

Randomly pivoted Cholesky: Practical approximation of a kernel matrix with few entry evaluations*

Yifan Chen[†], Ethan N. Epperly[†], Joel A. Tropp[†], and Robert J. Webber[†]

Abstract. Randomly pivoted Cholesky (RPCHOLESKY) is a natural algorithm for computing a rank- k approximation of an $N \times N$ positive semidefinite (psd) matrix. RPCHOLESKY can be implemented with just a few lines of code; it requires only $(k+1)N$ entry evaluations and $\mathcal{O}(k^2N)$ additional arithmetic operations. This paper offers the first serious investigation of its experimental and theoretical behavior. Empirically, RPCHOLESKY matches or improves on the performance of alternative algorithms for low-rank psd approximation. Furthermore, RPCHOLESKY provably achieves near-optimal approximation guarantees. The simplicity, effectiveness, and robustness of this algorithm strongly support its use in scientific computing and machine learning applications.

Key words. kernel method, low-rank approximation, Nyström approximation

MSC codes. 65F55, 65C99, 68T05

1. Introduction. Kernel methods [27] are a class of machine learning algorithms that are often used for prediction and clustering tasks. Kernel methods require manipulating a positive semidefinite (psd) *kernel matrix* $\mathbf{A} \in \mathbb{R}^{N \times N}$. The entries $a_{ij} = K(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$ of the kernel matrix are determined by a positive definite *kernel function* $K(\cdot, \cdot)$ that measures the similarity between pairs of data points [27, Def. 2.5]. The dimension N of the kernel matrix is equal to the number of data points, which can easily grow into the millions or higher. The size of modern data sets poses a challenge for kernel methods: even generating all N^2 entries of \mathbf{A} may be infeasible when $N \geq 10^5$.

To speed up kernel computations, a standard practice is to replace the kernel matrix \mathbf{A} with a low-rank approximation. Remarkably, in the common setting in which the $N \times N$ kernel matrix has rapidly decaying eigenvalues, it is possible to generate an accurate approximation using only $\mathcal{O}(N)$ entry evaluations and $\mathcal{O}(N)$ arithmetic operations. Nevertheless, it has been formidable to design a practical and provably justified algorithm that achieves this goal.

The *column Nyström approximation* [19, §19.2] is the most widely used technique for con-

*Submitted to the editors July 13, 2022.

Funding: YC acknowledges support from the Caltech Kortschak Scholar program. ENE acknowledges support from the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0021110. JAT and RJW acknowledge support from the Office of Naval Research through BRC Award N00014-18-1-2363 and from the National Science Foundation through FRG Award 1952777.

Disclaimer: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

[†]Division of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125 USA (yifanc@caltech.edu, epperly@caltech.edu, jtropp@caltech.edu, rwebber@caltech.edu).

structuring a low-rank psd approximation $\widehat{\mathbf{A}}$ of a kernel matrix \mathbf{A} . Given a set $S = \{s_1, \dots, s_k\} \subseteq \{1, \dots, N\}$ of k column indices, the associated Nyström approximation is the rank- k psd matrix

$$(1.1) \quad \widehat{\mathbf{A}} = \mathbf{A}(:, S) \mathbf{A}(S, S)^\dagger \mathbf{A}(S, :) \in \mathbb{R}^{N \times N}.$$

In this expression, $\mathbf{A}(:, S)$ is the submatrix with the selected columns, $\mathbf{A}(S, :)$ is the submatrix with the selected rows, and $\mathbf{A}(S, S)^\dagger$ is the Moore–Penrose pseudoinverse of the submatrix with the selected rows and columns. It can be shown [18, Eq. (1.2)] that the matrix (1.1) is the best psd approximation of \mathbf{A} such that $\text{range}(\widehat{\mathbf{A}}) \subseteq \text{range}(\mathbf{A}(:, S))$ and the residual $\mathbf{A} - \widehat{\mathbf{A}}$ is also psd.

Let us emphasize that the column Nyström approximation (1.1) only requires evaluating kN entries of the matrix \mathbf{A} . As a consequence, this approximation is particularly well-suited for matrices that must be accessed by evaluating entry-wise formulas, such as kernel matrices [27] and Hamiltonian matrices from quantum chemistry [29, Ch. 2].

From a mathematical point of view, algorithms for column Nyström decomposition (1.1) differ primarily in how they select the set S of k columns. The challenge is to identify columns that capture as much action of the matrix as possible. This paper studies a natural approach, which we call the *randomly pivoted Cholesky* (RPCHOLESKY) algorithm (subsection 1.1). Similar approaches have been used in the past to approximate a rectangular matrix by a subset of its rows [8] or to sample from a determinantal point process [22], and a remark of [21] mentions RPCHOLESKY as a possibility for psd low-rank approximation. However, RPCHOLESKY has not been investigated as a serious strategy for psd matrix approximation before now.

Our work provides the first in-depth exploration of the empirical and theoretical behavior of RPCHOLESKY. Across a range of test problems, we find that RPCHOLESKY matches or improves over alternative methods for selecting columns in the Nyström approximation. To explain this surprising observation, we prove that RPCHOLESKY achieves a rigorous error bound that is nearly optimal within the class of Nyström algorithms. Altogether, these results make a strong case that RPCHOLESKY is the method of choice for low-rank approximation of a psd matrix under a constraint on the number of entry evaluations.

1.1. Randomly pivoted Cholesky. We can best understand RPCHOLESKY as a member of a family of algorithms based on *partial Cholesky decomposition*. To that end, observe that the column Nyström approximation (1.1) can be constructed incrementally as follows. Beginning with $\widehat{\mathbf{A}}^{(0)} := \mathbf{0}$ and $\mathbf{A}^{(0)} := \mathbf{A}$, the partial Cholesky decomposition performs the following operations at each step $i = 1, 2, \dots, k$:

1. Select a pivot index $s_i \in \{1, \dots, N\}$.
2. Update the Nyström approximation:

$$\widehat{\mathbf{A}}^{(i)} = \widehat{\mathbf{A}}^{(i-1)} + \frac{\mathbf{A}^{(i-1)}(:, s_i) \mathbf{A}^{(i-1)}(s_i, :)}{a_{s_i s_i}^{(i-1)}}.$$

3. Update the residual (Schur complement):

$$\mathbf{A}^{(i)} = \mathbf{A}^{(i-1)} - \frac{\mathbf{A}^{(i-1)}(:, s_i) \mathbf{A}^{(i-1)}(s_i, :)}{a_{s_i s_i}^{(i-1)}}.$$

By induction, each approximation $\widehat{\mathbf{A}}^{(i)}$ and each residual $\mathbf{A}^{(i)}$ remain psd. After k steps, the partial

Cholesky decomposition $\widehat{\mathbf{A}}^{(k)}$ coincides with the Nyström approximation (1.1) defined by the index set $S = \{s_1, \dots, s_k\}$. With a careful implementation (Algorithm 2.1), steps 2 and 3 involve just kN entry evaluations and $\mathcal{O}(k^2 N)$ arithmetic operations.

There are several standard methods for choosing the set S of pivots. A classical strategy is the *greedy method* [9], also known as diagonal or complete pivoting [16]. At each step i , the pivot s_i is the index of the largest diagonal entry of the current residual matrix, with ties broken arbitrarily:

$$(1.2) \quad s_i \in \operatorname{argmax}\{a_{jj}^{(i-1)} : 1 \leq j \leq N\}.$$

Another common strategy is to select each pivot s_i *uniformly at random* from $\{1, \dots, N\}$, either with or without replacement [17, 34]. However, both the greedy method and the uniform sampling method can fail to find salient columns. The greedy method cannot explore the small diagonal entries of the residual matrix, while the uniform method cannot take advantage of the largest diagonal entries (subsections 4.1 and 4.2).

To circumvent this problem, we consider RPCHOLESKY as an alternative strategy for selecting the pivots. At each step i , this algorithm draws a random pivot s_i according to the probability distribution

$$(1.3) \quad \mathbb{P}\{s_i = j\} = \frac{a_{jj}^{(i-1)}}{\operatorname{tr} \mathbf{A}^{(i-1)}} \quad \text{for each } j = 1, \dots, N.$$

In other words, each pivot is sampled in proportion to the diagonal entries of the current residual matrix. RPCHOLESKY thus strikes a balance between exploring small diagonal entries and exploiting large ones. RPCHOLESKY was mentioned in [21] as a possible strategy for psd low-rank approximation, but we are not aware of any earnest investigation of its behavior.

Algorithm 1.1 describes a naive implementation of RPCHOLESKY, while subsection 2.1 presents a more efficient implementation.

Algorithm 1.1 RPCHOLESKY: Naive implementation

Input: Psd matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, iteration count k or tolerance η

Output: Rank- k psd approximation $\widehat{\mathbf{A}}^{(k)}$

Initialize $\widehat{\mathbf{A}}^{(0)} \leftarrow \mathbf{0}_{N \times N}$ and $\mathbf{A}_0 \leftarrow \mathbf{A}$

for $i = 1$ to k **do**

 Sample $s_i \sim \operatorname{diag} \mathbf{A}^{(i-1)} / \operatorname{tr} \mathbf{A}^{(i-1)}$

$\widehat{\mathbf{A}}^{(i)} \leftarrow \widehat{\mathbf{A}}^{(i-1)} + \mathbf{A}^{(i-1)}(:, s_i) \mathbf{A}^{(i-1)}(s_i, :)/a_{s_i s_i}^{(i-1)}$

$\mathbf{A}^{(i)} \leftarrow \mathbf{A}^{(i-1)} - \mathbf{A}^{(i-1)}(:, s_i) \mathbf{A}^{(i-1)}(s_i, :)/a_{s_i s_i}^{(i-1)}$

end for

▷ Alternately, run until $\operatorname{tr} \mathbf{A}^{(i)} \leq \eta \operatorname{tr} \mathbf{A}$

 ▷ Sample new pivot

▷ Update low-rank approximation

▷ Update Schur complement

1.2. Numerical results. Given that RPCHOLESKY is a natural algorithm, we were perplexed by the lack of attention that it has received in the literature on randomized matrix computations. Therefore, we undertook a suite of experiments to compare RPCHOLESKY with other column Nyström approximations on a range of test problems. Surprisingly, we discovered that RPCHOLESKY reliably produces some of the best matrix approximations using the fewest number of entry evaluations. This work is documented in section 2.

For now, let us offer a visually appealing example that is representative of our broader findings. We aim to approximate the kernel matrices associated with two data sets:

1. **Smile:** A Gaussian kernel matrix constructed from 10^4 data points depicting a smile in \mathbb{R}^2 . The eyes are constructed from just 10^2 points, making them easy to miss for certain sampling methods.
2. **Outliers:** A Gaussian kernel matrix constructed from a point cloud of 10^4 points in \mathbb{R}^{20} . Of these points, 50 are outliers that are positioned very far away from the others.

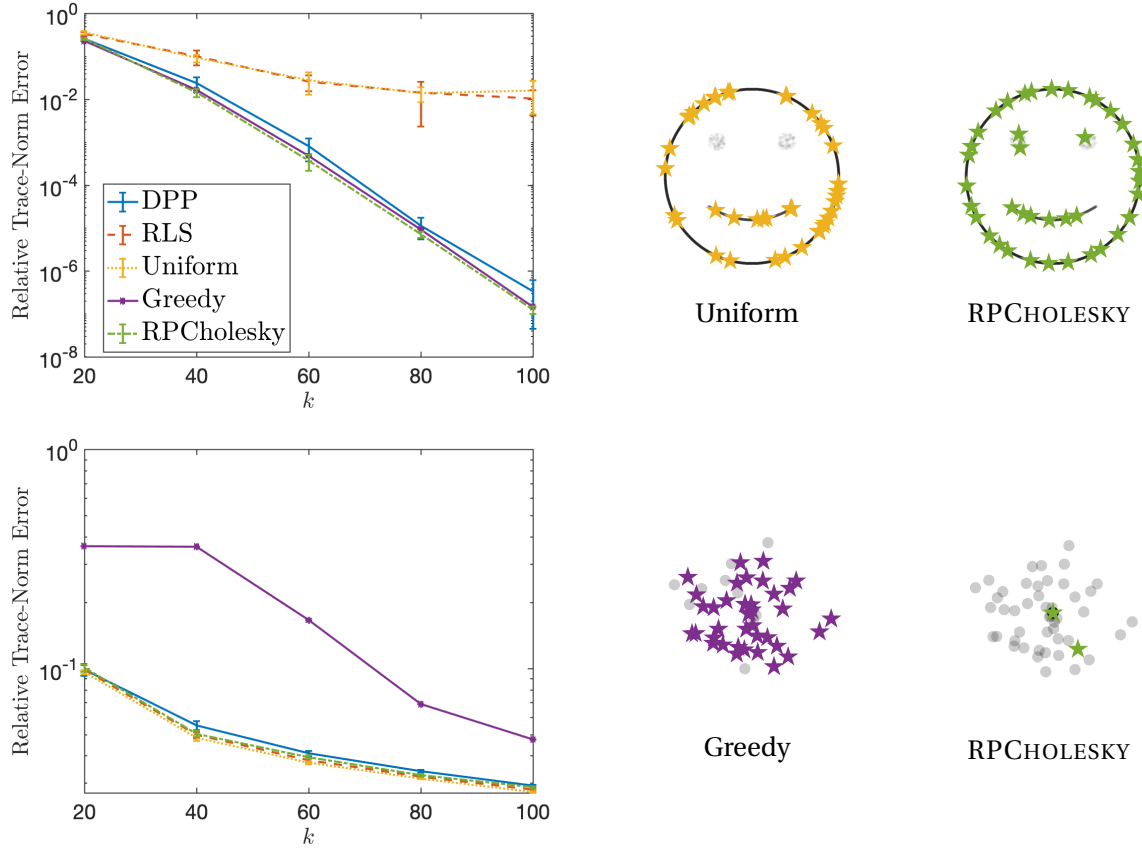


Figure 1: *Left:* Mean relative trace-norm error $\text{tr}(A - \widehat{A}^{(k)}) / \text{tr} A$ and standard deviation error bars for different Nyström methods. *Right:* Selected pivots (colored stars) and data points (grey circles) for the **Smile** and **Outliers** data, with **Outliers** data projected onto a random plane. The greedy method chooses many pivots from the 50 outliers, while RPCHOLESKY mostly avoids the outliers.

