

EMBRACE REJECTION: KERNEL MATRIX APPROXIMATION BY ACCELERATED RANDOMLY PIVOTED CHOLESKY*

ETHAN N. EPPERLY[†], JOEL A. TROPP[†], AND ROBERT J. WEBBER[‡]

Abstract. Randomly pivoted Cholesky (RPCHOLESKY) is an algorithm for constructing a low-rank approximation of a positive-semidefinite matrix using a small number of columns. This paper develops an accelerated version of RPCHOLESKY that employs block matrix computations and rejection sampling to efficiently simulate the execution of the original algorithm. For the task of approximating a kernel matrix, the accelerated algorithm can run over $40\times$ faster. The paper contains implementation details, theoretical guarantees, experiments on benchmark data sets, and an application to computational chemistry.

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

MSC codes. 65F55, 65C99, 68T05

DOI. 10.1137/24M1699048



See reproducibility of computational results at end of the article.

1. Introduction. This paper treats a core task in computational linear algebra:

Low-rank psd approximation: Given a positive-semidefinite (psd) matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$ and a rank parameter $1 \leq k < N$, compute a matrix $\mathbf{F} \in \mathbb{C}^{N \times k}$ such that $\mathbf{A} \approx \mathbf{F}\mathbf{F}^*$.

Randomly pivoted Cholesky (RPCHOLESKY) is a randomized variant of the pivoted partial Cholesky method that stands among the best algorithms for constructing a rank- k psd approximation from a set of k columns [4]. The present work introduces **accelerated RPCHOLESKY**, a new algorithm that runs up to $40\times$ faster than RPCHOLESKY, while producing approximations with the same quality. The accelerated method simulates the behavior of the simpler method by means of block-matrix computations and rejection sampling. Its benefits are most pronounced when the matrix dimension N and the approximation rank k are both moderately large (say, $N \geq 10^5$ and $k \geq 10^3$), a regime where existing algorithms struggle.

1.1. Access models for matrices. When designing an algorithm for low-rank approximation of a psd matrix \mathbf{A} , one must consider how the algorithm interacts

*Received by the editors October 7, 2024; accepted for publication (in revised form) by A. K. Saibaba July 30, 2025; published electronically November 17, 2025.

<https://doi.org/10.1137/24M1699048>

Funding: The first author 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 DE-SC0021110. The second and third authors acknowledge support from the Office of Naval Research through Awards N00014-18-1-2363 and N00014-24-1-2223, from the National Science Foundation through FRG Award 1952777, and from the Caltech Carver Mead New Adventures Fund.

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

[‡]Department of Mathematics, University of California San Diego, La Jolla, CA 92093 USA (rwebber@ucsd.edu).

with the matrix. For example, in the *matvec access model*, the primitive operation is the matrix–vector product $\mathbf{v} \mapsto \mathbf{A}\mathbf{v}$, and the algorithm designer aims to construct the low-rank approximation with the minimum number of matrix–vector products [26, 32, 36]. The matvec access model can arise from the discretization of differential and integral operators, among other applications.

In the *entry access model*, the primitive operation is the evaluation of the matrix entry $\mathbf{A}(i, j)$ for any pair (i, j) of row and column indices. The papers [4, 14, 16, 27, 30, 35] contain efficient algorithms for low-rank psd approximation in the entry access model. These methods are appealing because they only access and manipulate a small fraction of the entries in the target matrix.

The entry access model does not always provide a satisfactory abstraction for scientific computing and machine learning because it fails to account for computational efficiencies observed in practice by accessing the matrix in a block-wise fashion. To remedy this deficiency, this paper works in a paradigm that will be called the *submatrix access model*. The primitive operation is extracting a submatrix:

$$(1.1) \quad \mathbf{A}(\mathbf{S}_1, \mathbf{S}_2) = [\mathbf{A}(i, j)]_{i \in \mathbf{S}_1, j \in \mathbf{S}_2} \quad \text{for } \mathbf{S}_1, \mathbf{S}_2 \subseteq \{1, \dots, N\}.$$

This access model describes the setting where the cost of generating a moderately large submatrix is comparable with the cost of evaluating a single matrix entry. As the next subsection explains, the submatrix access model reflects the computational challenge that arises from kernel methods in machine learning.

1.2. Kernel matrices and the submatrix access model. Recall that a positive-definite kernel function $\kappa : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ can be used to define a similarity metric between pairs of data points [31]. Popular kernel functions include the Gaussian and ℓ_1 Laplace kernels:

$$(1.2) \quad \kappa(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{x}'\|^2\right), \quad \text{[Gaussian]}$$

$$(1.3) \quad \kappa(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{1}{\sigma} \sum_{i=1}^d |\mathbf{x}(i) - \mathbf{x}'(i)|\right). \quad \text{[}\ell_1 \text{ Laplace]}$$

Here, $\|\cdot\|$ is the ℓ_2 norm and $\sigma > 0$ is the bandwidth, which sets the scale of interactions. Given data points $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^d$, the kernel matrix $\mathbf{A} = [\kappa(\mathbf{x}_i, \mathbf{x}_j)]_{1 \leq i, j \leq N}$ is a psd matrix that tabulates the similarity between each pair of points. Kernel methods perform dense linear algebra with the kernel matrix to accomplish machine learning tasks, such as regression and clustering.

On modern computers, the runtime of numerical algorithms is often dominated by time moving data between cache and memory rather than time performing arithmetic operations [13, pp. 13, 581–582]. Each time columns of a kernel matrix are generated, whether a single column or a large block of $b \gg 1$ columns, all N data points $\{\mathbf{x}_i\}_{i=1}^N$ must be moved from memory to cache. When data movement dominates the runtime, generating columns in blocks of size b can be up to $b \times$ as fast as generating columns one-by-one. See section A for discussion of other reasons why blocking can increase the speed of matrix algorithms.

Figure 1 provides empirical evidence that the submatrix access model is the appropriate abstraction for kernel matrix computations. The figure charts the number of columns formed per second when evaluating blocks of b columns in a $10^5 \times 10^5$ kernel matrix with varying block sizes $1 \leq b \leq 10^3$. The per-column efficiency increases with b , and it saturates once the block size b is large enough to maximally reap the

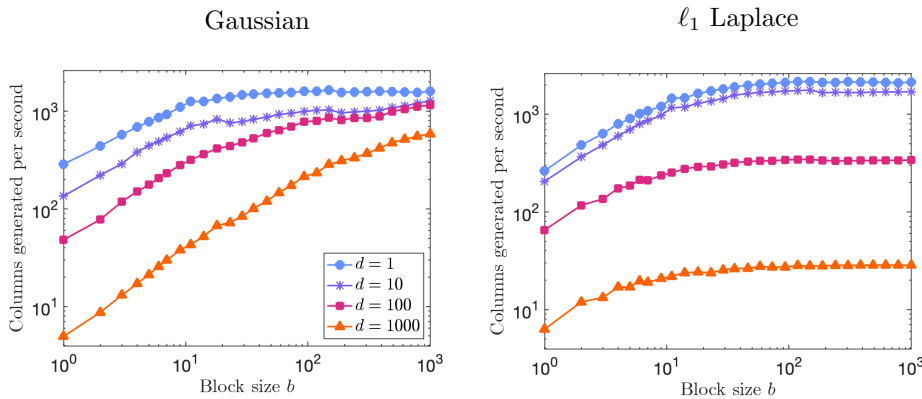


FIG. 1. Columns generated per second for Gaussian (left) and ℓ_1 Laplace (right) kernel matrices with bandwidth $\sigma = 1$. The data consists of $N = 10^5$ standard Gaussian points with dimension $d \in \{1, 10, 100, 1000\}$.

benefits of block-wise access. This critical block size and the maximum speedup is sensitive both to the dimension d of the data and the choice of kernel. In the most extreme case (Gaussian kernel with $d = 1000$), generating columns in blocks of size $b = 1000$ is $100\times$ as efficient as generating columns one at a time.

For moderately large data sets (say, $N \geq 10^5$), it is prohibitively expensive to generate and store the full kernel matrix. Fortunately, the kernel matrix often has rapidly decaying eigenvalues, so it can be approximated by a low-rank matrix constructed from a subset of the columns. When properly implemented with block-wise access, this approximation makes it possible to scale kernel methods to billions of data points [25]. The resulting machine learning methods can achieve cutting-edge speed and interpretability for tasks in scientific machine learning [2].

1.3. Efficient low-rank approximation in the submatrix access model.

Given the large speedups observed in Figure 1, it is natural to ask a question:

What algorithm for low-rank psd approximation is most efficient in the submatrix access model?

This paper answers the question by proposing the accelerated RPCHOLESKY algorithm. To contextualize this new method, the rest of this subsection reviews simple RPCHOLESKY (subsection 1.3.1) and block RPCHOLESKY (subsection 1.3.2). Last, subsection 1.3.3 takes a first look at accelerated RPCHOLESKY.

