8.46%	6.60%	6.44%	6.49%
years	128	0.474%	0.473%	0.473%	0.473%
adult	64	16.21%	15.57%	15.87%	15.45%
mnist	256	7.33%	7.72%	7.96%	7.71%
covtype	128	23.36%	22.88%	23.13%	22.67%

Table 2: Supervised learning errors. M is number of random features. Error rate is reported, the lower the better.

set (cpu, census, adult, and covtype) our method performs very well in terms of kernel approximation, and it also yields lower errors in most of the cases of the supervised learning tasks. Note that we only present partial experiment results on some selected number of random features. However, we also observe that the generalization error of our proposed method may be higher than that of other state-of-the-art methods at some other number of random features that we did not present here. This situation is also observed in other works [Avron *et al.*, 2016b]: better kernel approximation does not always result in better supervised learning performance. Under certain conditions, less number of features (worse kernel approximation) can play a similar role of regularization to avoid overfitting [Rudi *et al.*, 2016]. Rudi *et al.* [2016] suggest to solve this issue by tuning number of random features.

5.4 Distributions of Feature Weights

We are interested in the distributions of random feature weights derived from BQ and SES. For different numbers of random features ($M = 32, 128, 512$), we plotted the corresponding histograms and compared them with the equal weight of $1/M$ used in MC and QMC, as shown in Figure 2. When the number of random features is small, the difference between BQ and SES is big. As the number of random features increases, the weights of each method gradually converge to a stable distribution, and get increasingly closer to the uniform weight spike. This empirical observation agrees with the theoretical analysis we found, namely with the role of the weight estimator in the bias-variance trade-off.

5.5 Behavior of Randomized Optimization Solver

In theory, our randomized solver only needs to sample $O(M)$ data in our proposed objective function (8) to have a suffi-

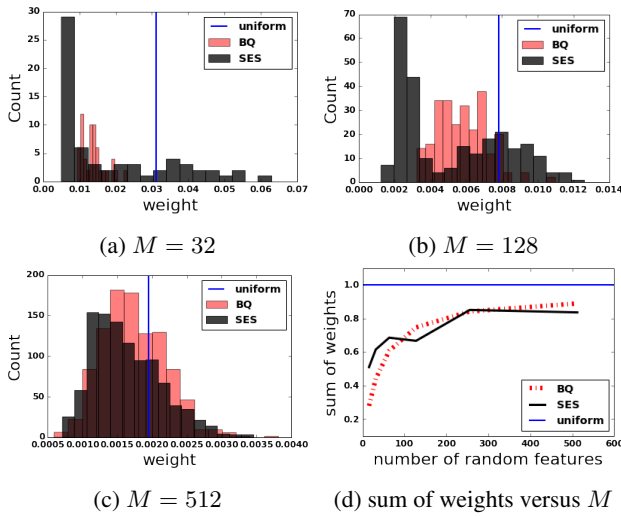


Figure 2: A comparison of distribution of weights on adult data set. M is the number of random features.

ciently accurate solution. We verify this assertion by varying the sampled data size r in sketching matrix S and examine it affects the relative kernel approximation errors. The results are shown in Figure 3. For census data set, we observe that as the number of sampled data r exceeds the number of random features M , the kernel approximation of our proposed method outperforms the state-of-the-art, QMC. Furthermore, the number of sampled data required by our proposed randomized solver is still the same order of M even for a much larger data set covtype, with 0.5 millions data points. On the other hand, BQ does not have consistent behavior regarding the sampled data size, which is used for tuning the hyperparameter σ_{GP} . This again illustrates that BQ is highly sensitive to the covariance matrix and is difficult to tune σ_{GP} on only a small subset of sampled data. In practice, we found that we usually only need to sample up to $2 \sim 4$ times the number of random features M to get a relatively stable and good minimizer β^* . This observation coincides with the complexity analysis of our proposed randomized solver in Section 4.