Figure 1 charts the accuracy of the column Nyström approximations obtained with different pivoting strategies. The left-hand panels show the average relative trace-norm error over 100 independent trials. The right-hand panels show the pivots selected by several Nyström methods with $k = 40$. For additional computational details, see [Appendix B](#).

These tests bring to light the failure modes of uniform sampling and the greedy method. Uni-

form sampling fails to select the pivots representing small, important features of the data, such as the eyes in the **Smile** example. The greedy method heavily emphasizes outliers, leading to poor approximation accuracy in the **Outliers** example. By contrast, RPCHOLESKY avoids these failure modes and achieves high accuracy for both test problems.

For completeness, we also compare RPCHOLESKY with two other column selection schemes for Nyström approximation: *determinantal point process* (DPP) sampling [7] and *ridge leverage score* (RLS) sampling [1, 4, 20, 26]. These methods offer strong theoretical guarantees, yet existing implementations are complicated and require far more entry evaluations than RPCHOLESKY to achieve the same approximation quality. See subsections 4.3 and 4.4 for full details about these approaches.

Figure 2 displays the computational cost of different column Nyström methods, as measured by the number of entry evaluations. This is the total number of entry evaluations required to first form the sampling distribution and then generate the sampled columns. Uniform sampling is the cheapest Nyström method, requiring just kN entry evaluations. The greedy method and RPCHOLESKY follow closely behind, requiring $(k+1)N$ entry evaluations. RLS sampling requires roughly $3kN$ entry evaluations and DPP sampling (using the vfx sampler [11]) requires between $80kN$ and $1000kN$ entry evaluations, making these latter two methods comparatively very expensive.

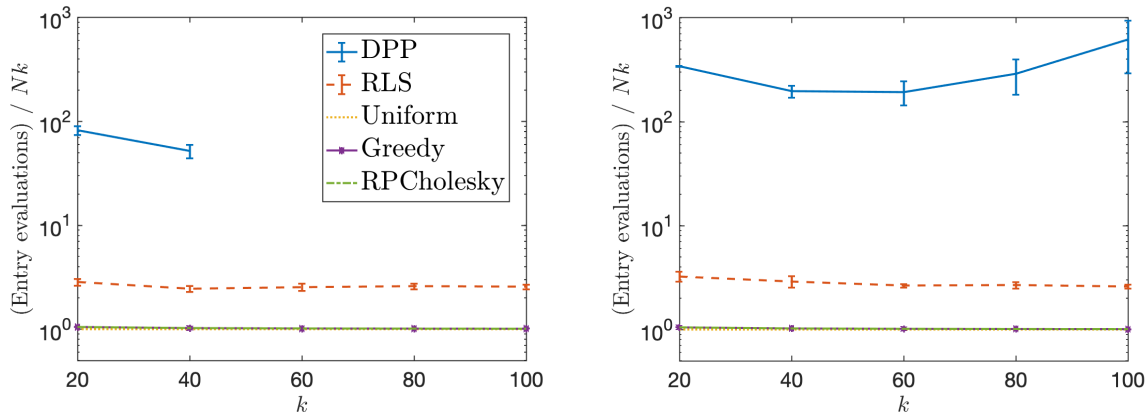


Figure 2: Entry evaluations for different Nyström methods applied to the **Smile** (left) and **Outliers** (right) matrices. The vfx DPP sampler fails for $k \geq 60$ on the **Smile** matrix.

In addition to being expensive, DPP sampling and RLS sampling exhibit other limitations. For the **Smile** kernel matrix with $k \geq 60$, DPP sampling using the vfx sampler [11] returns an error message due to the matrix A being close to low-rank and fails to produce any output. It is necessary to use a slower GS sampler [11] based on a complete eigendecomposition of the matrix A . For the **Smile** kernel matrix, RLS sampling provides a low-quality approximation, often failing to include data points from one or both eyes. The poor RLS sampling results are consistent with the limited theoretical error bounds for this method (see subsection 4.4), which necessitate using hundreds of columns to guarantee an accurate low-rank approximation for any psd matrix.

1.3. Theoretical results. After observing the excellent performance of the RPCHOLESKY algorithm empirically, we sought a more rigorous explanation. This paper proves that RPCHOLESKY

attains error bounds that are nearly optimal within the class of column Nyström approximations. [Section 3](#) studies the theoretical performance of RPCHOLESKY, and [section 4](#) outlines results for competing algorithms. We summarize the comparison here.

Let \mathbf{A} be a psd matrix, and let $\widehat{\mathbf{A}}$ be a column Nyström approximation [\(1.1\)](#) of \mathbf{A} . We consider the approximation error as $\text{tr}(\mathbf{A} - \widehat{\mathbf{A}})$, which is variously known as the trace norm, nuclear norm, or Schatten 1-norm of the residual $\mathbf{A} - \widehat{\mathbf{A}}$. We compare against the error attained by a best rank- r psd approximation $\llbracket \mathbf{A} \rrbracket_r$, obtained from an r -truncated eigenvalue decomposition.

A randomized column Nyström approximation $\widehat{\mathbf{A}}$ is called an (r, ε) -approximation of \mathbf{A} when

$$(1.4) \quad \mathbb{E} \text{tr}(\mathbf{A} - \widehat{\mathbf{A}}) \leq (1 + \varepsilon) \cdot \text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)$$

for parameters $r \in \mathbb{N}$ and $\varepsilon > 0$. The expectation averages over the random choices in the algorithm (e.g., the columns that are sampled). The efficiency of a particular sampling scheme depends on how large we must take the number k of columns. Here is the key theoretical question:

Question 1.1. *How many columns k are sufficient to guarantee that a randomized column Nyström method attains an (r, ε) -approximation [\(1.4\)](#) for every $N \times N$ psd input matrix?*

To achieve an (r, ε) -approximation for the worst-case matrix, a column Nyström approximation must use at least $k \geq r/\varepsilon$ columns. See [Theorem 4.3\(c\)](#) for an example. We will establish the remarkable fact that RPCHOLESKY achieves an (r, ε) -approximation for every psd matrix after accessing only a modest multiple of r/ε columns.

[Table 1](#) presents the best available upper bounds on the number k of columns for RPCHOLESKY and other column Nyström approximation methods to achieve [\(1.4\)](#). These bounds are expressed terms of r , ε , N , and the relative approximation error

$$\eta := \frac{\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)}{\text{tr} \mathbf{A}} \in [0, 1].$$

Method	Number k of columns	Reference
Greedy method	$(1 - (1 + \varepsilon)\eta)N$	Theorem 4.1
Uniform sampling*	$\frac{r-1}{\varepsilon\eta} + \frac{1}{\varepsilon}$	Theorem 4.2
DPP sampling	$\frac{r}{\varepsilon} + r - 1$	Theorem 4.3
RLS sampling**	$136 \frac{r}{\varepsilon} \log\left(\frac{3r}{\varepsilon\eta}\right)$	Corollary 4.5
RPCHOLESKY	$\frac{r}{\varepsilon} + r \log_+\left(\frac{1}{\varepsilon\eta}\right)$	Theorem 3.1
RPCHOLESKY	$\frac{r}{\varepsilon} + r + r \log_+\left(\frac{2^r}{\varepsilon}\right)$	Theorem 3.1

Table 1: Upper bounds on the number of columns for column Nyström approximation methods to produce an (r, ε) -approximation [\(1.4\)](#). *The result for uniform sampling assumes the diagonal entries of \mathbf{A} are set to one. **The result for RLS sampling assumes that $\varepsilon < \frac{1}{4}$.

DPP sampling has the strongest error bounds of all the methods in [Table 1](#). To achieve an (r, ε) -

approximation with DPP sampling, it is sufficient to take the number k of columns as

$$(1.5) \quad k \geq \frac{r}{\varepsilon} + r - 1.$$

RPCHOLESKY has the second strongest error bounds, requiring no more columns than

$$(1.6) \quad k \geq \frac{r}{\varepsilon} + r \log_+ \left(\frac{1}{\varepsilon \eta} \right),$$

where $\log_+(x) := \max\{\log x, 0\}$ for $x > 0$. The main difference between the DPP result (1.5) and the RPCHOLESKY result (1.6) is the multiplicative factor $\log(1/\eta)$ present in the latter. However, because the relative error η appears inside the logarithm, this factor has only a modest impact on the computational scaling. Indeed, $\log(1/\eta) < 37$ when $\eta > 10^{-16}$, the machine precision in double-precision arithmetic. See section 3 for additional RPCHOLESKY error bounds, which consider the limit as $\eta \rightarrow 0$.

The error bounds for RPCHOLESKY are *significantly stronger* than the bounds for the greedy method and uniform sampling. For a worst-case matrix \mathbf{A} , the greedy method requires $\Theta(N)$ columns to approach the best rank- r approximation error, and the uniform sampling method requires $\Theta(r/\eta)$ columns (subsections 4.1 and 4.2). In contrast, RPCHOLESKY requires a number of columns that is independent of the dimension N and depends only logarithmically on the relative error η . These results rigorously justify why RPCHOLESKY does not exhibit the same failure modes as the greedy method and the uniform sampling method.

Our theoretical work confirms that RPCHOLESKY achieves a given approximation error with roughly the same number of columns as the best theoretical approach for column Nyström approximation. At the same time, RPCHOLESKY uses fewer entry evaluations than other more complicated algorithms to attain this level of error. Given the simplicity, efficiency, and reliability of RPCHOLESKY, we believe that there is a strong case for deploying this algorithm in applications.

1.4. Plan for paper. The rest of the paper is organized as follows. Section 2 describes a fast implementation of RPCHOLESKY and presents numerical demonstrations of RPCHOLESKY in machine learning tasks. Section 3 presents the theoretical analysis of RPCHOLESKY, while section 4 discusses the performance of alternative methods. Section 5 surveys future directions.

1.5. Notation. The symbol \preceq denotes the *psd order* on symmetric matrices: $\mathbf{H} \preceq \mathbf{A}$ if and only if $\mathbf{A} - \mathbf{H}$ is psd. Submatrices are expressed using MATLAB notation, for example, $\mathbf{A}(:, i)$ represents the i th column of \mathbf{A} and $\mathbf{A}(S, :)$ denotes the submatrix of \mathbf{A} with rows indexed by the set S . RPCHOLESKY and its analysis extend to *complex-valued* psd matrices without significant modification, but for simplicity this work focuses on the real setting. The transpose of \mathbf{A} is denoted \mathbf{A}^* , which becomes the *conjugate transpose* for complex matrices.

2. Algorithms and experiments. This section gives efficient pseudocode for RPCHOLESKY (subsection 2.1) and presents a range of numerical examples (subsections 2.2 and 2.3).

2.1. Speeding up RPCholesky. To accelerate the naive implementation of RPCHOLESKY in Algorithm 1.1, we can apply the following techniques:

- The Nyström approximation $\hat{\mathbf{A}}$ can be stored and updated in factored form $\hat{\mathbf{A}} = \mathbf{F}\mathbf{F}^*$ where $\mathbf{F} \in \mathbb{R}^{N \times k}$ is a thin matrix.

- Rather than forming and maintaining the residual matrix $\mathbf{A}^{(i)}$, it suffices to generate the single column $\mathbf{A}^{(i)}(:, s_i)$ at every iteration $i = 1, 2, \dots, k$.

The resulting RPCHOLESKY algorithm requires $(k + 1)N$ entry evaluations and $2k^2N/3$ arithmetic operations; see [Algorithm 2.1](#).