1.3.1. Simple RPCholesky. Randomly pivoted Cholesky (RPCHOLESKY) is an algorithm for low-rank psd approximation that is designed for the entry access model [4]. After extracting the diagonal of the matrix, the algorithm forms a rank- k approximation by adaptively evaluating and manipulating k columns of the matrix.

Starting with the trivial approximation $\hat{\mathbf{A}}^{(0)} := \mathbf{0}$ and the initial residual $\mathbf{A}^{(0)} := \mathbf{A}$, RPCHOLESKY builds up a rank- k approximation $\hat{\mathbf{A}}^{(k)}$ by taking the following steps. For $i = 0, \dots, k - 1$,

1. **Choose a random pivot.** Select a random pivot $s_{i+1} \in \{1, \dots, N\}$ according to the probability distribution

$$(1.4) \quad \mathbb{P}\{s_{i+1} = j\} = \frac{\mathbf{A}^{(i)}(j, j)}{\text{tr } \mathbf{A}^{(i)}}.$$

2. **Update.** Evaluate the column $\mathbf{A}^{(i)}(:, s_{i+1})$ of the residual indexed by the selected pivot s_{i+1} . Update the approximation and the residual:

$$\begin{aligned} \widehat{\mathbf{A}}^{(i+1)} &:= \widehat{\mathbf{A}}^{(i)} + \frac{\mathbf{A}^{(i)}(:, s_{i+1})\mathbf{A}^{(i)}(s_{i+1}, :)}{\mathbf{A}^{(i)}(s_{i+1}, s_{i+1})}; \\ \mathbf{A}^{(i+1)} &:= \mathbf{A}^{(i)} - \frac{\mathbf{A}^{(i)}(:, s_{i+1})\mathbf{A}^{(i)}(s_{i+1}, :)}{\mathbf{A}^{(i)}(s_{i+1}, s_{i+1})}. \end{aligned}$$

In this algorithm and those that follow, “pivot” refers to the index of the selected column. Pivots are chosen by means of an adaptive, randomized selection rule. Additionally, $\mathbf{A}(\dots)$ indicates access of \mathbf{A} by MATLAB-style indexing notation; see subsection 1.6 for a discussion of notation.

With an optimized implementation (Algorithm D.1), RPCHOLESKY produces the approximation $\widehat{\mathbf{A}} = \widehat{\mathbf{A}}^{(k)}$ in factored form after $\mathcal{O}(k^2N)$ arithmetic operations. The history of RPCHOLESKY and related algorithms [4, 9, 10, 28] is surveyed in [4, sec. 3].

1.3.2. Block RPCholesky. RPCHOLESKY is less efficient than it could be because the algorithm evaluates columns one by one, failing to take advantage of the submatrix access model. For this reason, large-scale machine learning applications [2, 11] often employ the block RPCHOLESKY algorithm [4], which is faster. This method draws $b > 1$ random pivots with replacement according to the distribution (1.4). Then it updates the approximation and residual using the b pivots. For target rank k , the block size $b = k/10$ is often used in applications. The pseudocode for block RPCHOLESKY is available in Algorithm D.2 in the appendix.

This paper resolves several open empirical and theoretical questions about the performance of block RPCHOLESKY. First, new experiments confirm that simple RPCHOLESKY and block RPCHOLESKY provide similar approximation accuracy on typical problems (Figure 3), but there are challenging examples where the simple method is over $100\times$ more accurate than the block method (Figures 2 and 3). Second, new error bounds guarantee the eventual success of block RPCHOLESKY for any block size b (Theorem 4.2). However, the bounds show that the efficiency of the block method can degrade when the target matrix has rapid eigenvalue decay.

1.3.3. Accelerated RPCholesky. The main contribution of this paper is a new accelerated RPCHOLESKY algorithm, which combines the best features of the simple and block methods. Accelerated RPCHOLESKY produces the same distribution of random pivots as simple RPCHOLESKY and enjoys the same accuracy guarantees, while taking advantage of the submatrix access model to improve the runtime.

At each step, accelerated RPCHOLESKY proposes a block of b random pivots. Then the algorithm thins the list of proposals using rejection sampling, and it updates the matrix approximation using the accepted pivots. As compared with block RPCHOLESKY, the new feature of accelerated RPCHOLESKY is the rejection sampling step, which simulates the pivot distribution of simple RPCHOLESKY.

Here is a summary of the accelerated RPCHOLESKY method; a comprehensive explanation appears in section 2. The user fixes a block size parameter $b \geq 1$. Starting with the approximation $\widehat{\mathbf{A}}^{(0)} := \mathbf{0}$ and the residual $\mathbf{A}^{(0)} := \mathbf{A}$, accelerated RPCHOLESKY builds up a low-rank approximation $\widehat{\mathbf{A}}^{(t)}$ with rank at most bt by means of the following procedure. For $i = 0, \dots, t - 1$,

1. **Propose a block of random pivots.** Propose a block of independent and identically distributed random pivots $\{s'_{bi+1}, \dots, s'_{b(i+1)}\}$ with

$$\mathbb{P}\{s'_\ell = j\} = \frac{\mathbf{A}^{(i)}(j, j)}{\text{tr } \mathbf{A}^{(i)}}.$$

2. **Rejection sampling.** Initialize the incremental pivot set $S_i \leftarrow \emptyset$. For each $\ell = bi + 1, \dots, b(i + 1)$, accept the pivot s'_ℓ with probability

$$p(s'_\ell) = \frac{\mathbf{A}^{(i)}(s'_\ell, s'_\ell) - \mathbf{A}^{(i)}(s'_\ell, S_i) \mathbf{A}^{(i)}(S_i, S_i)^\dagger \mathbf{A}^{(i)}(S_i, s'_\ell)}{\mathbf{A}^{(i)}(s'_\ell, s'_\ell)},$$

and update $S_i \leftarrow S_i \cup \{s'_\ell\}$. Otherwise, reject the pivot and do nothing.

3. **Update.** Update the approximation:

$$\begin{aligned} \widehat{\mathbf{A}}^{(i+1)} &:= \widehat{\mathbf{A}}^{(i)} + \mathbf{A}^{(i)}(:, S_i) \mathbf{A}^{(i)}(S_i, S_i)^\dagger \mathbf{A}^{(i)}(S_i, :); \\ \mathbf{A}^{(i+1)} &:= \mathbf{A}^{(i)} - \mathbf{A}^{(i)}(:, S_i) \mathbf{A}^{(i)}(S_i, S_i)^\dagger \mathbf{A}^{(i)}(S_i, :). \end{aligned}$$

For numerical stability and efficiency, the pseudoinverse $\mathbf{A}^{(i)}(S_i, S_i)^\dagger$ should be computed and applied via a Cholesky decomposition of $\mathbf{A}^{(i)}(S_i, S_i)$.

Because of the possibility of rejections, this algorithm uses slightly more entry look-ups than simple RPCHOLESKY to generate a rank- k approximation. However, it is much faster in practice due to blockwise computations. Algorithm 2.2 provides an optimized implementation.

1.4. Illustrative comparison. Figure 2 presents a stylized example where simple RPCHOLESKY, block RPCHOLESKY, and accelerated RPCHOLESKY are applied to a Gaussian kernel matrix, which is generated from $N = 10^5$ data points forming a smile in \mathbb{R}^2 of radius 10. The bandwidth is set to $\sigma = 0.2$. This example is only intended to illustrate the differences between the RPCHOLESKY-type algorithms; subsection 3.1 provides more extensive comparisons on a testbed of synthetic data, random point clouds, and real-world data.

As shown in the left panel of Figure 2, both block RPCHOLESKY and accelerated RPCHOLESKY are $4\times$ to $5\times$ faster than simple RPCHOLESKY on this example. The right panel of Figure 2 shows the relative trace error




FIG. 2. Runtime (left) and relative trace norm error (1.5) (right) of RPCHOLESKY, block RPCHOLESKY, and accelerated RPCHOLESKY applied to a Gaussian kernel matrix \mathbf{A} associated with 10^5 points in \mathbb{R}^2 forming a smile (inset, right). The blocked methods use block size $b = 120$.

$$(1.5) \quad \frac{\operatorname{tr}(\mathbf{A} - \widehat{\mathbf{A}})}{\operatorname{tr}(\mathbf{A})}$$

of the methods. Accelerated RPCHOLESKY and simple RPCHOLESKY give the same accuracy, so the error curves coincide with each other. Block RPCHOLESKY fares much worse, attaining $800\times$ higher error than the other methods at rank $k = 1000$. This example supports the basic takeaway message: **accelerated RPCholesky produces the same distribution of random outputs as simple RPCholesky while being much faster.**

1.5. Outline. This paper explores the theoretical and empirical properties of accelerated RPCHOLESKY and compares it to simple and block RPCHOLESKY. Section 2 introduces accelerated RPCHOLESKY and discusses efficient implementations. Section 3 presents numerical experiments. Section 4 analyzes accelerated RPCHOLESKY and block RPCHOLESKY. Section 5 extends these methods and proposes an accelerated randomly pivoted QR method.