5.6 Uniform v.s. Non-uniform Shrinkage Weight

We conducted experiments with SES using a uniform shrinkage weight versus using non-uniform shrinkage weights on the cpu and census data sets. The results in Figures 4 confirm our conjecture that non-uniform shrinkage weights are more beneficial, although both enjoy the lower empirical risk from the Stein effect. We also observed similar behavior on other data sets, and we omit those results.

6 Discussion and Conclusion

Nyström methods [Williams and Seeger, 2001; Drineas and Mahoney, 2005] is yet another line of research for kernel approximation. Specifically, eigendecomposition of the kernel matrix is achieved via low rank factorization with various sampling techniques [Wang and Zhang, 2013]. We refer interested readers to the comprehensive study [Yang *et al.*, 2012] for more detailed comparisons. Due to the page limits, the

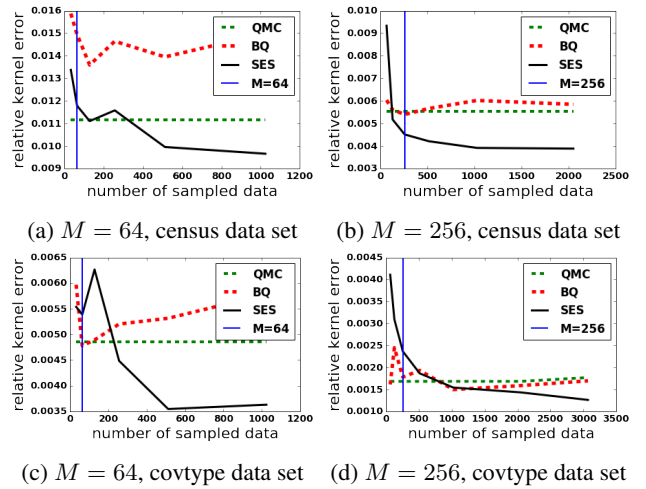


Figure 3: A comparison of different sampled size training data in solving objective function (8).




Figure 4: Comparison of uniform and non-uniform shrinkage weight.

scope of this paper only focus on the study of random Fourier features.

Finally, we notice [Sinha and Duchi, 2016] proposes a weighting scheme for random features which matches the label correlation matrix with the goal to optimize the supervise objective (e.g. classification errors). In contrast, our work aims to optimize the kernel approximation, which does not require the label information to solve weights and can be applied to different tasks after solving the weights once.

To conclude, we propose a novel shrinkage estimator, which enjoys the benefits of Stein effect w.r.t. lowering the empirical risk compared to empirical mean estimators. Our estimator induces non-uniform weights on the random features with desirable data-driven properties. We further provide an efficient randomized solver to optimize this estimation problem for real-world large-scale applications. The extensive experimental results demonstrate the effectiveness of our proposed method.

Acknowledgments

We thank the reviewers for their helpful comments. This work is supported in part by the National Science Foundation (NSF) under grants IIS-1546329 and IIS-1563887.