Algorithm 2.1 RPCHOLESKY: Efficient implementation

Input: Psd matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and iteration count k or tolerance η

Output: Factor matrix $\mathbf{F} \in \mathbb{R}^{N \times k}$ defining Nyström approximation $\hat{\mathbf{A}} = \mathbf{F}\mathbf{F}^*$ and selected pivots S

Initialize $\mathbf{F} \leftarrow \mathbf{0}_{N \times k}$, $S \leftarrow \emptyset$, and $\mathbf{d} \leftarrow \text{diag } \mathbf{A}$

for $i = 1$ to k **do**

▷ Alternately, run until $\sum_{j=1}^N d_j \leq \eta \text{tr } \mathbf{A}$

 Sample $s \sim \mathbf{d} / \sum_{j=1}^N d_j$

$S \leftarrow S \cup \{s\}$

$\mathbf{g} \leftarrow \mathbf{A}(:, s) - \mathbf{F}(:, 1:i-1)\mathbf{F}(s, 1:i-1)^*$

$\mathbf{F}(:, i) \leftarrow \mathbf{g} / \sqrt{\mathbf{g}_s}$

$\mathbf{d} \leftarrow \mathbf{d} - |\mathbf{F}(:, i)|^2$

▷ $|\mathbf{F}(:, i)|^2$ is the vector with j th entry $|\mathbf{F}(j, i)|^2$

end for

To provide an even greater speedup in some settings, RPCHOLESKY can be modified to sample multiple columns in a single iteration, as in [Algorithm 2.2](#). In experiments, this *blocked* version of RPCHOLESKY produces similar quality approximations to standard RPCHOLESKY, and it reduces the runtime when generating multiple columns of \mathbf{A} is faster than generating columns one at a time. For a discussion of these runtime considerations, see [Appendix B](#).

Algorithm 2.2 RPCHOLESKY: Blocked variant

Input: Psd matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, iteration count k , and block size B

Output: Factor matrix $\mathbf{F} \in \mathbb{R}^{N \times k}$ defining Nyström approximation $\hat{\mathbf{A}} = \mathbf{F}\mathbf{F}^*$ and selected pivots S

Initialize $\mathbf{F} \leftarrow \mathbf{0}_{N \times k}$, $S \leftarrow \emptyset$, $\mathbf{d} \leftarrow \text{diag } \mathbf{A}$, and $i \leftarrow 0$

while $i < k$ **do**

 Sample $s_1, s_2, \dots, s_{\min(B, k-i)} \stackrel{\text{iid}}{\sim} \mathbf{d} / \sum_{j=1}^N d_j$

$U \leftarrow \text{UNIQUE}(s_1, s_2, \dots, s_{\min(B, k-i)})$

▷ Identify unique elements

$S \leftarrow S \cup U$

$\mathbf{G} \leftarrow \mathbf{A}(:, U) - \mathbf{F}(:, 1:i)[\mathbf{F}(U, 1:i)]^*$

$\mathbf{R} \leftarrow \text{CHOL}(\mathbf{G}(U, :))$

▷ $\mathbf{G}(U, :) = \mathbf{R}^* \mathbf{R}$

$\mathbf{F}(:, i:i+|U|) \leftarrow \mathbf{G}\mathbf{R}^{-1}$

$\mathbf{d} \leftarrow \mathbf{d} - \text{SQUAREDROWNORMS}(\mathbf{F}(:, i:i+|U|))$

$i \leftarrow i + |U|$

end while

2.2. Kernel ridge regression. One powerful application for RPCHOLESKY is to accelerate *kernel ridge regression* (KRR) [27, §4.9.1]. KRR is a nonlinear extension of least-squares regression that approximates an unknown input–output map using a positive definite kernel function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ and a set of input–output pairs $(\mathbf{x}^{(1)}, y_1), \dots, (\mathbf{x}^{(N)}, y_N)$. As output, KRR provides a prediction

function of the form

$$(2.1) \quad f(\cdot; \boldsymbol{\beta}) := \sum_{i=1}^N \beta_i K(\mathbf{x}^{(i)}, \cdot),$$

with the coefficient vector $\boldsymbol{\beta}$ chosen to minimize a regularized least-squares loss:

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^N} \frac{1}{N} \sum_{j=1}^N |f(\mathbf{x}^{(j)}; \boldsymbol{\beta}) - y_j|^2 + \lambda \sum_{i,j=1}^N \beta_i \beta_j K(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}).$$

Explicitly, the vector $\boldsymbol{\beta}$ is the solution to a linear system

$$(2.2) \quad \boldsymbol{\beta} = (\mathbf{A} + \lambda \mathbf{N} \mathbf{I})^{-1} \mathbf{y},$$

where \mathbf{A} is the $N \times N$ kernel matrix and \mathbf{y} is the vector of output values.

Directly computing the $\boldsymbol{\beta}$ coefficients by (2.2) would require solving a dense linear system at $\mathcal{O}(N^3)$ cost, but RPCHOLESKY provides a faster alternative. First, we apply RPCHOLESKY to the kernel matrix \mathbf{A} to generate a random set of pivots $S = \{s_1, \dots, s_k\}$. Next, we identify a restricted prediction function

$$(2.3) \quad \hat{f}(\cdot; \hat{\boldsymbol{\beta}}) = \sum_{i=1}^k \hat{\beta}_i K(\mathbf{x}^{(s_i)}, \cdot)$$

that minimizes the loss

$$\min_{\hat{\boldsymbol{\beta}} \in \mathbb{R}^k} \frac{1}{N} \sum_{j=1}^N |\hat{f}(\mathbf{x}^{(j)}; \hat{\boldsymbol{\beta}}) - y_j|^2 + \lambda \sum_{i,j=1}^k \hat{\beta}_i \hat{\beta}_j K(\mathbf{x}^{(s_i)}, \mathbf{x}^{(s_j)}).$$

The coefficient vector $\hat{\boldsymbol{\beta}} \in \mathbb{R}^k$ satisfies a smaller linear system

$$(2.4) \quad \hat{\boldsymbol{\beta}} = (\mathbf{A}(S, :) \mathbf{A}(:, S) + \lambda N \mathbf{A}(S, S))^{-1} \mathbf{A}(S, :) \mathbf{y}.$$

involving a $k \times k$ matrix, which makes it relatively inexpensive to solve.

Algorithm 2.3 provides pseudocode for RPCHOLESKY-accelerated KRR. The most expensive steps are the RPCHOLESKY step and evaluation of the matrix–matrix product $\mathbf{A}(S, :) \mathbf{A}(:, S)$; each requires $\mathcal{O}(k^2 N)$ operations. After performing these steps, the coefficient vector $\hat{\boldsymbol{\beta}}$ can be computed in $\mathcal{O}(k^3)$ operations via a standard dense linear solver. Evaluating the prediction function (2.3) from RPCHOLESKY-accelerated KRR requires just $\mathcal{O}(k)$ operations, which improves over the $\mathcal{O}(N)$ cost of evaluating (2.1).

To demonstrate the effectiveness of RPCHOLESKY-accelerated KRR, we use **Algorithm 2.3** to predict the energy of the highest occupied molecular orbital (HOMO) for small organic molecules from the QM9 data set [24, 25]. The HOMO energy quantifies the electron-donating capacity of a molecule, and it is typically obtained by expensive first-principles calculations. Estimating HOMO energies and related properties from data has been a major goal of chemical machine learning [28].

To represent molecules as vectors for KRR, we apply the Coulomb matrix representation with ℓ_2 row norm sorting [28, §III.A], a standard approach based on Coulomb repulsions between atomic

Algorithm 2.3 RPCHOLESKY-accelerated kernel ridge regression

Input: Data points $\mathbf{X} = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\} \subseteq \mathbb{R}^d$; output values $\mathbf{y} \in \mathbb{R}^N$; approximation rank k ; regularization parameter $\lambda > 0$

Output: Pivots \mathbf{S} and coefficients $\hat{\boldsymbol{\beta}}$ defining a prediction function $\hat{f}(\cdot)$ by (2.3)

$$\begin{aligned} \mathbf{A} &\leftarrow \text{KERNELMATRIX}(\mathbf{X}) \\ (\sim, \mathbf{S}) &\leftarrow \text{RPCHOLESKY}(\mathbf{A}, k) \\ \hat{\boldsymbol{\beta}} &\leftarrow \left(\mathbf{A}(\mathbf{S}, :) \mathbf{A}(:, \mathbf{S}) + \lambda N \mathbf{A}(\mathbf{S}, \mathbf{S}) \right)^{-1} \mathbf{A}(\mathbf{S}, :) \mathbf{y} \end{aligned}$$

nuclei. We standardize the data and, following [28], evaluate the similarity between data points using the ℓ_1 Laplace kernel:

$$(2.5) \quad K(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{1}{\sigma} \sum_{j=1}^d |x_j - y_j|\right).$$

We divide the data into 100,000 data points for training and roughly 33,000 data points for testing and set the bandwidth σ and ridge parameter λ using cross-validation. We measure out-of-sample prediction errors using the symmetric mean absolute percentage error (SMAPE)

$$(2.6) \quad \text{SMAPE} = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \frac{|y_i^{\text{test}} - \hat{f}(\mathbf{x}^{(i), \text{test}}; \hat{\boldsymbol{\beta}})|}{(|y_i^{\text{test}}| + |\hat{f}(\mathbf{x}^{(i), \text{test}}; \hat{\boldsymbol{\beta}})|)/2}.$$

Forming and storing the full kernel matrix for the QM9 data set would require 40 GB and 4 trillion arithmetic operations ($20,000 \times$ the operation count for the **Smile** matrix from subsection 1.2), making direct KRR highly expensive. Prior work [28] applied KRR to this data set using a random subsample of $N \leq 64,000$ training points. Here, RPCHOLESKY allows us to use take advantage of more of the data ($N = 100,000$ training points) at a reduced computational cost by using a small approximation rank $k \leq 1000 \ll N$.

Figure 3 displays the results. The prediction accuracies for RPCHOLESKY, uniform sampling, and RLS sampling are very close, with uniform sampling being the slight favorite. However, an examination shows that **uniform sampling has 18 \times higher prediction errors** than RPCHOLESKY for the molecules in the data set with the largest number of atoms (29); see Table 2. This observation suggests that RPCHOLESKY can be better than uniform sampling at representing less populated regions of data space, as seen earlier in the **Smile** example (subsection 1.2). The greedy method has notably worse performance than the three other methods, and we were unable to use DPP sampling because of the large values of $N = 10^5$ and $k = 10^3$.

2.3. Spectral clustering. We can also use RPCHOLESKY to accelerate *kernel spectral clustering* [33]. Kernel spectral clustering is a machine learning approach that uses a nonnegative, positive definite kernel function $K: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ to identify features in a data set $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}$ and applies clustering based on those features.

In a standard formulation of kernel spectral clustering [33], the features are defined by the dominant left eigenvectors $\mathbf{v}^{(1)}, \dots, \mathbf{v}^{(m)}$ of the matrix $\mathbf{D}^{-1} \mathbf{A}$. In this expression, $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the kernel matrix and $\mathbf{D} \in \mathbb{R}^{N \times N}$ is the diagonal matrix that lists the row sums of \mathbf{A} . The eigenvectors are




Figure 3: *Left*: Prediction error (2.6) for different Nyström algorithms. *Right*: Relative trace-norm error for different Nyström algorithms.

Table 2: Prediction error (2.6) for the largest molecules in the QM9 dataset

Compound #	Composition	Uniform	RPCHOLESKY
1,996	CC(C)CC1C01	0.519	0.030
8,664	CC(CO)(C=O)C=O	0.582	0.016
13,812	CC(O)C1CC1C0	0.523	0.033
64,333	CC1(C)CCOC(=N)O1	0.525	0.027
81,711	OC1C2CC1OCCO2	0.576	0.016
109,816	CCC1C(O)C1CC#C	0.514	0.023
118,229	COCC(C)OC(C)C	0.575	0.075
122,340	CC1C(CCCO)N1C	0.557	0.017
131,819	OCCCN1C=NC=N1	0.570	0.030

normalized such that $(\mathbf{v}^{(i)})^* \mathbf{D} \mathbf{v}^{(j)} = \delta_{ij}$, and each data point $\mathbf{x}^{(i)}$ is projected as

$$\tilde{\mathbf{x}}^{(i)} = (\mathbf{v}_i^{(1)}, \mathbf{v}_i^{(2)}, \dots, \mathbf{v}_i^{(m)}).$$

The projected data points $\tilde{\mathbf{x}}^{(1)}, \dots, \tilde{\mathbf{x}}^{(N)}$ are clustered using any standard method, such as k-means with m centers.