1.6. Notation. Entries and submatrices of \mathbf{A} are expressed using MATLAB notation. For example, $\mathbf{A}(:, i)$ represents the i th column of \mathbf{A} and $\mathbf{A}(\mathbf{S}, :)$ denotes the submatrix of \mathbf{A} with rows indexed by the set \mathbf{S} . The conjugate transpose of a matrix \mathbf{F} is denoted as \mathbf{F}^* , and the Moore–Penrose pseudoinverse is \mathbf{F}^\dagger . The inverse of the conjugate transpose of a nonsingular matrix is denoted $\mathbf{F}^{-*} := (\mathbf{F}^{-1})^* = (\mathbf{F}^*)^{-1}$. The matrix $\mathbf{\Pi}_{\mathbf{F}}$ is the orthogonal projector onto the column span of \mathbf{F} . The function $\lambda_i(\mathbf{A})$ outputs the i th largest eigenvalue of a psd matrix \mathbf{A} . The symbol \preceq denotes the psd order on Hermitian matrices: $\mathbf{H} \preceq \mathbf{A}$ if and only if $\mathbf{A} - \mathbf{H}$ is psd. A matrix \mathbf{A} is described as “rank- r ” when $\operatorname{rank} \mathbf{A} \leq r$. Last, $\|\cdot\|$ denotes the vector ℓ_2 norm.

2. Faster RPCholesky by rejection sampling. This section explains how RPCHOLESKY can be implemented efficiently with rejection sampling (subsections 2.1 and 2.2). Afterward, subsection 2.3 introduces and provides the pseudocode for accelerated RPCHOLESKY. Subsection 2.4 explains how to implement accelerated RPCHOLESKY with more operations but lower memory requirements.

2.1. Implementing RPCholesky with rejection sampling. Rejection sampling is a standard approach for sampling from a challenging probability distribution [24, Sec. 2.2]. This section shows how to use rejection sampling to sample the pivots s_1, \dots, s_k according to the RPCHOLESKY distribution. Rejection sampling was originally applied to RPCHOLESKY in [15], building on related ideas from [29].

Consider the following setting. Suppose that one has already sampled pivots $\mathbf{S}_i = \{s_1, \dots, s_i\}$, which define the rank- i approximation $\widehat{\mathbf{A}}^{(i)} = \mathbf{A}(:, \mathbf{S}_i)\mathbf{A}(\mathbf{S}_i, \mathbf{S}_i)^\dagger\mathbf{A}(\mathbf{S}_i, :)$ and the residual $\mathbf{A}^{(i)} = \mathbf{A} - \widehat{\mathbf{A}}^{(i)}$. In addition, suppose that one has access to an upper bound \mathbf{u} on the diagonal of the residual $\mathbf{A}^{(i)}$:

$$\mathbf{u}(j) \geq \mathbf{A}^{(i)}(j, j) \quad \text{for each } j = 1, \dots, N.$$

The diagonal of the residual decreases at each step of the RPCHOLESKY procedure: $\mathbf{A}^{(i)}(j, j) \leq \mathbf{A}^{(i-1)}(j, j) \leq \dots \leq \mathbf{A}(j, j)$ for each index j . Thus, $\mathbf{u} = \operatorname{diag}(\mathbf{A})$ is one possible choice for the upper bound.

To select the next pivot s_{i+1} , rejection sampling executes the following loop:

Rejection sampling loop:

1. **Propose.** Draw a random pivot s according to the *proposal distribution* \mathbf{u} , given by $\mathbb{P}\{s = j\} = \mathbf{u}(j) / \sum_{k=1}^n \mathbf{u}(k)$.
2. **Accept/reject.** With probability $\mathbf{A}^{(i)}(s, s) / \mathbf{u}(s)$, accept the pivot and set $s_{i+1} \leftarrow s$. Otherwise, reject the pivot and return to step 1.

Notice that steps 1 and 2 are repeated until a pivot s_{i+1} is accepted. Thus, the procedure generates a set $S_{i+1} \leftarrow S_i \cup \{s_{i+1}\}$ of $i + 1$ pivots, defining a rank- $(i + 1)$ approximation $\hat{\mathbf{A}}^{(i+1)}$. Continuing in this way, one can sample pivots s_{i+2}, s_{i+3}, \dots and construct a low-rank approximation of the desired rank k .

The advantage of rejection sampling over the standard RPCHOLESKY implementation (Algorithm D.1) is that rejection sampling only evaluates the diagonal entries $\text{diag } \mathbf{A}^{(i)}(s, s)$ at the proposed pivot indices s , whereas Algorithm D.1 evaluates all the entries of $\text{diag } \mathbf{A}^{(i)}$ at every step. Nonetheless, the efficiency of rejection sampling depends on the typical size of the acceptance probability $\mathbf{A}^{(i)}(s, s) / \mathbf{u}(s)$. A small acceptance probability leads to many executions of the loop 1–2 and thus a slow algorithm. Accelerated RPCHOLESKY avoids this issue by periodically updating \mathbf{u} to keep the acceptance rate from becoming too low.

2.2. Efficient rejection sampling in the submatrix access model. This section describes how to implement RPCHOLESKY with rejection sampling in the submatrix access model. The resulting procedure, called REJECTIONSAMPLESUBMATRIX, serves as a subroutine in the full accelerated RPCHOLESKY algorithm, which is developed in subsection 2.3.

REJECTIONSAMPLESUBMATRIX works as follows. Assume the proposal distribution $\mathbf{u} := \text{diag } \mathbf{A}$, and assume the initial pivot set is empty: $S_0 = \emptyset$. To take advantage of the submatrix access model, begin by sampling a block $S' = \{s'_1, \dots, s'_b\}$ of $b \geq 1$ proposals independently from the distribution

$$\mathbb{P}\{s' = j\} = \frac{\mathbf{A}(j, j)}{\text{tr}(\mathbf{A})} \quad \text{for } j = 1, \dots, N.$$

For efficiency, generate the entire submatrix $\mathbf{A}(S', S')$. Finally, perform b pivot accept or reject steps with probabilities given by the diagonal submatrix elements to obtain a set $S \subseteq S'$ of accepted pivots. Each time a pivot is accepted, the procedure performs a Cholesky step to eliminate the accepted column from the submatrix and to update the acceptance probabilities. The pseudocode for REJECTIONSAMPLESUBMATRIX is given in Algorithm 2.1. As inputs, it takes the proposed pivots S' and the submatrix $\mathbf{H} := \mathbf{A}(S', S')$, and it returns the set $S \subseteq S'$ of accepted pivots and a Cholesky factorization $\mathbf{A}(S, S) = \mathbf{L}\mathbf{L}^*$.

In contrast to the procedure in the previous section, REJECTIONSAMPLESUBMATRIX only makes b proposals in total, and the set $S \subseteq S'$ of accepted pivots has random cardinality $1 \leq |S| \leq b$. The first pivot s'_1 is always accepted because the proposal distribution $\mathbf{u} = \text{diag}(\mathbf{A})$. Conditional on the number of acceptances $|S|$, the accepted pivots $s_1, \dots, s_{|S|}$ have the same distribution as $|S|$ pivots selected by the standard RPCHOLESKY algorithm (Algorithm D.1) applied to the full matrix \mathbf{A} .

2.3. Accelerated RPCholesky. The shortcoming of the block rejection sampling approach in the previous section is that the acceptance rate declines as the

Algorithm 2.1. REJECTIONSAMPLESUBMATRIX subroutine.

Input: Psd matrix $\mathbf{H} \in \mathbb{C}^{b \times b}$ and proposed pivots $S' = \{s'_1, \dots, s'_b\}$
Output: Accepted pivots $S \subseteq S'$; Cholesky factor \mathbf{L}

```

 $\mathbf{L} \leftarrow \mathbf{0}_{b \times b}$ ,  $\mathsf{T} \leftarrow \emptyset$ 
 $\mathbf{u} \leftarrow \text{diag}(\mathbf{H})$ 
for  $i = 1, \dots, b$  do
  if  $u_i \cdot \text{RAND} < h_{ii}$  then ▷ Accept or reject
     $\mathsf{T} \leftarrow \mathsf{T} \cup \{i\}$  ▷ If accept, induct pivot
     $\mathbf{L}(i : b, i) \leftarrow \mathbf{H}(i : b, i) / \sqrt{\mathbf{H}(i, i)}$  ▷ Compute Cholesky factor
     $\mathbf{H}(i + 1 : b, i + 1 : b) \leftarrow \mathbf{H}(i + 1 : b, i + 1 : b) - \mathbf{L}(i + 1 : b, i)\mathbf{L}(i + 1 : b, i)^*$ 
  end if
end for
 $\mathbf{L} \leftarrow \mathbf{L}(\mathsf{T}, \mathsf{T})$  ▷ Delete rejected rows and columns
 $S \leftarrow \{s'_i : i \in \mathsf{T}\}$  ▷ Get accepted pivots

```

Algorithm 2.2. Accelerated RPCHOLESKY: Standard implementation.