References

- [Avron *et al.*, 2013] Haim Avron, Vikas Sindhwani, and David Woodruff. Sketching structured matrices for faster nonlinear regression. In *NIPS*, 2013.
- [Avron *et al.*, 2016a] Haim Avron, Kenneth L Clarkson, and David P Woodruff. Faster kernel ridge regression using sketching and preconditioning. *arXiv*, 2016.
- [Avron *et al.*, 2016b] Haim Avron, Vikas Sindhwani, Jiyan Yang, and Michael W Mahoney. Quasi-monte carlo feature maps for shift-invariant kernels. *JMLR*, 2016.
- [Bach, 2015] Francis Bach. On the equivalence between quadrature rules and random features. *arXiv*, 2015.
- [Caflish, 1998] Russel E Caflish. Monte carlo and quasi-monte carlo methods. *Acta numerica*, 1998.
- [Chang and Lin, 2011] Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2011.
- [Chitta *et al.*, 2012] Radha Chitta, Rong Jin, and Anil K Jain. Efficient kernel clustering using random fourier features. In *ICDM*, 2012.
- [Dai *et al.*, 2014] Bo Dai, Bo Xie, Niao He, Yingyu Liang, Anant Raj, Maria-Florina F Balcan, and Le Song. Scalable kernel methods via doubly stochastic gradients. In *NIPS*, 2014.
- [Dick *et al.*, 2013] Josef Dick, Frances Y Kuo, and Ian H Sloan. High-dimensional integration: the quasi-monte carlo way. *Acta Numerica*, 2013.
- [Drineas and Mahoney, 2005] Petros Drineas and Michael W Mahoney. On the nyström method for approximating a gram matrix for improved kernel-based learning. *JMLR*, 2005.
- [Ghahramani and Rasmussen, 2002] Zoubin Ghahramani and Carl E Rasmussen. Bayesian monte carlo. In *NIPS*, 2002.
- [Huang *et al.*, 2014] Po-Sen Huang, Haim Avron, Tara N Sainath, Vikas Sindhwani, and Bhuvana Ramabhadran. Kernel methods match deep neural networks on timit. In *ICASSP*, 2014.
- [Huszár and Duvenaud, 2012] Ferenc Huszár and David Duvenaud. Optimally-weighted herding is bayesian quadrature. In *UAI*, 2012.
- [Kivinen *et al.*, 2004] Jyrki Kivinen, Alexander J Smola, and Robert C Williamson. Online learning with kernels. *IEEE transactions on signal processing*, 52(8), 2004.
- [Le *et al.*, 2013] Quoc Le, Tamás Sarló, and Alexander Smola. Fastfood-computing hilbert space expansions in loglinear time. In *ICML*, 2013.
- [Li and Póczos, 2016] Chun-Liang Li and Barnabás Póczos. Utilize old coordinates: Faster doubly stochastic gradients for kernel methods. In *UAI*, 2016.
- [Muandet *et al.*, 2014] K. Muandet, K. Fukumizu, B. Sriperumbudur, A. Gretton, and B. Schölkopf. Kernel mean estimation and stein effect. In *ICML*, 2014.
- [Rahimi and Recht, 2007] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *NIPS*, 2007.
- [Rahimi and Recht, 2009] Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *NIPS*, 2009.
- [Rasmussen, 2006] Carl Edward Rasmussen. Gaussian processes for machine learning. 2006.
- [Rudi *et al.*, 2016] Alessandro Rudi, Raffaello Camoriano, and Lorenzo Rosasco. Generalization properties of learning with random features. *arXiv*, 2016.
- [Rudin, 2011] Walter Rudin. *Fourier analysis on groups*. John Wiley & Sons, 2011.
- [Sinha and Duchi, 2016] Aman Sinha and John C Duchi. Learning kernels with random features. In *NIPS*, 2016.
- [Steinwart and Christmann, 2008] Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.
- [Wang and Zhang, 2013] Shusen Wang and Zhihua Zhang. Improving cur matrix decomposition and the nyström approximation via adaptive sampling. *JMLR*, 2013.
- [Wang *et al.*, 2017] Shusen Wang, Alex Gittens, and Michael W Mahoney. Sketched ridge regression: Optimization perspective, statistical perspective, and model averaging. *arXiv*, 2017.
- [Williams and Seeger, 2001] Christopher Williams and Matthias Seeger. Using the nyström method to speed up kernel machines. In *NIPS*, 2001.
- [Woodruff and others, 2014] David P Woodruff et al. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 2014.
- [Yang *et al.*, 2012] 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 *NIPS*, 2012.
- [Yang *et al.*, 2014] Jiyan Yang, Vikas Sindhwani, Haim Avron, and Michael W Mahoney. Quasi-monte carlo feature maps for shift-invariant kernels. In *ICML*, 2014.