Directly computing the eigendecomposition of $\mathbf{D}^{-1} \mathbf{A}$ would require $\mathcal{O}(N^3)$ operations. However, RPCHOLESKY provides a cheaper alternative (Algorithm 2.4) that involves just $\mathcal{O}(k^2 N)$ operations. In the RPCHOLESKY-accelerated approach, a low-rank approximation $\hat{\mathbf{A}}^{(k)}$ is used in place of the kernel matrix \mathbf{A} , and the diagonal matrix $\hat{\mathbf{D}}$ containing the row sums of $\hat{\mathbf{A}}^{(k)}$ is used in place of \mathbf{D} . The data is projected onto the dominant normalized left eigenvectors of $\hat{\mathbf{D}}^{-1} \hat{\mathbf{A}}^{(k)}$ and the projected data points are clustered using k -means.

Algorithm 2.4 RPCHOLESKY-accelerated spectral clustering**Input:** Data points $X = \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}\} \subseteq \mathbb{R}^d$; number m of clusters; approximation rank k **Output:** Partition of X into m clusters $X = X_1 \cup \dots \cup X_m$

```

 $A \leftarrow \text{KERNELMATRIX}(X)$ 
 $F \leftarrow \text{RPCHOLESKY}(A, k)$ 
 $\hat{D} \leftarrow \text{diag}(F(F^* \mathbb{1}))$ 
 $G \leftarrow \hat{D}^{-1/2} F$ 
 $(Q, R) \leftarrow \text{QR}(G, \text{'econ'})$ 
 $(U, \sim, \sim) \leftarrow \text{SVD}(R)$ 
 $V \leftarrow \hat{D}^{-1/2} Q U$ 
 $\tilde{X} \leftarrow \{V(i, 1:m) : 1 \leq i \leq |X|\}$ 
 $X_1, \dots, X_m \leftarrow \text{K-MEANS}(\tilde{X}, m)$ 

```

In [Figure 4](#), we apply RPCHOLESKY-accelerated spectral clustering to separate the letters “RPC”, written in papyrus font. After reading 140 of the 2187 columns of the kernel matrix, RPCHOLESKY reliably produces a perfect clustering as shown in [Figure 4](#) (top left). In contrast, given the same approximation rank, RLS and uniform sampling frequently produce an incorrect clustering. The fraction of perfect clusterings over 100 independent trials is given at the bottom of [Figure 4](#). For this example, the greedy method outperforms all the randomized sampling methods, but we have already seen ([subsections 1.2](#) and [2.2](#)) that the greedy method can be unreliable and has serious failure modes.

3. Theoretical analysis of RPCholesky. Given the appealing computational profile of the RPCHOLESKY algorithm, we would like to understand when it is guaranteed to produce an accurate low-rank approximation. In this section, we establish the following new result:

Theorem 3.1 (Randomly pivoted Cholesky). *Fix $r \in \mathbb{N}$ and $\delta \in (0, 1)$, and let A be a psd matrix. The column Nyström approximation $\hat{A}^{(k)}$ produced by RPCHOLESKY ([Algorithm 2.1](#)) attains the bound*

$$(3.1) \quad \mathbb{E} \text{tr}(A - \hat{A}^{(k)}) \leq \text{tr}(A - \llbracket A \rrbracket_r) + \delta \text{tr} A,$$

provided that the number k of columns satisfies

$$(3.2) \quad k \geq \frac{r\eta}{\delta} + \min \left\{ r \log \left(\frac{1}{\delta} \right), r + r \log_+ \left(\frac{2^r \eta}{\delta} \right) \right\}.$$

The relative error η is defined by $\eta := \text{tr}(A - \llbracket A \rrbracket_r) / \text{tr}(A)$. As usual, $\log_+(x) := \max\{\log x, 0\}$ for $x > 0$.

The major takeaway from [Theorem 3.1](#) is that the number k of columns depends *logarithmically* on the tolerance $1/\delta$, as long as $\delta \geq \eta$. The error bounds for RPCHOLESKY become more powerful as $\eta \rightarrow 0$ and RPCHOLESKY achieves the theoretically optimal error

$$\mathbb{E} \text{tr}(A - \hat{A}^{(k)}) = \text{tr}(A - \llbracket A \rrbracket_r)$$

with just $k = r$ columns when $\eta = 0$.




Figure 4: *Top*: A correct clustering produced by RPCHOLESKY along with selected incorrect clusterings produced by RLS, DPP, and uniform sampling with rank $k = 140$. *Bottom*: Fraction of correct clusterings produced by Nyström algorithms over 100 independent trials.

By setting $\delta = \varepsilon\eta$ for $\varepsilon > 0$, [Theorem 3.1](#) guarantees that RPCHOLESKY must produce an (r, ε) -approximation as soon as the number k of columns satisfies

$$k \geq \frac{r}{\varepsilon} + r \log_+ \left(\frac{1}{\varepsilon\eta} \right).$$

In practice, the logarithmic factor is typically a modest constant, so the bound is comparable with the minimal cost of $k \geq \frac{r}{\varepsilon}$ columns. Thus, RPCHOLESKY improves on the greedy method, in which the number k of columns may be proportional to the dimension N in the worst case. It also improves on the uniform sampling method, in which the number k of columns can depend *polynomially* on the relative error $1/\eta$. See [subsections 4.1](#) and [4.2](#) for details.

Additionally, [Theorem 3.1](#) demonstrates that RPCHOLESKY produces an (r, ε) -approximation

when the number k of columns satisfies

$$(3.3) \quad k \geq \frac{r}{\varepsilon} + r + r \log_+ \left(\frac{2^r}{\varepsilon} \right).$$

This alternative error bound is significant because it completely eliminates the dependence on the relative error η . When r is small, (3.3) is competitive with the cost of $k \geq \frac{r}{\varepsilon} + r - 1$ columns for DPP sampling. However, when r is large, (3.3) deteriorates due to the quadratic dependence on r .

Although some error bounds for RPCHOLESKY can be deduced from earlier work [8, Thms. 4–5 & Prop. 1], [Theorem 3.1](#) gives bounds that are strictly sharper.

3.1. Proof of [Theorem 3.1](#). The proof of [Theorem 3.1](#) is based on the properties of the *expected residual function*:

$$(3.4) \quad \Phi(\mathbf{A}) := \mathbb{E}[\mathbf{A}^{(1)} \mid \mathbf{A}] = \mathbf{A} - \frac{\mathbf{A}^2}{\text{tr } \mathbf{A}}.$$

This function returns the expectation of the residual $\mathbf{A}^{(1)}$ after applying one step of RPCHOLESKY to \mathbf{A} . The equality (3.4) follows from a short computation using the sampling distribution (1.3) and the definition of the residual ([Algorithm 1.1](#)). Note that Φ is defined on psd matrices of any dimension.

The first lemma describes some basic facts about the expected residual function Φ . We postpone the proof to [subsection 3.2](#).

Lemma 3.2 (Expected residual). *The expected residual map Φ defined in (3.4) is positive, monotone, and concave with respect to the psd order. That is, for all psd \mathbf{A}, \mathbf{H} with the same dimensions,*

$$(3.5) \quad \mathbf{0} \leq \Phi(\mathbf{A}) \leq \Phi(\mathbf{A} + \mathbf{H});$$

$$(3.6) \quad \theta \Phi(\mathbf{A}) + (1 - \theta) \Phi(\mathbf{H}) \leq \Phi(\theta \mathbf{A} + (1 - \theta) \mathbf{H}) \quad \text{for all } \theta \in [0, 1].$$

The second lemma describes how the trace of the residual declines after multiple steps of the RPCHOLESKY procedure. The proof appears in [subsection 3.3](#).

Lemma 3.3 (Contraction rate). *Consider the k -fold composition of the expected residual (3.4):*

$$\Phi^{\circ k} := \underbrace{\Phi \circ \Phi \circ \dots \circ \Phi}_{k \text{ times}} \quad \text{for each } k \in \mathbb{N}.$$

Fix $r \in \mathbb{N}$. For each psd matrix \mathbf{H} and each $\Delta > 0$,

$$\text{tr } \Phi^{\circ k}(\mathbf{H}) \leq \text{tr}(\mathbf{H} - \llbracket \mathbf{H} \rrbracket_r) + \Delta \text{tr } \mathbf{H} \quad \text{when } k \geq \frac{r \text{tr}(\mathbf{H} - \llbracket \mathbf{H} \rrbracket_r)}{\Delta \text{tr } \mathbf{H}} + r \log_+ \left(\frac{1}{\Delta} \right).$$

Last, we present a bound which shows that the error after k steps of RPCHOLESKY is comparable with the error in the best rank- k approximation. The proof appears in [subsection 3.4](#).

Lemma 3.4 (Error doubling). *For each psd matrix \mathbf{A} , the residual matrix $\mathbf{A}^{(k)}$ after applying k steps of RPCHOLESKY satisfies*

$$\mathbb{E} \text{tr } \mathbf{A}^{(k)} \leq 2^k \text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_k) \quad \text{for each } k \in \mathbb{N}.$$

With these results at hand, we quickly establish the main error bound for `RPCHOLESKY`.

Proof of Theorem 3.1. Let \mathbf{A} be a psd matrix. By the concavity (3.6) of the expected residual map (3.4) and the matrix version of Jensen's inequality [5, Thm. 4.16], the residual matrices satisfy

$$\mathbb{E} \mathbf{A}^{(j)} = \mathbb{E} \left[\mathbb{E} \left[\mathbf{A}^{(j)} \mid \mathbf{A}^{(j-1)} \right] \right] = \mathbb{E} \Phi(\mathbf{A}^{(j-1)}) \leq \Phi(\mathbb{E} \mathbf{A}^{(j-1)}) \quad \text{for each } j \in \{1, \dots, k\}.$$

Next, by monotonicity (3.5) of Φ and the last display,

$$\mathbb{E} \mathbf{A}^{(k)} \leq \Phi(\mathbb{E} \mathbf{A}^{(k-1)}) \leq \Phi \circ \Phi(\mathbb{E} \mathbf{A}^{(k-2)}) \leq \dots \leq \Phi^{\circ k}(\mathbf{A}).$$

Using the fact that the trace is linear and preserves the psd order,

$$\mathbb{E} \operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)}) = \mathbb{E} \operatorname{tr} \mathbf{A}^{(k)} \leq \operatorname{tr} \Phi^{\circ k}(\mathbf{A}).$$

Next, we apply Lemma 3.3 with $\mathbf{H} = \mathbf{A}$ and $\Delta = \delta$ to see that

$$\mathbb{E} \operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)}) \leq \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r) + \delta \operatorname{tr} \mathbf{A} \quad \text{when } k \geq \frac{r\eta}{\delta} + r \log\left(\frac{1}{\delta}\right),$$

where we have recognized the relative error η .

By a similar argument,

$$\mathbb{E} \operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)}) = \mathbb{E} \operatorname{tr} \mathbf{A}^{(k)} \leq \operatorname{tr} \Phi^{\circ(k-r)}(\mathbb{E} \mathbf{A}^{(r)}).$$

We can apply Lemma 3.3 with $\mathbf{H} = \mathbb{E} \mathbf{A}^{(r)}$ and $\Delta = \delta \cdot \operatorname{tr} \mathbf{A} / \mathbb{E} \operatorname{tr} \mathbf{A}^{(r)}$ to see that

$$(3.7) \quad \mathbb{E} \operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)}) \leq \operatorname{tr}(\mathbb{E} \mathbf{A}^{(r)} - \llbracket \mathbb{E} \mathbf{A}^{(r)} \rrbracket_r) + \delta \operatorname{tr} \mathbf{A}$$

provided the number of columns satisfies

$$(3.8) \quad k - r \geq \frac{r \operatorname{tr}(\mathbb{E} \mathbf{A}^{(r)} - \llbracket \mathbb{E} \mathbf{A}^{(r)} \rrbracket_r)}{\delta \operatorname{tr} \mathbf{A}} + r \log\left(\frac{\mathbb{E} \operatorname{tr} \mathbf{A}^{(r)}}{\delta \operatorname{tr} \mathbf{A}}\right).$$

This bound can be simplified as follows. Observe that the (random) residual matrix $\mathbf{A}^{(r)}$ is a Schur complement of \mathbf{A} , so it satisfies $\mathbb{E} \mathbf{A}^{(r)} \leq \mathbf{A}$. By the Ky Fan variational principle [36, Thm. 8.17], the best rank- r approximation error in the trace norm is monotone with respect to the psd order, so

$$\operatorname{tr}(\mathbb{E} \mathbf{A}^{(r)} - \llbracket \mathbb{E} \mathbf{A}^{(r)} \rrbracket_r) \leq \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r).$$

Additionally, Lemma 3.4 guarantees that $\mathbb{E} \operatorname{tr} \mathbf{A}^{(r)} \leq 2^r \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)$. Using these facts, we can simplify (3.7)–(3.8) to show that

$$\mathbb{E} \operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)}) \leq \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r) + \delta \operatorname{tr} \mathbf{A} \quad \text{when } k \geq r + \frac{r\eta}{\delta} + r \log\left(\frac{2^r \eta}{\delta}\right),$$

This completes the proof of the second bound. ■

3.2. Proof of Lemma 3.2. Let A, H be psd matrices, and recall the definition (3.4) of the expected residual map Φ . First, to prove that Φ is positive, note that

$$(3.9) \quad \Phi(H) = \left(\mathbf{I} - \frac{H}{\text{tr } H} \right) H \geq \mathbf{0}.$$

Next, to establish concavity, we choose $\theta \in [0, 1]$ and make the calculation

$$\Phi(\theta A + (1 - \theta)H) - \theta\Phi(A) - (1 - \theta)\Phi(H) = \frac{\theta(1 - \theta)}{\theta \text{tr } A + (1 - \theta) \text{tr } H} \left(\sqrt{\frac{\text{tr } H}{\text{tr } A}} A - \sqrt{\frac{\text{tr } A}{\text{tr } H}} H \right)^2 \geq \mathbf{0}.$$

Last, to establish monotonicity, observe that Φ is positive homogeneous; that is, $\Phi(\tau A) = \tau\Phi(A)$ for $\tau \geq 0$. Invoking the concavity property (3.6) with $\theta = 1/2$,

$$\Phi(A + H) = 2\Phi\left(\frac{A + H}{2}\right) \geq \Phi(A) + \Phi(H) \geq \Phi(A).$$

We have used the positivity of $\Phi(H)$ in the last step. ■

3.3. Proof of Lemma 3.3. First, we show that it suffices to consider the case of a diagonal matrix. Let H be an $N \times N$ psd matrix with eigendecomposition $H = V\Lambda V^*$. The definition (3.4) of the expected residual map implies that

$$\Phi(H) = V\Phi(\Lambda)V^*.$$

By iteration, the same relation holds with $\Phi^{\circ k}$ in place of Φ . In particular, $\text{tr } \Phi^{\circ k}(H) = \text{tr } \Phi^{\circ k}(\Lambda)$. Therefore, we may restrict our attention to the diagonal case where $H = \Lambda$.