Input: Psd matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$; block size b ; number of rounds t
Output: Matrix $\mathbf{F} \in \mathbb{C}^{N \times k}$ defining low-rank approximation $\widehat{\mathbf{A}} = \mathbf{F}\mathbf{F}^*$; pivot set S

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

```

for  $i = 0$  to  $t - 1$  do
  Step 1: Propose a block of pivots
  Sample  $s'_{ib+1}, \dots, s'_{(i+1)b} \stackrel{\text{iid}}{\sim} \mathbf{u}$  and set  $S'_i \leftarrow \{s'_{ib+1}, \dots, s'_{(i+1)b}\}$ 
  Step 2: Rejection sampling
   $\mathbf{H} \leftarrow \mathbf{A}(S'_i, S'_i) - \mathbf{F}(S'_i, :)\mathbf{F}(S'_i, :)^*$ 
   $S_i, \mathbf{L} \leftarrow \text{REJECTIONSAMPLESUBMATRIX}(S'_i, \mathbf{H})$ 
   $S \leftarrow S \cup S_i$  ▷ Update pivots
  Step 3: Update low-rank approximation and proposal distribution
   $\mathbf{G} \leftarrow (\mathbf{A}(:, S_i) - \mathbf{F}\mathbf{F}(S_i, :)^*)\mathbf{L}^{-*}$ 
   $\mathbf{F} \leftarrow [\mathbf{F} \ \mathbf{G}]$ 
   $\mathbf{u} \leftarrow \mathbf{u} - \text{SQUAREDROWNORMS}\mathbf{G}$  ▷ Update diagonal
   $\mathbf{u} \leftarrow \max\{\mathbf{u}, \mathbf{0}\}$  ▷ Helpful in floating point arithmetic
end for

```

diagonal of the residual matrix decreases. Accelerated RPCHOLESKY remedies this weakness by employing a multiround sampling procedure. At the beginning of each round, the sampling distribution $\mathbf{u} \leftarrow \text{diag}(\mathbf{A} - \widehat{\mathbf{A}})$ is updated using the current low-rank approximation $\widehat{\mathbf{A}}$.

In detail, accelerated RPCHOLESKY works as follows. Fix a number of rounds $t \geq 1$ and block size $b \geq 1$, and initialize $\widehat{\mathbf{A}} \leftarrow \mathbf{0}$ and $\mathbf{u} \leftarrow \text{diag}(\mathbf{A})$. For $i = 0, 1, \dots, t - 1$, perform the following three steps:

1. **Propose a block of pivots.** Sample b pivots independently with replacement according to the distribution $\mathbb{P}\{s = j\} = \mathbf{u}(j) / \sum_{k=1}^n \mathbf{u}(k)$.
2. **Rejection sampling.** Use REJECTIONSAMPLESUBMATRIX (Algorithm 2.1) to downsample the proposed pivots.
3. **Update.** Update the low-rank approximation $\widehat{\mathbf{A}}$ and the diagonal vector $\mathbf{u} \leftarrow \text{diag}(\mathbf{A} - \widehat{\mathbf{A}})$ to reflect the newly accepted pivots.

The optimized pseudocode for accelerated RPCHOLESKY is given in Algorithm 2.2. As output, Algorithm 2.2 returns an approximation $\widehat{\mathbf{A}} \approx \mathbf{A}$ of rank $t \leq k \leq bt$, where k is the total number of accepted pivots. The rank k is at least t since the first pivot of each round is always accepted. For convenience and computational efficiency, $\widehat{\mathbf{A}}$ is reported in factored form $\widehat{\mathbf{A}} = \mathbf{F}\mathbf{F}^*$.

The computational timing of Algorithm 2.2 can be summarized as follows. Step 1 makes a negligible contribution to the overall timing, requiring $\mathcal{O}(N + b)$ operations. Step 2 is also relatively fast since it requires just $\mathcal{O}((b+k)b^2)$ operations, independent of the dimension N of the matrix. Step 3 is typically the slowest, as it generates and processes a large block of $|\mathbf{S}_i| \leq b$ columns in $\mathcal{O}(k|\mathbf{S}_i|N)$ operations. However, the user should take care that the acceptance rate does not fall below the level $\mathcal{O}(b/N)$ because then step 2 also becomes a computational bottleneck.

2.4. Low-memory implementation. As a limitation, Algorithm 2.2 requires storing an $N \times k$ factor matrix, which is prohibitively expensive when the approximation rank and the matrix dimensions are very large, say, $k \geq 10^4$ and $N \geq 10^6$. For a problem of this particular size, storing the factor matrix requires ≥ 80 GB of memory in double-precision floating-point format, which exceeds the available memory on many machines.

Algorithm 2.3 addresses the memory bottleneck with an alternative implementation of accelerated RPCHOLESKY that only requires $\mathcal{O}(N + k^2)$ storage and repeatedly regenerates entries of \mathbf{A} on an as-needed basis. The method produces the same distribution of pivots and acceptance probabilities as Algorithm 2.2, and the computational cost is still $\mathcal{O}(Nk^2)$ operations. However, the number of kernel matrix evaluations

Algorithm 2.3. Accelerated RPCHOLESKY: Low-memory implementation.

Input: Psd matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$; block size b ; number of rounds t

Output: Pivot set \mathbf{S} defining low-rank approximation $\widehat{\mathbf{A}} = \mathbf{A}(:, \mathbf{S})\mathbf{A}(\mathbf{S}, \mathbf{S})^\dagger \mathbf{A}(\mathbf{S}, :)$;

Cholesky factor \mathbf{L} with $\mathbf{A}(\mathbf{S}, \mathbf{S}) = \mathbf{L}\mathbf{L}^*$

Initialize $k \leftarrow t \cdot b$, $\mathbf{L} \leftarrow \mathbf{0}_{0 \times 0}$, $\mathbf{S} \leftarrow \emptyset$, and $\mathbf{u} \leftarrow \text{diag } \mathbf{A}$

for $i = 0$ to $t - 1$ **do**

Step 1: Propose a block of pivots

Sample $s'_{ib+1}, \dots, s'_{(i+1)b} \stackrel{\text{iid}}{\sim} \mathbf{u}$ and set $\mathbf{S}'_i \leftarrow \{s'_{ib+1}, \dots, s'_{(i+1)b}\}$

Step 2: Rejection sampling

$\mathbf{H} \leftarrow \mathbf{A}(\mathbf{S}'_i, \mathbf{S}'_i) - \mathbf{A}(\mathbf{S}'_i, \mathbf{S})\mathbf{L}^{-*}\mathbf{L}^{-1}\mathbf{A}(\mathbf{S}'_i, \mathbf{S})$

$\mathbf{S}_i, \mathbf{L}_i \leftarrow \text{REJECTIONSAMPLESUBMATRIX}(\mathbf{H}, \mathbf{S}'_i)$

$\mathbf{S} \leftarrow \mathbf{S} \cup \mathbf{S}_i$

▷ Update pivots

Step 3: Update low-rank approximation and proposal distribution

for $j = 0$ to $t - 1$ **do**

▷ Update diagonal, k entries at a time

$\mathbf{R} \leftarrow \{jb + 1, \dots, (j + 1)b\}$

▷ Chunk of b rows

Evaluate the submatrix $\mathbf{A}(\mathbf{R}, \mathbf{S})$

$\mathbf{G} \leftarrow (\mathbf{A}(\mathbf{R}, \mathbf{S}_i) - \mathbf{A}(\mathbf{R}, \mathbf{S} \setminus \mathbf{S}_i)\mathbf{L}^{-*}\mathbf{L}^{-1}\mathbf{A}(\mathbf{S} \setminus \mathbf{S}_i, \mathbf{S}_i))\mathbf{L}_i^{-*}$

$\mathbf{u}(\mathbf{R}) \leftarrow \mathbf{u}(\mathbf{R}) - \text{SQUAREDROWNORMS}(\mathbf{G})$

end for

$\mathbf{L} \leftarrow \begin{bmatrix} \mathbf{L} & \mathbf{0} \\ \mathbf{A}(\mathbf{S}_i, \mathbf{S} \setminus \mathbf{S}_i)\mathbf{L}_i^{-*} & \mathbf{L}_i \end{bmatrix}$

▷ Update Cholesky factor

$\mathbf{u} \leftarrow \max\{\mathbf{u}, \mathbf{0}\}$

▷ Ensure \mathbf{u} remains nonnegative

end for

increases from $\mathcal{O}(Nk)$ to $\mathcal{O}(Nk^2)$ because the kernel matrix entries are regenerated when they are needed.

The low-memory implementation in Algorithm 2.3 does not explicitly store an $N \times k$ factor matrix. Rather, it maintains a Cholesky factorization $\mathbf{A}(\mathbf{S}, \mathbf{S}) = \mathbf{L}\mathbf{L}^* \in \mathbb{C}^{k \times k}$ that implicitly defines the low-rank approximation:

$$(2.1) \quad \hat{\mathbf{A}} = \mathbf{A}(:, \mathbf{S})\mathbf{A}(\mathbf{S}, \mathbf{S})^{-1}\mathbf{A}(\mathbf{S}, :) = \mathbf{A}(:, \mathbf{S})\mathbf{L}^{-*}\mathbf{L}^{-1}\mathbf{A}(\mathbf{S}, :).$$

The submatrix $\mathbf{A}(:, \mathbf{S})$ is not stored, and it is regenerated as needed. The implicit low-rank approximation $\hat{\mathbf{A}}$ can be accessed via matrix–vector products, and it can be employed, for example, to speed up kernel machine learning algorithms [11].

2.5. Related work: Robust blockwise random pivoting. In recent work [12], Dong, Chen, Martinsson, and Pearce introduced robust blockwise random pivoting, a variant of block randomly pivoted QR for rectangular matrix low-rank approximation; this method can also be extended to provide a partial pivoted Cholesky decomposition for psd low-rank approximation. Similar to accelerated RPCHOLESKY, RBRP works by first sampling a block of random pivots and downsampling them to a smaller set of representatives. A comparison of our method with their approach appears in section C.

3. Numerical results. This section presents numerical experiments highlighting the speed and accuracy of accelerated RPCHOLESKY, both for a testbed of 125 kernel matrices (subsection 3.1) and for 8 potential energy calculations from molecular chemistry (subsection 3.2). The code to replicate the experiments can be found in the `block.experiments` folder of <https://github.com/epperly/Randomly-Pivoted-Cholesky>. In all these experiments, a target rank k is set, and block RPCHOLESKY and accelerated RPCHOLESKY are run for as many rounds t as needed to select k pivots. The full-memory implementation of accelerated RPCHOLESKY (Algorithm 2.2) is used in all experiments.

3.1. Testbed of 125 kernel matrices. The first round of experiments evaluates the accuracy and speed of accelerated RPCHOLESKY using over 100 kernel matrices, although some matrices were constructed from the same data with different kernel functions κ . The target matrices were generated from the following synthetic, random, and real-world data sets:

- **Challenging synthetic examples.** Three synthetic data were constructed to provide challenging examples of outliers and imbalanced classes of data points. The first two data sets depict a smile and a spiral in \mathbb{R}^2 , scaled to three different length scales. The third data set represents a Gaussian point cloud in \mathbb{R}^{20} , with 50 or 500 or 5000 points perturbed to become outliers. These synthetic data sets include $N = 10^5$ points, and the kernel function is Gaussian with bandwidth $\sigma = 1$.
- **Random point clouds.** Four random data sets were simulated, containing $N = 10^5$ standard Gaussian data points in dimension $d \in \{2, 10, 100, 1000\}$. The kernel function is Matérn-3/2 or ℓ_1 -Laplace with bandwidth $\sigma \in \{\sqrt{d}/2, \sqrt{d}, 2\sqrt{d}\}$ or Gaussian with bandwidth $\sigma \in \{\sqrt{d}/20, \sqrt{d}/10, \sqrt{d}/5\}$.
- **Real-world data.** 28 real-world data sets were downloaded from LIBSVM, OpenML, and UCI. They include the 20 examples from [11, Tab. 1], supplemented with `a9a`, `cadata`, `Click_prediction_small`, `phishing`, `MiniBoonNE`, `santander`, `skin_nonskin`, and `SUSY`. These data sets range in size from $N = 4.3 \times 10^4$ to $N = 1.3 \times 10^6$ points, but the larger data sets were randomly

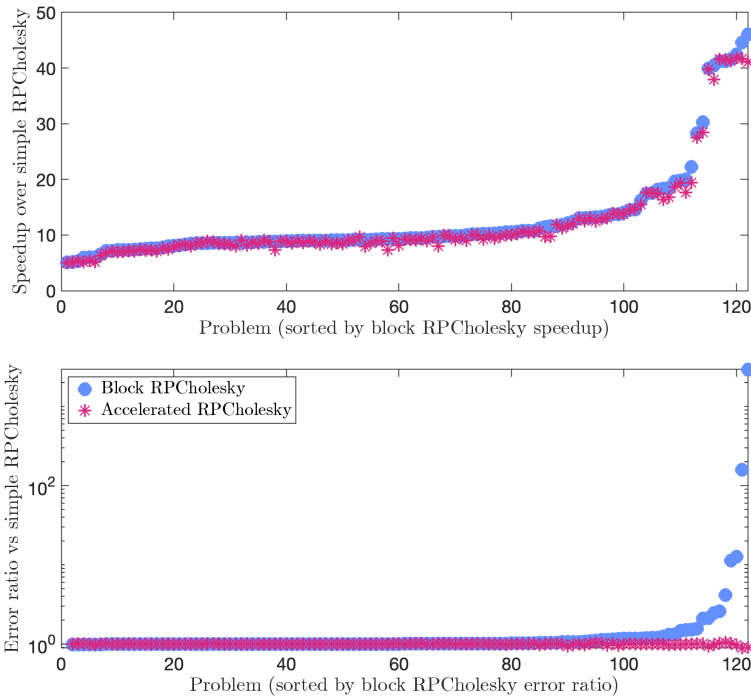


FIG. 3. Speed-up factor (top) and trace error ratio (eq. (3.1), bottom) of block and accelerated RPCHOLESKY compared to simple RPCHOLESKY on the testbed of 125 kernel matrices. Examples are sorted by the block RPCHOLESKY speed-up factor (top) or the block RPCHOLESKY error ratio (bottom).

subsampled to a maximum size of $N = 10^5$ points. The kernel function is either Gaussian, Matérn-3/2, or ℓ_1 -Laplace. The bandwidth is the median pairwise distance (ℓ_2 or ℓ_1 , as appropriate) from a random subsample of 10^3 points.

For each kernel matrix, the approximation rank is $k = 1000$ and the block size is $b = 150$. A small number of examples were discarded which possess a rank- k approximation of relative trace error $< 10^{-13}$; after this filtering process, there were 125 examples.

Results are shown in Figure 3. The top panel shows the speed-up factor of block and accelerated RPCHOLESKY over simple RPCHOLESKY. The speed-up factor is at least $5\times$ for all kernel matrices, and it reaches a level of $40\times$ to $50\times$ for 5% of the matrices. Consistent with Figure 1, the largest speedups are for high-dimensional data sets (real or synthetic) with the Gaussian or Matérn kernels. On most examples, the acceptance rate for accelerated RPCHOLESKY remains at or above 50%. However, for a few problems, like COMET_MC_SAMPLE with a Gaussian kernel, the acceptance rate for accelerated RPCHOLESKY drops as low as 3%. As might be expected, the examples where accelerated RPCHOLESKY has a high rejection rate coincide with the problems where block RPCHOLESKY is much less accurate than accelerated RPCHOLESKY, demonstrating the importance of filtering the pivot set for these examples.

The bottom panel shows the ratio of trace norm errors

$$(3.1) \quad \frac{\text{tr}(\mathbf{A} - \widehat{\mathbf{A}})}{\text{tr}(\mathbf{A} - \widehat{\mathbf{A}}_{\text{RPC}})},$$

where $\hat{\mathbf{A}}$ is computed by block or accelerated RPCHOLESKY and $\hat{\mathbf{A}}_{\text{RPC}}$ is produced by simple RPCHOLESKY. As expected, accelerated RPCHOLESKY and RPCHOLESKY have the same approximation quality. For 100 out of 125 kernel matrices, block RPCHOLESKY leads to $1.00\times$ to $1.16\times$ the simple RPCHOLESKY error. For the remaining matrices, block RPCHOLESKY produces a significantly worse approximation than simple RPCHOLESKY; the error is $3000\times$ worse for the most extreme problem.

3.2. Potential energy surfaces in molecular chemistry. Molecular dynamics simulations rely on a fast-to-compute model for the potential energy of an atomic system. Recent work has explored fitting potential energies with machine learning methods such as kernel ridge regression (KRR) [5, 6, 7, 33]. However, despite the ubiquity of KRR in molecular chemistry, there is a surprising lack of consensus over the best numerical method for KRR optimization.

Responding to this uncertainty, the present paper demonstrates how *low-rank preconditioners* can accelerate KRR for potential energy modeling and it shows that accelerated RPCHOLESKY is the best available low-rank approximation method for this use case. These conclusions are supported by experiments on a benchmark data set of eight molecules: aspirin, benzene, ethanol, malonaldehyde, naphthalene, salicylic acid, toluene, and uracil.

The experiments in this section use density functional theory data released by Chmiela et al. [6], available at sgdml.org. This data was prepared as follows:

- *Data size:* The measurements for each molecule were randomly subsampled to $N = 10^5$ training points and $N_{\text{test}} = 10^4$ testing points.
- *Input data features:* The input data consists of positions $\mathbf{r}_i \in \mathbb{R}^3$ for each atom $i = 1, \dots, n_{\text{atoms}}$. Since the potential energy is invariant under rigid motions, the data is transformed into a vector \mathbf{x} listing the inverse pairwise distances $1/\|\mathbf{r}_i - \mathbf{r}_j\|$ for $1 \leq i < j \leq n_{\text{atoms}}$ [6]. Following the usual procedure, these features \mathbf{x} were standardized by subtracting the mean and dividing by the standard deviation. The dimensionality of the input data is thus $d = \binom{n_{\text{atoms}}}{2} \in [36, 210]$, depending on the molecule.
- *Output data features:* The output data is potential energy measurements in units of kcal mol^{-1} . The potential energies were centered by subtracting the mean, but they were *not* divided by the standard deviation. The standard deviations are thus $2.3\text{--}6.0 \text{ kcal mol}^{-1}$, depending on the molecule.