Second, we obtain an upper bound on $\text{tr } \Phi^{\circ k}(\Lambda)$. Because the map $\Lambda \mapsto \text{tr } \Phi^{\circ k}(\Lambda)$ is concave and invariant under permutations of its arguments, averaging together some of the eigenvalues $\lambda_{11}, \dots, \lambda_{NN}$ can only increase the value of $\text{tr } \Phi^{\circ k}(\Lambda)$. Therefore, by introducing the function

$$f_k(a, b, r, q) := \text{tr } \Phi^{\circ k} \left(\text{diag} \left(\underbrace{\frac{a}{r}, \dots, \frac{a}{r}}_{r \text{ times}}, \underbrace{\frac{b}{q}, \dots, \frac{b}{q}}_{q \text{ times}} \right) \right),$$

we obtain the upper bound

$$(3.10) \quad \text{tr } \Phi^{\circ k}(\Lambda) \leq f_k(\text{tr}[\Lambda]_r, \text{tr}(\Lambda - [\Lambda]_r), r, N - r).$$

For further reference, we note that $f_k(a, b, r, q)$ is weakly increasing as a function of q . Indeed, for every tuple (a, b, r, q) ,

$$\begin{aligned} f_k(a, b, r, q) &= \text{tr } \Phi^{\circ k} \left(\text{diag} \left(\underbrace{\frac{a}{r}, \dots, \frac{a}{r}}_{r \text{ times}}, \underbrace{\frac{b}{q}, \dots, \frac{b}{q}}_{q \text{ times}} \right) \right) = \text{tr } \Phi^{\circ k} \left(\text{diag} \left(\underbrace{\frac{a}{r}, \dots, \frac{a}{r}}_{r \text{ times}}, \underbrace{\frac{b}{q}, \dots, \frac{b}{q}}_{q \text{ times}}, 0 \right) \right) \\ &\leq \text{tr } \Phi^{\circ k} \left(\text{diag} \left(\underbrace{\frac{a}{r}, \dots, \frac{a}{r}}_{r \text{ times}}, \underbrace{\frac{b}{q+1}, \dots, \frac{b}{q+1}}_{q+1 \text{ times}} \right) \right) = f_k(a, b, r, q+1). \end{aligned}$$

We have exploited the fact that Φ is defined for matrices of every dimension.

Next, we derive a worst-case expression for the error $f_k(a, b, r, q)$. For each $k = 0, 1, 2, \dots$, define the nonnegative quantities $a^{(k)}$ and $b^{(k)}$ via the relation

$$\Phi^{\circ k} \left(\text{diag} \left(\underbrace{\frac{a}{r}, \dots, \frac{a}{r}}_{r \text{ times}}, \underbrace{\frac{b}{q}, \dots, \frac{b}{q}}_{q \text{ times}} \right) \right) =: \text{diag} \left(\underbrace{\frac{a^{(k)}}{r}, \dots, \frac{a^{(k)}}{r}}_{r \text{ times}}, \underbrace{\frac{b^{(k)}}{q}, \dots, \frac{b^{(k)}}{q}}_{q \text{ times}} \right).$$

By the definition (3.4) of the expected residual map Φ , the quantities $a^{(k)}$ and $b^{(k)}$ satisfy the recurrence relations

$$a^{(k)} - a^{(k-1)} = -\frac{(a^{(k-1)})^2}{r(a^{(k-1)} + b^{(k-1)})} \quad \text{and} \quad b^{(k)} - b^{(k-1)} = -\frac{(b^{(k-1)})^2}{q(a^{(k-1)} + b^{(k-1)})}$$

with initial conditions $a^{(0)} = a$ and $b^{(0)} = b$. This construction guarantees $f_k(a, b, r, q) = a^{(k)} + b^{(k)}$. Additionally, the quantities $a^{(k)}$ and $b^{(k)}$ depend continuously on q , and they converge as $q \rightarrow \infty$ to limiting values $\bar{a}^{(k)}$ and $\bar{b}^{(k)} \equiv b$, where the sequence $\bar{a}^{(k)}$ satisfies

$$\bar{a}^{(k)} - \bar{a}^{(k-1)} = -\frac{(\bar{a}^{(k-1)})^2}{r(\bar{a}^{(k-1)} + b)} \quad \text{with initial condition } \bar{a}^{(0)} = a.$$

It follows that

$$f_k(a, b, r, q) \leq \bar{a}^{(k)} + b \leq a + b.$$

We have used the facts that $f_k(a, b, r, q)$ is increasing in q and $\bar{a}^{(k)}$ is decreasing in k .

All that remains is to determine how quickly $\bar{a}^{(k)}$ decreases as a function of k . To that end, we pass from discrete time to continuous time. At each instant $t = 0, 1, 2, \dots$, the discrete-time process $\bar{a}^{(t)}$ is bounded from above by the continuous-time process $x(t)$ satisfying the ODE

$$\frac{d}{dt} x(t) = -\frac{x(t)^2}{r(x(t) + b)} \quad \text{with initial condition } x(0) = a.$$

The comparison between discrete- and continuous-time processes holds because $x \mapsto -x^2/(rx + rb)$ is decreasing over the range $x \in (0, \infty)$. Next, assuming $\Delta(a + b) \leq a$, we can use separation of variables to solve for the time t_\star at which $x(t_\star) = \Delta(a + b)$:

$$t_\star = \int_{x=a}^{\Delta(a+b)} -\frac{r(x+b)}{x^2} dx = rb \left(\frac{1}{\Delta(a+b)} - \frac{1}{a} \right) + r \log \left(\frac{a}{\Delta(a+b)} \right) \leq \frac{rb}{\Delta(a+b)} + r \log \left(\frac{1}{\Delta} \right).$$

It follows (even for $\Delta(a + b) > a$) that

$$\bar{a}^{(k)} \leq \Delta(a + b) \quad \text{when} \quad k \geq \frac{rb}{\Delta(a+b)} + r \log_+ \left(\frac{1}{\Delta} \right).$$

To conclude, substitute $a = \text{tr}[\mathbf{\Lambda}]_r$ and $b = \text{tr}(\mathbf{\Lambda} - [\mathbf{\Lambda}]_r)$ and combine with (3.10). ■

3.4. Proof of Lemma 3.4. Let \mathbf{P} denote the orthogonal projection onto the k dominant eigenvectors of \mathbf{A} and set $\mathbf{P}_\perp := \mathbf{I} - \mathbf{P}$. Consider a hierarchical sampling procedure for the pivot index s_1 :

- With probability $\text{tr}(\mathbf{A}\mathbf{P})/\text{tr}\mathbf{A}$, the “good” event \mathbf{G} occurs and the pivot index is chosen as

$$\mathbb{P}(s_1 = i \mid \mathbf{G}) = \frac{(\mathbf{A}\mathbf{P})_{ii}}{\text{tr}(\mathbf{A}\mathbf{P})}.$$

- With probability $\text{tr}(\mathbf{A}\mathbf{P}_\perp)/\text{tr}\mathbf{A}$, the “bad” event $\mathbf{B} = \mathbf{G}^c$ occurs and the pivot is selected as

$$\mathbb{P}(s_1 = i \mid \mathbf{B}) = \frac{(\mathbf{A}\mathbf{P}_\perp)_{ii}}{\text{tr}(\mathbf{A}\mathbf{P}_\perp)}.$$

It suffices to analyze this two-step sampling procedure. Indeed, a short calculation confirms that this procedure leads to the same sampling distribution for s_1 as in RPCHOLESKY.

On the bad event \mathbf{B} , apply the crude bound

$$(3.11) \quad \text{tr}(\mathbf{A}^{(1)} - \llbracket \mathbf{A}^{(1)} \rrbracket_{k-1}) = \sum_{j=k}^N \lambda_j(\mathbf{A}^{(1)}) \leq \lambda_k(\mathbf{A}) + \text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_k).$$

This relation holds because $\mathbf{A}^{(1)} \leq \mathbf{A}$, and the Weyl monotonicity principle [36, Thm. 8.11] yields $\lambda_j(\mathbf{A}_1) \leq \lambda_j(\mathbf{A})$ for $j = k, \dots, N$.

On the good event \mathbf{G} , apply the Ky Fan variational principle [36, Thm. 8.17] to write

$$\text{tr}(\mathbf{A}^{(1)} - \llbracket \mathbf{A}^{(1)} \rrbracket_{k-1}) = \min \left\{ \sum_{j=k}^N \mathbf{u}_j^* \mathbf{A}^{(1)} \mathbf{u}_j : \mathbf{u}_k, \dots, \mathbf{u}_N \text{ orthonormal} \right\}.$$

Choose the vectors $\mathbf{u}_{k+1}, \dots, \mathbf{u}_N$ to be unit-norm eigenvectors $\mathbf{v}_{k+1}(\mathbf{A}), \dots, \mathbf{v}_N(\mathbf{A})$ associated with the smallest eigenvalues. Set $\mathbf{u}_k := \mathbf{P}\mathbf{e}_{s_1} / \|\mathbf{P}\mathbf{e}_{s_1}\|$, which is orthonormal to the other vectors by the choice of \mathbf{P} . This gives the bound

$$(3.12) \quad \text{tr}(\mathbf{A}^{(1)} - \llbracket \mathbf{A}^{(1)} \rrbracket_{k-1}) \leq \sum_{j=k+1}^N \mathbf{v}_j(\mathbf{A})^* \mathbf{A}^{(1)} \mathbf{v}_j(\mathbf{A}) + \frac{\mathbf{e}_{s_1}^* \mathbf{P} \mathbf{A}^{(1)} \mathbf{P} \mathbf{e}_{s_1}}{\mathbf{e}_{s_1}^* \mathbf{P} \mathbf{e}_{s_1}} \leq \text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_k) + \frac{(\mathbf{P} \mathbf{A}^{(1)} \mathbf{P})_{s_1 s_1}}{p_{s_1 s_1}}.$$

Using the fact that $\mathbf{A}^{(1)} = \mathbf{A} - \mathbf{A} \mathbf{e}_{s_1} \mathbf{e}_{s_1}^* \mathbf{A} / a_{s_1 s_1}$, it follows

$$(\mathbf{P} \mathbf{A}^{(1)} \mathbf{P})_{s_1 s_1} = (\mathbf{A}\mathbf{P})_{s_1 s_1} - \frac{((\mathbf{A}\mathbf{P})_{s_1 s_1})^2}{a_{s_1 s_1}} = \frac{(\mathbf{A}\mathbf{P})_{s_1 s_1} (\mathbf{A}\mathbf{P}_\perp)_{s_1 s_1}}{a_{s_1 s_1}}.$$

Therefore, taking the conditional expectation of both sides of (3.12),

$$(3.13) \quad \begin{aligned} \mathbb{E}[\text{tr}(\mathbf{A}^{(1)} - \llbracket \mathbf{A}^{(1)} \rrbracket_{k-1}) \mid \mathbf{G}] &\leq \text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_k) + \sum_{i=1}^N \frac{(\mathbf{A}\mathbf{P})_{ii}}{\text{tr}(\mathbf{A}\mathbf{P})} \cdot \frac{(\mathbf{A}\mathbf{P})_{ii} (\mathbf{A}\mathbf{P}_\perp)_{ii}}{a_{ii} p_{ii}} \\ &\leq \left(1 + \frac{\lambda_1(\mathbf{A})}{\text{tr}(\mathbf{A}\mathbf{P})} \right) \text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_k). \end{aligned}$$

The last line uses the inequalities $(\mathbf{A}\mathbf{P})_{ii} \leq a_{ii}$ and $(\mathbf{A}\mathbf{P})_{ii} \leq \lambda_1(\mathbf{A}) p_{ii}$.