This paper employs a standard KRR model [23], in which the potential energy for each molecule is approximated by a linear combination of kernel functions,

$$(3.2) \quad f(\mathbf{x}) = \sum_{i=1}^N \beta_i \kappa(\mathbf{x}_i, \mathbf{x}),$$

that are centered on the training data points $(\mathbf{x}_i)_{1 \leq i \leq N}$. The coefficients $\beta_i \in \mathbb{R}$ are chosen as the solution to the linear system:

$$(3.3) \quad (\mathbf{A} + \mu \mathbf{I})\boldsymbol{\beta} = \mathbf{y}.$$

Here, $\mathbf{y} \in \mathbb{R}^n$ is the vector of potential energy output values, $\mathbf{A} = (\kappa(\mathbf{x}_i, \mathbf{x}_j))_{1 \leq i, j \leq N}$ is the kernel matrix, and $\mu > 0$ is the regularization parameter. Following [6], the present study chooses κ to be a Matérn-5/2 kernel

$$\kappa(\mathbf{x}, \mathbf{x}') = \left(1 + \sqrt{5} \cdot \frac{\|\mathbf{x} - \mathbf{x}'\|}{\sigma} + \frac{5}{3} \cdot \frac{\|\mathbf{x} - \mathbf{x}'\|^2}{\sigma^2} \right) \exp \left(-\sqrt{5} \cdot \frac{\|\mathbf{x} - \mathbf{x}'\|}{\sigma} \right)$$

with bandwidth $\sigma = \sqrt{d}$. The regularization parameter is set to $\mu = 10^{-9}N$.




FIG. 4. Runtime needed to generate the preconditioner (prep time) and run the necessary number of CG iterations until $\|\mathbf{A}\boldsymbol{\beta}^{(t)} - \mathbf{y}\|/\|\mathbf{y}\| < 10^{-3}$ (PCG time). Each calculation is performed once and the results are averaged over the eight molecules.

For large N , the standard approach for solving (3.3) in both the scientific [5, 6, 7, 33] and mathematical [1, 3, 8, 11, 17, 18, 34] literatures is preconditioned conjugate gradient (PCG) [20, Alg. 11.5.1]. PCG is an iterative algorithm, and each iteration requires a matrix–vector product with $\mathbf{A} + \mu\mathbf{I}$ and a linear solve with a preconditioner matrix \mathbf{P} . To control the number of iterations and the total runtime, it is essential to find an easy-to-invert preconditioner satisfying $\mathbf{P} \approx \mathbf{A} + \mu\mathbf{I}$. Recent work has promoted the use of *low-rank preconditioners*, which take the form

$$(3.4) \quad \mathbf{P} = \mathbf{F}\mathbf{F}^* + \mu\mathbf{I} \quad \text{for } \mathbf{F} \in \mathbb{R}^{N \times k},$$

where $\mathbf{F}\mathbf{F}^* \approx \mathbf{A}$ is a rank- k approximation. The community has vigorously debated the best algorithm to compute \mathbf{F} in practice and in theory [1, 3, 8, 11, 17, 18, 27, 30, 34].

Figure 4 charts the average runtime needed to perform potential energy prediction with a low-rank preconditioner generated by accelerated RPCHOLESKY, simple RPCHOLESKY, ridge leverage score sampling (using the RRLS algorithm [27]), uniform Nyström approximation [8, 17], and greedy Nyström approximation [18, 34]. The results show that accelerated RPCHOLESKY achieves the fastest runtimes, improving on the other low-rank approximation algorithms by 4% to 17%. Accelerated RPCHOLESKY has the twin advantages that it quickly generates a preconditioner (fast prep time), and it leads to the minimal number of PCG iterations (fast PCG time). Another important feature is the consistent performance across problems: accelerated RPCHOLESKY always produces near-optimal low-rank approximations, whereas other methods can struggle on difficult examples; see Table 1.

The potential energy model in this paper is a direct application of standard KRR, yet it leads to higher accuracies for four of eight molecules than the more complicated kernel model in [6]. On the most extreme example (uracil), the simple model is $10\times$ more accurate than the existing model. See Table 2 for a table including all eight molecules. These experiments demonstrate that simple kernel methods combined with scalable algorithms based on accelerated RPCHOLESKY can be competitive for problems in scientific machine learning.

4. Theoretical results. This section compares accelerated RPCHOLESKY with simple RPCHOLESKY using rigorous error analysis. On the surface, the comparison is straightforward. Since accelerated RPCHOLESKY generates the same distribution

TABLE 1

Relative error $\text{tr}(\mathbf{A} - \mathbf{F}\mathbf{F}^*) / \text{tr}(\mathbf{A})$ of the low-rank approximation for each kernel matrix. Accelerated and simple RPCHOLESKY have the same error up to ± 0.001 random variations. Best results are indicated in bold for each molecule.

Molecule	Accelerated RPC	Simple RPC	Greedy	Uniform	Ridge leverage scores
Aspirin	0.091	0.090	0.085	0.097	0.099
Benzene	0.064	0.064	0.069	0.067	0.070
Ethanol	0.078	0.078	0.087	0.079	0.080
Malonaldehyde	0.073	0.073	0.084	0.074	0.075
Naphthalene	0.106	0.106	0.108	0.109	0.111
Salicylic acid	0.093	0.093	0.095	0.096	0.098
Toluene	0.104	0.103	0.109	0.106	0.107
Uracil	0.056	0.056	0.051	0.062	0.066
Average (8 molecules)	0.083	0.083	0.086	0.086	0.088

TABLE 2

Mean absolute error (kcal / mol) for the existing potential energy model [6] versus the new model (3.2). Best results are indicated in bold for each molecule.

Molecule	Formula	Old energy error	New energy error
Aspirin	C ₉ H ₈ O ₄	0.127	0.040
Benzene	C ₆ H ₆	0.069	0.048
Ethanol	C ₂ H ₆ O	0.052	0.089
Malonaldehyde	C ₃ H ₄ O ₂	0.074	0.133
Naphthalene	C ₁₀ H ₈	0.113	0.126
Salicylic acid	C ₇ H ₆ O ₃	0.105	0.102
Toluene	C ₆ H ₅ CH ₃	0.092	0.146
Uracil	C ₄ H ₄ N ₂ O ₂	0.103	0.010

of pivots as simple RPCHOLESKY, it satisfies the main RPCHOLESKY error bound [4, Thm. 5.1].

THEOREM 4.1 (sufficient pivots for simple and accelerated RPCHOLESKY [4]). Consider a psd matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$, and let $\mathbf{A}^{(k)}$ denote the random residual after applying simple or accelerated RPCHOLESKY with k random pivots. Then

$$\mathbb{E}[\text{tr}(\mathbf{A}^{(k)})] \leq (1 + \varepsilon) \cdot \sum_{i=r+1}^N \lambda_i(\mathbf{A})$$

as soon as

$$k \geq \frac{r}{\varepsilon} + r \log \left(\frac{1}{\varepsilon \eta} \right), \quad \text{where } \eta = \frac{\sum_{i=r+1}^N \lambda_i(\mathbf{A})}{\sum_{i=1}^r \lambda_i(\mathbf{A})}.$$

Theorem 4.1 bounds the sufficient number of pivots in accelerated RPCHOLESKY, yet it does not bound the sufficient number of iterations needed to produce the pivots via rejection sampling. The following new result completes the theoretical understanding of accelerated RPCHOLESKY by bounding the number of iterations. It also provides error bounds for block RPCHOLESKY. This new result will be proved in subsections 4.3 to 4.6.

THEOREM 4.2 (sufficient iterations for accelerated and block RPCHOLESKY). Consider a psd matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$, and let $\mathbf{A}^{(t)}$ denote the random residual after applying t rounds of accelerated RPCHOLESKY or block RPCHOLESKY with a block size $b \geq 1$. Then

$$\mathbb{E}[\text{tr}(\mathbf{A}^{(t)})] \leq (1 + \varepsilon) \cdot \sum_{i=r+1}^N \lambda_i(\mathbf{A})$$

as soon as the number of proposed pivots bt satisfies

$$bt \geq \frac{r}{\varepsilon} + (r + b) \log\left(\frac{1}{\varepsilon\eta}\right), \quad \text{where } \eta = \frac{\sum_{i=r+1}^N \lambda_i(\mathbf{A})}{\sum_{i=1}^r \lambda_i(\mathbf{A})}.$$

The comparison between accelerated and block RPCHOLESKY follows from a coupling argument in which block RPCHOLESKY accepts a greater number of pivots than accelerated RPCHOLESKY in each round, leading to a higher approximation quality. However, in this comparison, accelerated RPCHOLESKY has the benefit of rejecting many of the proposed pivots and therefore producing a similarly accurate approximation with a smaller approximation rank. When the methods are run with the same approximation rank, accelerated RPCHOLESKY can be much more accurate, as reflected in the numerical experiments of section 3. The failure modes of block RPCHOLESKY can be seen in Figures 2 and 3.

The bounds in Theorems 4.1 and 4.2 hold for all accuracy levels $\varepsilon > 0$, and they each contain a term r/ε , which is the minimal number of pivots needed to guarantee an error level $(1 + \varepsilon) \cdot \sum_{i=r+1}^N \lambda_i(\mathbf{A})$ in any low-rank approximation produced by a pivoted partial Cholesky decomposition applied to a worst-case matrix [4, Thm. C.1]. The bounds also contain a term $r \log(1/(\varepsilon\eta))$ or $(r + b) \log(1/(\varepsilon\eta))$, which represents the number of pivots that need to be accepted or proposed for the algorithm to improve the approximation quality from an initial error level $\text{tr}(\mathbf{A}) = \eta^{-1} \cdot \sum_{i=r+1}^N \lambda_i(\mathbf{A})$ to a lower error level $(1 + \varepsilon) \cdot \sum_{i=r+1}^N \lambda_i(\mathbf{A})$. Indeed, the only difference between Theorems 4.1 and 4.2 is the quantity $b \log(1/(\varepsilon\eta))$, present in the latter. This difference becomes significant when the theorems are evaluated for a small target rank $r < b$ and a tiny approximation error $\eta \ll 1$. This extreme setting occurs when the matrix \mathbf{A} is nearly rank- r . Then, accelerated RPCHOLESKY accepts very few proposals per round, so it takes many rounds to converge to high accuracy.

Last, Theorem 4.2 can be compared with the bound of Deshpande, Rademacher, Vempala, and Wang [9]. Their bound was originally derived for block randomly pivoted QR, but it can be extended to block RPCHOLESKY. Here is a simplified version of their result, adapted from [4, sec. 3.6].

THEOREM 4.3 (sufficient iterations for block RPCHOLESKY with a large block size [9]). *Fix $\varepsilon \leq 3/2$ and a large block size $b \geq 3r/\varepsilon$. Consider a psd matrix \mathbf{A} , and let $\mathbf{A}^{(t)}$ denote the random residual after applying t rounds of block RPCHOLESKY. Then*

$$\mathbb{E}[\text{tr}(\mathbf{A}^{(t)})] \leq (1 + \varepsilon) \cdot \sum_{i=r+1}^N \lambda_i(\mathbf{A})$$

as soon as

$$t \geq 1 + \frac{\log(1/\eta + 1)}{\log(3/\varepsilon)}, \quad \text{where } \eta = \frac{\sum_{i=r+1}^N \lambda_i(\mathbf{A})}{\sum_{i=1}^r \lambda_i(\mathbf{A})}.$$

When Theorems 4.2 and 4.3 are both applicable, they specify a similar number of sufficient iterations for block RPCHOLESKY to converge to high accuracy. However, Theorem 4.3 has a more limited range of applicability since it requires a large block

size $b \geq 3r/\varepsilon$. This prescription for a large block size is exactly the reverse of block RPCHOLESKY's behavior in practice, in which the method provides more accurate approximations given a smaller block size. In particular, Theorem 4.3 has limited applicability in the typical setting where the user tunes the block size b for a given computing setup to achieve the fastest runtime. The final weakness of Theorem 4.3 is that it does not extend to accelerated RPCHOLESKY.

4.1. Method of proof. The proof uses many probabilistic and linear algebraic tools that were introduced in the analysis of simple RPCHOLESKY [4, Thm. 5.1], and it follows the same core approach:

- (a) Evaluate the function that gives the expected value of the residual matrix after generating and processing a single pivot.
- (b) Prove monotonicity and concavity properties for the expected residual function.
- (c) Use monotonicity and concavity to identify a worst-case matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$.
- (d) Analyze the worst-case matrix to produce error bounds.

As the core difficulty, accelerated RPCHOLESKY leads to an expected residual function that does not satisfy the matrix monotonicity property. Therefore, it is necessary to develop more subtle concavity and monotonicity properties, and the proof is twice as long as before.

The expected residual function for accelerated RPCHOLESKY can be derived as follows. Let $\widehat{\mathbf{A}}^{(1,i)}$ be the low-rank approximation generated from the first $i \leq b$ proposals, and let $\mathbf{A}^{(1,i)} = \mathbf{A} - \widehat{\mathbf{A}}^{(1,i)}$ be the corresponding residual matrix. Accelerated RPCHOLESKY updates the residual matrix in three steps:

1. A pivot s'_{i+1} is sampled according to $\mathbb{P}\{s'_{i+1} = j\} = \mathbf{A}(j, j) / \text{tr} \mathbf{A}$.
2. The pivot is accepted with probability $\mathbf{A}^{(1,i)}(j, j) / \mathbf{A}(j, j)$.
3. Conditional on acceptance, the residual matrix is updated:

$$\mathbf{A}^{(1,i+1)} = \mathbf{A}^{(1,i)} - \mathbf{A}^{(1,i)}(:, j) \mathbf{A}^{(1,i)}(j, :) / \mathbf{A}^{(1,i)}(j, j).$$

By combining steps 1–3 above, the expected residual matrix is

$$\begin{aligned} \mathbb{E}[\mathbf{A}^{(1,i+1)} | \mathbf{A}^{(1,i)}] &= \mathbf{A}^{(1,i)} - \sum_{j=1}^N \frac{\mathbf{A}(j, j)}{\text{tr}(\mathbf{A})} \cdot \frac{\mathbf{A}^{(1,i)}(j, j)}{\mathbf{A}(j, j)} \cdot \frac{\mathbf{A}^{(1,i)}(:, j) \mathbf{A}^{(1,i)}(j, :)}{\mathbf{A}^{(1,i)}(j, j)} \\ &= \mathbf{A}^{(1,i)} - \frac{1}{\text{tr}(\mathbf{A})} (\mathbf{A}^{(1,i)})^2. \end{aligned}$$

Therefore, the expected residual function can be defined as

$$\phi_\alpha(x) = x - \frac{1}{\alpha} x^2, \quad \text{where } \alpha = \text{tr}(\mathbf{A}) \text{ and } x = \mathbf{A}^{(1,i)}.$$

This function depends on a scalar parameter $\alpha > 0$, and it can be evaluated for inputs x that are scalars or matrices. By repeated function composition, this function bounds the expected residual matrix of accelerated RPCHOLESKY after any number of steps.

Bounding the error in this way requires an expected residual function that is concave and monotone. However, the function $\phi_\alpha(x)$ barely has enough monotonicity for the argument to work. In particular, the mapping $x \mapsto \phi_\alpha(x)$ is only monotone for scalars $x \in [0, \alpha/2]$, and it fails to be monotone for large classes of psd matrices.

The analysis pushes through the technical difficulties and produces an upper bound for accelerated RPCHOLESKY depending on the expected residual function. The bound becomes worse when the eigenvalues are averaged together or increased. Hence, the worst-case matrix takes the form

$$\mathbf{A} = \text{diag} \left\{ \underbrace{\frac{\gamma}{r}, \dots, \frac{\gamma}{r}}_{r \text{ times}}, \underbrace{\frac{\beta}{N-r}, \dots, \frac{\beta}{N-r}}_{N-r \text{ times}} \right\}, \quad \text{where} \quad \begin{cases} \gamma = \sum_{i=1}^r \lambda_i(\mathbf{A}), \\ \beta = \sum_{i=r+1}^N \lambda_i(\mathbf{A}). \end{cases}$$

It is a two-cluster matrix with a small block of leading eigenvalues and a large block of trailing eigenvalues. The identification of this two-cluster matrix is not new: the same worst-case matrix was already encountered in the proof of [4, Thm. 5.1]. Applying the expected residual function to this matrix completes the proof of Theorem 4.2.

4.2. Organization of proof. Here is the overall organization for the proof of Theorem 4.2, which will be carried out over the next four subsections.

Part I: Properties of the expected residual function. The first part of the analysis (subsection 4.3) establishes concavity and monotonicity properties for the expected residual function. The concavity and monotonicity properties are more detailed than what was proved in [4, Thm. 5.1].

Part II: Accelerated RPCHOLESKY error. The second part of the analysis (subsection 4.4) establishes the following error bound.