Finally, combine the bound (3.11) for the bad event and the bound (3.13) for the good event to verify that

$$\begin{aligned} \mathbb{E} \operatorname{tr}(\mathbf{A}^{(1)} - \llbracket \mathbf{A}^{(1)} \rrbracket_{k-1}) &= \mathbb{E}[\operatorname{tr}(\mathbf{A}^{(1)} - \llbracket \mathbf{A}^{(1)} \rrbracket_{k-1}) \mid \mathbf{G}] \cdot \mathbb{P}(\mathbf{G}) + \mathbb{E}[\operatorname{tr}(\mathbf{A}^{(1)} - \llbracket \mathbf{A}^{(1)} \rrbracket_{k-1}) \mid \mathbf{B}] \cdot \mathbb{P}(\mathbf{B}) \\ &\leq \left(1 + \frac{\lambda_1(\mathbf{A}) + \lambda_k(\mathbf{A})}{\operatorname{tr} \mathbf{A}}\right) \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_k). \end{aligned}$$

Since $\lambda_1(\mathbf{A}) + \lambda_k(\mathbf{A}) \leq \operatorname{tr} \mathbf{A}$, this bound can be weakened to give

$$\mathbb{E} \operatorname{tr}(\mathbf{A}^{(1)} - \llbracket \mathbf{A}^{(1)} \rrbracket_{k-1}) \leq 2 \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_k).$$

By iterating, we conclude that

$$\mathbb{E} \operatorname{tr} \mathbf{A}^{(k)} \leq 2 \mathbb{E} \operatorname{tr}(\mathbf{A}^{(k-1)} - \llbracket \mathbf{A}^{(k-1)} \rrbracket_1) \leq 4 \mathbb{E} \operatorname{tr}(\mathbf{A}^{(k-2)} - \llbracket \mathbf{A}^{(k-2)} \rrbracket_2) \leq \dots \leq 2^k \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_k).$$

This estimate is the statement of the doubling bound. ■

4. Comparison with other methods. This section compares R_PCHOLESKY with other psd low-rank approximation methods. We cover the greedy method (subsection 4.1), uniform sampling (subsection 4.2), DPP sampling (subsection 4.3), RLS sampling (subsection 4.4), and alternative algorithms that are not based on the column Nyström approximation (subsection 4.5). For brevity, the proofs are deferred to Appendix A.

4.1. The greedy method. The *greedy method* (1.2) is a column Nyström approximation with a long history in numerical analysis [16, §10.3] under the name *complete pivoting* or *diagonal pivoting*. Fine and Scheinberg [9] later popularized the method's use for kernel computations. Despite its popularity, the greedy method is known to fail dramatically when applied to some worst-case matrices \mathbf{A} , as discussed by Higham [16, §10.3.1]. Further supporting this conclusion, the following theorem analyzes the large number of columns needed for the greedy method to provide a (r, ε) -approximation.

Theorem 4.1 (Greedy method). *The greedy method has the following error properties:*

- (a) *If $k \geq (1 - (1 + \varepsilon)\eta)N$, the Nyström approximation $\widehat{\mathbf{A}}^{(k)}$ produced by the greedy method is a (r, ε) -approximation (1.4).*
- (b) *For any $N \geq r \geq 1$, $\varepsilon > 0$, and $\eta \in [0, 1]$, there exists a psd matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ such that $\operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r) = \eta \operatorname{tr} \mathbf{A}$ and $\widehat{\mathbf{A}}^{(k)}$ produced by the greedy method is not a (r, ε) -approximation for $k \leq \min\{\eta N, (1 - (1 + \varepsilon)\eta)N\} - 1$.*

Proof. See Appendix A. ■

Theorem 4.1 shows that the greedy method requires accessing $\Omega(N)$ columns to make progress in approximating a worst-case matrix \mathbf{A} . This potential for slow convergence is a major concern in applications, and subsections 1.2 and 2.2 provide examples of kernel matrices for which the greedy method produces worse-quality approximations than other Nyström methods. To address the deficiencies of the greedy method, Gu and Miranian [13] introduced a *strong rank-revealing* pivoting strategy which produces higher-quality approximations on worst-case examples. However, this approach is sufficiently complicated that it has not seen widespread use.

4.2. Uniform sampling. Another popular column Nyström approximation scheme selects pivots *uniformly at random*, either with or without replacement. Uniform sampling without replacement leads to a larger number of unique columns selections, hence, a higher-quality Nyström approximation [17]. Uniform sampling can be effective when the dominant r eigenvectors have mass that is spread out equally over all the coordinates (a property known as “incoherence”) [12]. However, uniform sampling fails when the dominant eigenvectors are heavily concentrated on just a few coordinates. This potential for bad behavior is reflected in the following worst-case error bounds:

Theorem 4.2 (Uniform sampling). *Suppose \mathbf{A} is a psd matrix with ones on the diagonal, and the pivots are sampled uniformly at random either with or without replacement.*

- (a) *If $k \geq \frac{r-1}{\varepsilon\eta} + \frac{1}{\varepsilon}$, the Nyström approximation $\hat{\mathbf{A}}^{(k)}$ produced by uniform sampling is a (r, ε) -approximation.*
- (b) *For any $r \geq 2$ and any $\varepsilon > 0$ and $\eta > 0$ satisfying $\sqrt{\frac{\varepsilon}{\varepsilon+1}} + \eta < 1$, there exists a psd matrix \mathbf{A} such that $\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r) = \eta \text{tr} \mathbf{A}$ and $\hat{\mathbf{A}}^{(k)}$ produced by uniform sampling is not a (r, ε) -approximation for $k < \frac{r-1}{\eta} (\sqrt{\varepsilon^{-1} + 1} - 1)^2$.*

Proof. See Appendix A. ■

Although uniform sampling is typically applied to a kernel matrix \mathbf{A} with ones on the diagonal, a more general *diagonal sampling* [10] method randomly selects the pivots with probabilities proportional to $\text{diag} \mathbf{A}$. Diagonal sampling attains exactly the same error bounds as uniform sampling, as seen from the proof of Theorem 4.2. RPCHOLESKY is equivalent to performing diagonal sampling iteratively on the residual matrix, and it leads to significantly stronger error bounds.

4.3. Determinantal point process sampling. Determinantal point process (DPP) sampling [7] is a column Nyström approximation method that randomly selects a pivot set $S \subseteq \{1, \dots, N\}$ with cardinality $|S| = k$ according to the distribution

$$(4.1) \quad \mathbb{P}\{\{s_1, \dots, s_k\} = S\} = \frac{\det \mathbf{A}(S, S)}{\sum_{|S'|=k} \det \mathbf{A}(S', S')}.$$

DPP sampling satisfies the following error bounds (see [3, Thm. 1], [14]):

Theorem 4.3 (k -DPP sampling). *The Nyström approximation produced by k -DPP sampling is nearly optimal in the following sense:*

- (a) *If $k \geq r/\varepsilon + r - 1$, the approximation $\hat{\mathbf{A}}^{(k)}$ produced by a k -DPP is an (r, ε) -approximation.*
- (b) *For any $r \geq 1$ and $\varepsilon > 0$, there exists a psd matrix \mathbf{A} such that $\hat{\mathbf{A}}^{(k)}$ produced by a k -DPP is not an (r, ε) -approximation for $k < r/\varepsilon + r - 1$.*
- (c) *For any $r \geq 1$ and $\varepsilon > 0$, there exists a matrix \mathbf{A} such that no rank- k column Nyström approximation is an (r, ε) -approximation for $k < r/\varepsilon$.*

Proof. Parts (a) and (b) were essentially proved in [3] and [14]. For part (c), see Appendix A. ■

The theory above demonstrates that DPP sampling has nearly optimal theoretical properties, but implementing DPP sampling for large k values remains expensive relative to peer methods [2, 6]. RPCHOLESKY is much cheaper to perform and is identical to performing 1-DPP sampling iteratively on the residual matrix.

4.4. Ridge leverage score sampling. Like DPP sampling, ridge leverage score (RLS) sampling [1, 4, 20, 26] is a Nyström approximation defined by a distribution on the pivot set S . However, the sampling distribution is potentially more tractable than in DPP sampling since it only depends on a single matrix inverse rather than numerous determinants of $k \times k$ submatrices. For any $\lambda > 0$, the vector of ridge leverage scores is defined as

$$\ell^\lambda = \text{diag}(\mathbf{A}(\mathbf{A} + \lambda \mathbf{I})^{-1}),$$

and the vector of sampling probabilities is defined by [20] to be

$$\mathbf{p} = \min \left\{ 1, 16 \ell^\lambda \log \left(\frac{1}{\delta} \sum_{j=1}^N \ell_j^\lambda \right) \right\}$$

where $\delta \in (0, 1/8)$ is a parameter called the “failure probability”. The coordinate set $S \subseteq \{1, \dots, N\}$ is generated by independently including each index i with probability \mathbf{p}_i . This form of RLS sampling has been analyzed by [20, Thm. 3], and their main result can be summarized as follows:

Theorem 4.4 (Ridge leverage score sampling). *With probability at least $1 - \delta$, the Nyström approximation $\hat{\mathbf{A}}$ produced by RLS sampling satisfies*

$$\|\mathbf{A} - \hat{\mathbf{A}}\| \leq \lambda$$

and

$$\text{rank } \hat{\mathbf{A}} \leq 32 \left(r + \frac{\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)}{\lambda} \right) \log \left(\frac{r}{\delta} + \frac{\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)}{\delta \lambda} \right).$$

Theorem 4.4 is different from the other error bounds presented in this work since it guarantees control in the spectral norm, rather than the trace norm, and it introduces a random approximation rank and a failure probability δ not seen in the other bounds. Due to these differences, the following calculation is needed to convert **Theorem 4.4** into a bound in the expected trace norm:

Corollary 4.5 (Ridge leverage score (r, ε) -approximation). *For any $\varepsilon < \frac{1}{4}$, let $\hat{\mathbf{A}}$ be the Nyström approximation obtained from RLS sampling with*

$$\lambda = \frac{\varepsilon}{2r} \text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r), \quad \delta = \frac{\varepsilon \eta}{2},$$

and set

$$\hat{\mathbf{A}}^{(k)} = \begin{cases} \hat{\mathbf{A}}, & \text{rank } \hat{\mathbf{A}} \leq k, \\ \mathbf{0}, & \text{rank } \hat{\mathbf{A}} > k. \end{cases}$$

Then, $\hat{\mathbf{A}}^{(k)}$ is an (r, ε) -approximation if $k \geq 136 \frac{r}{\varepsilon} \log \left(\frac{3r}{\varepsilon \eta} \right)$.

Proof. See [Appendix A](#). ■

Corollary 4.5 guarantees better worst-case performance for RLS sampling than for the greedy method or uniform sampling, at least when N and η^{-1} are very large. However, the available implementations of RLS sampling [4, 20, 26] are fairly complicated, as they require selecting parameters

λ and δ and approximating the resulting RLS sampling distribution. Because of these preprocessing steps, RLS sampling requires a higher number of entry evaluations than RPCHOLESKY for a fixed approximation rank k .

4.5. Algorithms that are not based on column Nyström approximation. In addition to column Nyström approximation, there are two other popular approaches for randomized low-rank approximation of psd matrices. First, *random embeddings* can produce a high-quality approximation (even higher quality than RPCHOLESKY) if it is possible to compute matrix–vector products with the input matrix. Yet, these methods may be too expensive to deploy in the entry-wise access model. Random embedding methods include the single-pass Nyström method [31], randomized subspace iteration [15], and randomized block Krylov iteration [30].

Second, *random features* (RF) methods can be used to approximate kernel matrices for many well-known kernel functions K [19, 23]. RF methods do not require any entry evaluations, and the computations are highly parallelizable. However, due to the dependence on Monte Carlo sampling, the accuracy is limited by the $1/\sqrt{k}$ convergence rate from the central limit theorem, regardless of the decay of the eigenvalues in the input matrix [35]. In contrast, RPCHOLESKY takes advantage of the eigenvalue decay, so it can attain much higher accuracy.

5. Conclusion. This work has demonstrated the power of RPCHOLESKY for the low-rank approximation of a positive semidefinite kernel matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$. RPCHOLESKY can be programmed in just a few lines of code, and it accelerates many kernel algorithms such as kernel ridge regression and spectral clustering. RPCHOLESKY reduces the computational cost of these algorithms from $\mathcal{O}(N^3)$ operations to just $\mathcal{O}(k^2N)$ operations, where the approximation rank k can be much smaller than the matrix dimension N .

Numerical experiments suggest that RPCHOLESKY dominates other column Nyström approximation algorithms in terms of efficiency and memory footprint. Given a fixed approximation rank k , RPCHOLESKY requires a very small number of entry evaluations, just $(k+1)N$. RPCHOLESKY typically produces approximations that match or improve on the greedy method, uniform sampling, RLS sampling, and DPP sampling. Moreover, theoretical error bounds guarantee that RPCHOLESKY converges nearly as fast as possible in the expected trace norm.

Taken as a whole, this work paves the way for greater use of RPCHOLESKY. The method can potentially be pushed even further, for example, through combinations with other accelerated methods for prediction, clustering, and other learning tasks. The future is indeed bright for this simple yet surprisingly effective algorithm.

Acknowledgements. We thank Mateo Díaz, Zachary Frangella, Marc Gilles, Eitan Levin, Eliza O’Reilly, and Jonathan Weare for helpful discussions and corrections.

Appendix A. More proofs. In this section, we prove error bounds for the greedy method (Theorem 4.1), uniform sampling (Theorem 4.2), DPP sampling (Theorem 4.3(c)), and RLS sampling (Corollary 4.5).

Proof of Theorem 4.1. At each iteration $1 \leq i \leq k$, there are at most $N - i + 1$ nonzero entries in the diagonal of the residual matrix $\mathbf{A}^{(i-1)}$, and the largest entry is incorporated into the pivot set S .

Consequently,

$$\operatorname{tr} \mathbf{A}^{(i)} \leq \left(1 - \frac{1}{N-i+1}\right) \operatorname{tr} \mathbf{A}^{(i-1)} = \frac{N-i}{N-i+1} \operatorname{tr} \mathbf{A}^{(i-1)}.$$

By induction, it follows that $\operatorname{tr} \mathbf{A}^{(k)} \leq \frac{N-k}{N} \operatorname{tr} \mathbf{A}$ and

$$\frac{\operatorname{tr} \mathbf{A}^{(k)}}{\operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)} \leq \left(1 - \frac{k}{N}\right) \frac{\operatorname{tr} \mathbf{A}}{\operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)} = \frac{1 - \frac{k}{N}}{\eta}.$$

If $k \geq (1 - (1 + \varepsilon)\eta)N$, the right-hand side is bounded by $1 + \varepsilon$, establishing part (a).

Next, let M be the smallest integer strictly greater than $(1 - \eta)N$. Set

$$\mathbf{A} = \begin{pmatrix} (1 + \delta)\mathbf{I}_{N-M} & \\ & \mathbf{B}_M \end{pmatrix},$$

where \mathbf{I}_{N-M} is the $(N - M) \times (N - M)$ identity matrix and \mathbf{B}_M is the $M \times M$ matrix of all ones. Set

$$\delta = \frac{M - (1 - \eta)N}{(1 - \eta)(N - M)},$$

so that the \mathbf{B}_M block is responsible for a $1 - \eta$ fraction of the trace. Because $1 + \delta > 1$, the greedy method chooses entries from the $(1 + \delta)\mathbf{I}_{N-M}$ block before the \mathbf{B}_M block. An explicit calculation shows

$$\frac{\operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)})}{\operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_1)} = \frac{1}{\eta} - \frac{k}{N - M} > 1 + \varepsilon$$

for each $k \leq \min\{\eta N, (1 - (1 + \varepsilon)\eta)N\} - 1$. This establishes part (b) for any $r \geq 1$. ■

Proof of Theorem 4.2. Part (a) is a modification of the error bound [10, Eq. (4)], and it is proved using similar techniques. Assume the sampling is conducted with replacement, because sampling with replacement leads to lower error. The trace-norm error takes the form

$$\operatorname{tr}(\mathbf{A} - \mathbf{A}(:, S)\mathbf{A}(S, S)^\dagger \mathbf{A}(S, :)) = \operatorname{tr}(\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)})\mathbf{A}^{1/2}),$$

where $\mathbf{P}_{\mathbf{A}^{1/2}(:, S)}$ is the orthogonal projector onto the range of $\mathbf{A}^{1/2}(:, S)$. Let \mathbf{v}_i denote the i th eigenvector of \mathbf{A} , and calculate

$$(A.1) \quad \operatorname{tr}(\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)})\mathbf{A}^{1/2}) = \operatorname{tr}((\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)})\mathbf{A}(\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)}))$$

$$(A.2) \quad = \sum_{i=1}^N \mathbf{v}_i^* (\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)})\mathbf{A}(\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)})\mathbf{v}_i$$

$$(A.3) \quad \leq \sum_{i=1}^r \mathbf{v}_i^* (\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)})\mathbf{A}(\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)})\mathbf{v}_i + \sum_{i=r+1}^N \mathbf{v}_i^* \mathbf{A}\mathbf{v}_i$$

$$(A.4) \quad = \sum_{i=1}^r \|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P}_{\mathbf{A}^{1/2}(:, S)})\mathbf{v}_i\|^2 + \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r),$$

where (A.1) uses the idempotency of $\mathbf{I} - \mathbf{P}_{A^{1/2}(:,S)}$ and the cyclicity of the trace, (A.2)–(A.3) uses the psd ordering $(\mathbf{I} - \mathbf{P}_{A^{1/2}(:,S)})\mathbf{A}(\mathbf{I} - \mathbf{P}_{A^{1/2}(:,S)}) \preceq \mathbf{A}$, and (A.3)–(A.4) uses the characterization of $\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)$ in terms of eigenvectors of \mathbf{A} .

Next, consider the error term $\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P}_{A^{1/2}(:,S)})\mathbf{v}_i\|^2$. Since $\mathbf{A}^{1/2}\mathbf{P}_{A^{1/2}(:,S)}$ is the optimal approximation to $\mathbf{A}^{1/2}$ in the row space of $\mathbf{A}^{1/2}(\mathbf{S}, :)$, the matrix $\mathbf{A}^{1/2}\mathbf{P}_{A^{1/2}(:,S)}$ can be replaced with any other matrix having the same row space and the error term will only increase or stay the same. In particular,

$$\|\mathbf{A}^{1/2}(\mathbf{I} - \mathbf{P}_{A^{1/2}(:,S)})\mathbf{v}_i\|^2 \leq \left\| \mathbf{A}^{1/2} \left(\mathbf{I} - \frac{1}{k} \sum_{j=1}^k \frac{\mathbf{e}_{s_j} \mathbf{e}_{s_j}^*}{a_{s_j s_j}} \right) \mathbf{v}_i \right\|^2,$$

where \mathbf{e}_{s_i} is the unit vector in the direction of the i th random pivot. Observe that the random vectors

$$\mathbf{A}^{1/2} \frac{\mathbf{e}_{s_i} \mathbf{e}_{s_i}^*}{a_{s_i s_i}} \mathbf{v}_i, \quad i = 1, \dots, k$$

are independent, they have mean $\mathbf{A}^{1/2} \mathbf{v}_i$, and the trace of their covariance matrix is $\text{tr} \mathbf{A} - \lambda_i(\mathbf{A})$. Hence,

$$\mathbb{E} \left\| \mathbf{A}^{1/2} \left(\mathbf{I} - \frac{1}{k} \sum_{j=1}^k \frac{\mathbf{e}_{s_j} \mathbf{e}_{s_j}^*}{a_{s_j s_j}} \right) \mathbf{v}_i \right\|^2 = \frac{\text{tr} \mathbf{A} - \lambda_i(\mathbf{A})}{k}.$$

Last, summing over $i = 1, \dots, r$ guarantees the error bound

$$\mathbb{E} \text{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)}) \leq \frac{r-1}{k} \text{tr} \mathbf{A} + \left(1 + \frac{1}{k}\right) \text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r),$$

which completes part (a) of the theorem.

To prove part (b), consider the $N \times N$ matrix

$$\mathbf{A} = \begin{pmatrix} \mathbf{B}_{N-M(r-1)} & & & \\ & \mathbf{C}_{M,\delta} & & \\ & & \ddots & \\ & & & \mathbf{C}_{M,\delta} \end{pmatrix},$$

where $\mathbf{B}_{N-M(r-1)}$ is the $(N - M(r - 1)) \times (N - M(r - 1))$ matrix of all ones and $\mathbf{C}_{M,\delta}$ is the $M \times M$ matrix with entries

$$\mathbf{C}_{M,\delta} = \begin{pmatrix} 1 & \delta & \cdots & \delta \\ \delta & 1 & \cdots & \delta \\ \vdots & \vdots & \ddots & \vdots \\ \delta & \delta & \cdots & 1 \end{pmatrix}.$$

Set

$$N = \left\lceil \left(1 - \sqrt{\frac{\varepsilon}{1 + \varepsilon}}\right) \frac{(r-1)(M-1)}{\eta} \right\rceil, \quad \delta = 1 - \frac{\eta N}{(r-1)(M-1)},$$

where $\lceil x \rceil$ denotes the smallest integer greater than or equal to x . For large enough M , this construction guarantees that the top r eigenvalues of \mathbf{A} are responsible for a $1 - \eta$ fraction of the trace, and δ lies very close to $\sqrt{\frac{\varepsilon}{1+\varepsilon}}$.

Next, consider the rank- k Nyström approximation $\widehat{\mathbf{A}}^{(k)}$ that selects k_1 columns from the first block, k_2 columns from the second block, etc. The Schur complement of $\mathbf{C}_{M,\delta}$ with respect to any k_i distinct columns has trace

$$(M - k_i) \left(1 + \frac{1}{\delta^{-1} + k_i - 1} \right) (1 - \delta).$$

Therefore, calculate

$$\frac{\text{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)})}{\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)} \geq \frac{1}{r-1} \sum_{i=2}^r \frac{M - k_i}{M-1} \left(1 + \frac{1}{\delta^{-1} + k_i - 1} \right).$$

Using the convexity of $f(x) = \frac{1}{x}$ and the fact that $\mathbb{E} k_i = \frac{Mk}{N}$ for $2 \leq i \leq r$, calculate

$$(A.5) \quad \frac{\mathbb{E} \text{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)})}{\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)} \geq \frac{M-k}{M-1} \left(1 + \frac{1}{\delta^{-1} + \frac{Mk}{N} - 1} \right).$$

For large enough M , (A.5) guarantees

$$\frac{\mathbb{E} \text{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)})}{\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)} > 1 + \varepsilon$$

for each $k < \frac{r-1}{\eta} (\sqrt{\varepsilon^{-1} + 1} - 1)^2$. ■

Proof of Theorem 4.3(c). The proof uses the same construction given in [14, Lem. 6.2]. Consider the $Mr \times Mr$ matrix

$$\mathbf{A} = \begin{pmatrix} \mathbf{C}_{M,\delta} & & & \\ & \mathbf{C}_{M,\delta} & & \\ & & \ddots & \\ & & & \mathbf{C}_{M,\delta} \end{pmatrix},$$

where $\mathbf{C}_{M,\delta}$ is the $M \times M$ matrix with entries

$$\mathbf{C}_{M,\delta} = \begin{pmatrix} 1 & \delta & \cdots & \delta \\ \delta & 1 & & \vdots \\ \vdots & & \ddots & \\ \delta & \cdots & & 1 \end{pmatrix},$$

and consider the rank- k Nyström approximation $\widehat{\mathbf{A}}^{(k)}$ that selects k_1 columns from the first block, k_2 columns from the second block, etc. Similar to the proof of Theorem 4.2, calculate

$$\frac{\text{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)})}{\text{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)} = \frac{1}{r} \sum_{i=1}^r \frac{M - k_i}{M-1} \left(1 + \frac{1}{\delta^{-1} + k_i - 1} \right) \geq \frac{M-k}{M-1} \left(1 + \frac{1}{\delta^{-1} + k/r - 1} \right)$$

Choose M large enough and δ close enough to 1 so that

$$\frac{\operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}^{(k)})}{\operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)} > 1 + \varepsilon$$

for each $k < \frac{r}{\varepsilon}$. ■

Proof of Corollary 4.5. Plugging in the chosen values for λ and δ , [Theorem 4.4](#) guarantees with probability at least $1 - \frac{\varepsilon\eta}{2}$,

$$(A.6) \quad \|\mathbf{A} - \widehat{\mathbf{A}}\| \leq \frac{\varepsilon}{2r} \operatorname{tr}(\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r)$$

and

$$\operatorname{rank} \widehat{\mathbf{A}} \leq 32 \left(r + \frac{2r}{\varepsilon} \right) \log \left(\frac{2r}{\varepsilon\eta} + \frac{4r}{\varepsilon^2\eta} \right).$$

Additionally, since $\widehat{\mathbf{A}} \geq \mathbf{0}$, the Weyl monotonicity principle [[36](#), Thm. 8.11] guarantees

$$(A.7) \quad \lambda_i(\mathbf{A} - \widehat{\mathbf{A}}) \leq \lambda_i(\mathbf{A}), \quad 1 \leq i \leq N.$$

(A.6) and (A.7) together imply

$$\operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}) \leq \left(1 + \frac{\varepsilon}{2} \right) \operatorname{tr} \mathbf{A}$$

With the remaining probability $\frac{\varepsilon\eta}{2}$, the trace-norm error is bounded by

$$\operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}}) \leq \operatorname{tr} \mathbf{A}.$$

Consequently, $\widehat{\mathbf{A}}^{(k)}$ is a (r, ε) -approximation for any

$$(A.8) \quad k \geq 32 \left(r + \frac{2r}{\varepsilon} \right) \log \left(\frac{2r}{\varepsilon\eta} + \frac{4r}{\varepsilon^2\eta} \right).$$

To complete the proof, bound the right-hand side of (A.8) using $136 \frac{r}{\varepsilon} \log \left(\frac{3r}{\varepsilon\eta} \right)$. ■

Appendix B. Details of numerical experiments. The numerical experiments in [subsection 1.2](#) and [section 2](#) are implemented in Python, with code available at <https://github.com/epperly/Randomly-Pivoted-Cholesky>. Below we provide further implementation details for some of the Nyström methods:

RPCHOLESKY. All RPCHOLESKY results in [subsection 1.2](#) and [section 2](#) use the standard algorithm ([Algorithm 2.1](#)). However, the blocked RPCHOLESKY algorithm ([Algorithm 2.2](#)) leads to near-identical accuracy when applied to the $10^5 \times 10^5$ kernel matrix from [subsection 2.2](#) with a block size $B = 50$. For this data set, blocked RPCHOLESKY yields a $5\times$ speedup with the ℓ_1 Laplace kernel ([2.5](#)) and a $20\times$ speedup after switching to the Gaussian kernel.

DPP sampling. DPP sampling is implemented using the DPPy Python package [[11](#)]. Yet, the samplers from this package all produce error messages when tried on certain inputs, particularly

when k is large and \mathbf{A} is nearly low-rank. The theoretically fastest alpha sampler cannot be used for our numerical comparisons since it gives error messages even when approximating the **Outliers** matrix with $k = 20$. Rather, **Figures 2** and **4** are produced using the `vfx` sampler introduced in [6], and **Figure 1** is produced using the comparatively slow GS sampler which requires a full eigen-decomposition of \mathbf{A} . In this work, all DPP samplers generate exact samples from the k -DPP distribution (4.1); we did not carefully study inexact samplers based on MCMC [2].

RLS sampling. Several RLS sampling algorithms have been introduced in the literature, including recursive ridge leverage scores (RRLS) [20], SQUEAK [4], and BLESS [26]. The analysis and experiments in this paper are based on the RRLS algorithm [20], which is implemented in the `recursiveNystrom` method [32] for Python. We find little difference in accuracy between the standard and “accelerated” `recursiveNystrom` methods for our experiments, so we always use the accelerated method with the superior runtime.

For comparison, we have also implemented RLS sampling by directly forming the matrix $\mathbf{A}(\mathbf{A} + \lambda \mathbf{I})^{-1}$ and sampling k indices without replacement using the normalized ridge leverage scores as probabilities. Even with the ridge leverage scores computed exactly in this way, the relative trace-norm error for the **Smile** example with $k = 100$ is still roughly 10^{-2} , five orders of magnitude larger than the error due to `RPCHOLESKY`. This indicates that the poor performance of RLS on this example is not due to errors in the approximation of the ridge leverage scores.

REFERENCES

- [1] A. ALAOUI AND M. W. MAHONEY, *Fast randomized kernel ridge regression with statistical guarantees*, in *Advances in Neural Information Processing Systems*, vol. 28, 2015, <https://proceedings.neurips.cc/paper/2015/file/f3f27a324736617f20abf2ffd806f6d-Paper.pdf>.
- [2] N. ANARI, S. OVEIS GHARAN, AND A. REZAEI, *Monte Carlo Markov chain algorithms for sampling strongly Rayleigh distributions and determinantal point processes*, in *29th Annual Conference on Learning Theory*, vol. 49 of *Proceedings of Machine Learning Research*, 2016, pp. 103–115, <https://proceedings.mlr.press/v49/anari16.html>.
- [3] M.-A. BELABBAS AND P. J. WOLFE, *Spectral methods in machine learning and new strategies for very large datasets*, *Proceedings of the National Academy of Sciences*, 106 (2009), pp. 369–374, <https://doi.org/10.1073/pnas.0810600105>.
- [4] D. CALANDRIELLO, A. LAZARIC, AND M. VALKO, *Distributed adaptive sampling for kernel matrix approximation*, in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, vol. 54 of *Proceedings of Machine Learning Research*, 2017, pp. 1421–1429, <https://proceedings.mlr.press/v54/calandriello17a.html>.
- [5] E. CARLEN, *Trace inequalities and quantum entropy: An introductory course*, in *Entropy and the Quantum*, R. Sims and D. Ueltschi, eds., vol. 529 of *Contemporary Mathematics*, Providence, RI, 2010, American Mathematical Society, pp. 73–140, <https://doi.org/http://dx.doi.org/10.1090/conm/529>.
- [6] M. DEREZINSKI, D. CALANDRIELLO, AND M. VALKO, *Exact sampling of determinantal point processes with sublinear time preprocessing*, in *Advances in Neural Information Processing Systems*, vol. 32, 2019, <https://proceedings.neurips.cc/paper/2019/file/fa3060edb66e6ff4507886f9912e1ab9-Paper.pdf>.
- [7] M. DEREZINSKI AND M. W. MAHONEY, *Determinantal point processes in randomized numerical linear algebra*, *Notices of the American Mathematical Society*, 68 (2021), p. 1, <https://doi.org/10.1090/noti2202>.
- [8] A. DESHPANDE AND S. VEMPALA, *Adaptive sampling and fast low-rank matrix approximation*, in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, J. Díaz, K. Jansen, J. D. P. Rolim, and U. Zwick, eds., *Lecture Notes in Computer Science*, Berlin, 2006, Springer, pp. 292–303, https://doi.org/10.1007/11830924_28.
- [9] S. FINE AND K. SCHEINBERG, *Efficient SVM training using low-rank kernel representations*, *Journal of Machine Learning Research*, 2 (2002), p. 243–264, <https://dl.acm.org/doi/10.5555/944790.944812>.

- [10] A. FRIEZE, R. KANNAN, AND S. VEMPALA, *Fast Monte-Carlo algorithms for finding low-rank approximations*, Journal of the ACM, 51 (2004), pp. 1025–1041, <https://doi.org/10.1145/1039488.1039494>.
- [11] G. GAUTIER, G. POLITO, R. BARDENET, AND M. VALKO, *DPPy: DPP sampling with Python*, Journal of Machine Learning Research, 20 (2019), pp. 1–7, <http://jmlr.org/papers/v20/19-179.html>.
- [12] A. GITTENS, *The spectral norm error of the naïve Nyström extension*, arXiv:1110.5305 [cs, math], (2011).
- [13] M. GU AND L. MIRANIAN, *Strong rank revealing Cholesky factorization*, Electronic Transactions on Numerical Analysis, 17 (2004), pp. 76–92, <http://eudml.org/doc/125035>.
- [14] V. GURUSWAMI AND A. K. SINOP, *Optimal column-based low-rank matrix reconstruction*, in Proceedings of the 2012 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2012, pp. 1207–1214, <https://doi.org/10.1137/1.9781611973099.95>.
- [15] N. HALKO, P. G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288, <https://doi.org/10.1137/090771806>.
- [16] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, 2nd ed. ed., 2002, <https://doi.org/10.1137/1.9780898718027>.
- [17] S. KUMAR, M. MOHRI, AND A. TALWALKAR, *Sampling techniques for the Nyström method*, in Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics, vol. 5 of Proceedings of Machine Learning Research, 2009, pp. 304–311, <https://proceedings.mlr.press/v5/kumar09a.html>.
- [18] C.-K. LI AND R. MATHIAS, *Extremal characterizations of the Schur complement and resulting inequalities*, SIAM Review, 42 (2000), pp. 233–246, <https://doi.org/10.1137/S0036144599337290>.
- [19] P.-G. MARTINSSON AND J. A. TROPP, *Randomized numerical linear algebra: Foundations and algorithms*, Acta Numerica, 29 (2020), p. 403–572, <https://doi.org/10.1017/S0962492920000021>.
- [20] C. MUSCO AND C. MUSCO, *Recursive sampling for the Nyström method*, in Advances in Neural Information Processing Systems, vol. 30, 2017, <https://proceedings.neurips.cc/paper/2017/file/a03fa30821986dff10fc66647c84c9c3-Paper.pdf>.
- [21] C. MUSCO AND D. P. WOODRUFF, *Sublinear time low-rank approximation of positive semidefinite matrices*, in 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), 2017, pp. 672–683, <https://doi.org/10.1109/FOCS.2017.68>.
- [22] J. POULSON, *High-performance sampling of generic determinantal point processes*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 378 (2020), p. 20190059, <https://doi.org/10.1098/rsta.2019.0059>.
- [23] A. RAHIMI AND B. RECHT, *Random features for large-scale kernel machines*, in Advances in Neural Information Processing Systems, vol. 20, 2007, <https://proceedings.neurips.cc/paper/2007/file/013a006f03dbc5392effeb8f18fda755-Paper.pdf>.
- [24] R. RAMAKRISHNAN, P. O. DRAL, M. RUPP, AND O. A. VON LILIENFELD, *Quantum chemistry structures and properties of 134 kilo molecules*, Scientific Data, 1 (2014), p. 140022, <https://doi.org/10.1038/sdata.2014.22>.
- [25] L. RUDDIGKEIT, R. VAN DEURSEN, L. C. BLUM, AND J.-L. REYMOND, *Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17*, Journal of Chemical Information and Modeling, 52 (2012), pp. 2864–2875, <https://doi.org/10.1021/ci300415d>.
- [26] A. RUDI, D. CALANDRIELLO, L. CARRATINO, AND L. ROSASCO, *On fast leverage score sampling and optimal learning*, in Advances in Neural Information Processing Systems, vol. 31, 2018, <https://proceedings.neurips.cc/paper/2018/file/56584778d5a8ab88d6393cc4cd11e090-Paper.pdf>.
- [27] B. SCHÖLKOPF AND A. J. SMOLA, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, The MIT Press, Cambridge, MA, 2018, <https://doi.org/10.7551/mitpress/4175.001.0001>.
- [28] A. STUKE, M. TODOROVIĆ, M. RUPP, C. KUNKEL, K. GHOSH, L. HIMANEN, AND P. RINKE, *Chemical diversity in molecular orbital energy predictions with kernel ridge regression*, The Journal of Chemical Physics, 150 (2019), p. 204121, <https://doi.org/10.1063/1.5086105>.
- [29] A. SZABO AND N. S. OSTLUND, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*, Dover Publications, Mineola, NY, 2017.
- [30] J. A. TROPP AND R. J. WEBBER, *Randomized algorithms for low-rank matrix approximation: Design, analysis, and applications*, In preparation, (2022).
- [31] J. A. TROPP, A. YURTSEVER, M. UDELL, AND V. CEVHER, *Fixed-rank approximation of a positive-semidefinite matrix from streaming data*, in Advances in Neural Information Processing Systems, vol. 30, 2017, <https://proceedings.neurips.cc/paper/2017/file/4558dbb6f6f8bb2e16d03b85bde76e2c-Paper.pdf>.

- [32] A. VANRAES, C. MUSCO, AND C. MUSCO, *Python implementation of the RLS-Nyström Method*, July 19 2019, <https://github.com/axelv/recursive-nystrom>.
- [33] U. VON LUXBURG, *A tutorial on spectral clustering*, *Statistics and Computing*, 17 (2007), pp. 395–416, <https://doi.org/10.1007/s11222-007-9033-z>.
- [34] C. WILLIAMS AND M. SEEGER, *Using the Nyström method to speed up kernel machines*, in *Advances in Neural Information Processing Systems*, vol. 13, 2000, <https://proceedings.neurips.cc/paper/2000/file/19de10adbaa1b2ee13f77f679fa1483a-Paper.pdf>.
- [35] T. YANG, Y.-F. LI, M. MAHDAVI, R. JIN, AND Z.-H. ZHOU, *Nyström method vs random Fourier features: A theoretical and empirical comparison*, in *Advances in Neural Information Processing Systems*, vol. 25, 2012, <https://proceedings.neurips.cc/paper/2012/file/621bf66ddb7c962aa0d22ac97d69b793-Paper.pdf>.
- [36] F. ZHANG, *Matrix Theory: Basic Results and Techniques*, Springer, New York, 2nd ed., 2011, <https://link.springer.com/book/10.1007/978-1-4614-1099-7>.