LEMMA 4.4 (central error bound). *Fix a block size $b \geq 1$ and a target psd matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$. Let $\mathbf{A}^{(t)}$ denote the residual matrix after applying t rounds of accelerated or block RPCHOLESKY with blocks of b proposals. Also introduce the matrix-valued function*

$$\Phi_b(\mathbf{M}) = \underbrace{\phi_{\text{tr}(\mathbf{M})} \circ \phi_{\text{tr}(\mathbf{M})} \circ \dots \circ \phi_{\text{tr}(\mathbf{M})}}_{b \text{ times}}(\mathbf{M}), \quad \text{where} \quad \phi_\alpha(x) = x - \frac{1}{\alpha}x^2,$$

and let $\Phi_b^{(t)} = \Phi_b \circ \Phi_b \circ \dots \circ \Phi_b$ denote the t -fold composition of Φ_b . Then, the eigenvalues of $\mathbb{E}[\mathbf{A}^{(t)}]$ are bounded by the eigenvalues of $\Phi_b^{(t)}(\mathbf{A})$ as follows:

$$\lambda_i(\mathbb{E}[\mathbf{A}^{(t)}]) \leq \lambda_i(\Phi_b^{(t)}(\mathbf{A})) \quad \text{for each } i = 1, \dots, N.$$

The proof of Lemma 4.4 is the most intricate argument of the paper. This lemma shows that the expected trace norm error is bounded by a deterministic quantity:

$$\mathbb{E} \text{tr}(\mathbf{A}^{(t)}) \leq \text{tr}(\Phi_b^{(t)}(\mathbf{A})).$$

This upper bound depends only on the eigenvalues of \mathbf{A} , independent of the eigenvectors.

Part III: Permutation averaging lemma. The next key idea is to average together the leading eigenvalues of \mathbf{A} and also average together the trailing eigenvalues. Subsection 4.5 proves the following helpful lemma, which shows that the averaging can only make the error worse.

LEMMA 4.5 (permutation averaging). *For any diagonal, psd matrix $\mathbf{\Lambda} \in \mathbb{R}^{N \times N}$ and any random permutation $\mathbf{P} \in \mathbb{R}^{N \times N}$,*

$$\text{tr}(\Phi_b^{(t)}(\mathbf{\Lambda})) \leq \text{tr}(\Phi_b^{(t)}(\mathbb{E}[\mathbf{P}\mathbf{\Lambda}\mathbf{P}^*])).$$

Part IV: Main error bound. The last part of the analysis in subsection 4.6 examines the worst-case matrix \mathbf{A} and obtains a simple recursive formula for $\text{tr}(\Phi_b^{(t)}(\mathbf{A}))$ in terms of t . Bounding the recursion leads to the proof of the main result, Theorem 4.2.

4.3. Part I: Properties of the expected residual function. This section establishes concavity and monotonicity properties for the expected residual function. The first lemma considers the application of this function to scalars.

LEMMA 4.6 (expected residual properties). *Let $\phi_\alpha : [0, \alpha] \rightarrow [0, \alpha]$ be the function $\phi_\alpha(x) = x - \frac{1}{\alpha}x^2$, and let*

$$\phi_\alpha^{(b)} = \underbrace{\phi_\alpha \circ \phi_\alpha \circ \dots \circ \phi_\alpha}_{b \text{ times}}$$

denote the b -fold composition of ϕ_α . Then the following holds for each $b \geq 1$:

- (a) **Uniform boundedness:** $\phi_\alpha^{(b)} \leq \alpha/4$.
- (b) **Monotonicity in x on $[0, \alpha/2]$:** $\phi_\alpha^{(b)}(x) \leq \phi_\alpha^{(b)}(x')$ if $x \leq x' \leq \alpha/2$.
- (c) **Joint positive homogeneity:** $\phi_{\beta \cdot \alpha}^{(b)}(\beta \cdot x) = \beta \cdot \phi_\alpha^{(b)}(x)$ for $\beta > 0$.
- (d) **Joint concavity:** The mapping $(x, \alpha) \mapsto \phi_\alpha^{(b)}(x)$ is concave:

$$\theta \phi_{\alpha_1}^{(b)}(x_1) + (1 - \theta) \phi_{\alpha_2}^{(b)}(x_2) \leq \phi_{\theta \alpha_1 + (1 - \theta) \alpha_2}^{(b)}(\theta x_1 + (1 - \theta) x_2) \quad \text{for } \theta \in (0, 1).$$

- (e) **Upper bound:** $\phi_\alpha^{(b)}(x) \leq 1/(b\alpha^{-1} + x^{-1})$.

Proof. See section B. □

The next lemma shows that the expected residual function continues to exhibit concavity and monotonicity when it is applied to certain classes of matrices.

LEMMA 4.7 (expected residual properties for matrices). *Define*

$$\Phi_b(\mathbf{A}) = \underbrace{\phi_{\text{tr}(\mathbf{A})} \circ \phi_{\text{tr}(\mathbf{A})} \circ \dots \circ \phi_{\text{tr}(\mathbf{A})}}_{b \text{ times}}(\mathbf{A}), \quad \text{where } \phi_\alpha(x) = x - \frac{1}{\alpha}x^2.$$

Let $\Phi_b^{(t)} = \Phi_b \circ \Phi_b \circ \dots \circ \Phi_b$ denote the t -fold composition of Φ_b . Then the following holds for any $t \geq 1$:

- (a) $\Phi_b^{(t)}$ is nondecreasing on the space of diagonal, psd matrices:

$$\Phi_b^{(t)}(\mathbf{\Lambda}_1) \preceq \Phi_b^{(t)}(\mathbf{\Lambda}_2) \quad \text{if } \mathbf{\Lambda}_1 \preceq \mathbf{\Lambda}_2.$$

- (b) $\Phi_b^{(t)}$ is concave on the space of diagonal, psd matrices:

$$\theta \Phi_b^{(t)}(\mathbf{\Lambda}_1) + (1 - \theta) \Phi_b^{(t)}(\mathbf{\Lambda}_2) \preceq \Phi_b^{(t)}(\theta \mathbf{\Lambda}_1 + (1 - \theta) \mathbf{\Lambda}_2) \quad \text{for } \theta \in (0, 1).$$

- (c) Φ_1 is concave on the space of all psd matrices:

$$\theta \Phi_1(\mathbf{A}_1) + (1 - \theta) \Phi_1(\mathbf{A}_2) \preceq \Phi_1(\theta \mathbf{A}_1 + (1 - \theta) \mathbf{A}_2) \quad \text{for } \theta \in (0, 1).$$

Proof. See section B. □

4.4. Part II: Accelerated RPCholesky error. This section analyzes the error of accelerated RPCHOLESKY and proves Lemma 4.4. The technical crux is the following proposition, which bounds the error of accelerated RPCHOLESKY or block RPCHOLESKY applied to a random psd matrix.

PROPOSITION 4.8 (block and accelerated RPCHOLESKY on a random matrix). Consider a random psd matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$, and let $\mathbf{A}^{(1,b)}$ denote the residual matrix after applying one round of accelerated RPCHOLESKY or block RPCHOLESKY to \mathbf{A} with $b \geq 1$ proposals. Also introduce the matrix-valued function

$$\Phi_b(\mathbf{M}) = \underbrace{\phi_{\text{tr}(\mathbf{M})} \circ \phi_{\text{tr}(\mathbf{M})} \circ \dots \circ \phi_{\text{tr}(\mathbf{M})}}_{b \text{ times}}(\mathbf{M}), \quad \text{where } \phi_\alpha(x) = x - \frac{1}{\alpha}x^2.$$

Then the eigenvalues of $\mathbb{E}[\mathbf{A}^{(1,b)}]$ are bounded by the eigenvalues of $\Phi_b^{(1,b)}(\mathbb{E}[\mathbf{A}])$:

$$\lambda_i(\mathbb{E}[\mathbf{A}^{(1,b)}]) \leq \lambda_i(\Phi_b(\mathbb{E}[\mathbf{A}])) \quad \text{for each } i = 1, \dots, N.$$

The proof of this result relies on the min–max principle [37, Thm. 8.9].

LEMMA 4.9 (min–max principle). The eigenvalues of a Hermitian matrix $\mathbf{M} \in \mathbb{C}^{N \times N}$ satisfy

$$\lambda_i(\mathbf{M}) = \min_{\dim(\mathcal{S})=N-i+1} \max_{\mathbf{v} \in \mathcal{S}, \|\mathbf{v}\|=1} \mathbf{v}^* \mathbf{M} \mathbf{v},$$

where the maximum is taken over all $(N - i + 1)$ -dimensional linear subspaces \mathcal{S} . The optimal \mathcal{S} is spanned by the eigenvectors of \mathbf{M} with the $N - i + 1$ smallest eigenvalues.

Proof of Proposition 4.8. The proof has three steps.

Step 1: Comparison of block and accelerated RPCholesky. Accelerated RPCHOLESKY and block RPCHOLESKY yield the same distribution of pivot proposals, so the proposals $\mathbf{S}' = \{s'_1, \dots, s'_b\}$ are assumed to be identical. Block RPCHOLESKY then accepts all the proposals, leading to a residual matrix

$$\tilde{\mathbf{A}}^{(1,b)} = \mathbf{A} - \mathbf{A}(:, \mathbf{S}') \mathbf{A}(\mathbf{S}', \mathbf{S}')^\dagger \mathbf{A}(\mathbf{S}', :) = \mathbf{A}^{1/2}(\mathbf{I} - \Pi_{\mathbf{A}^{1/2}(:, \mathbf{S}')}) \mathbf{A}^{1/2},$$

where

$$\Pi_{\mathbf{A}^{1/2}(:, \mathbf{S}')} = \mathbf{A}^{1/2}(:, \mathbf{S}') \mathbf{A}(\mathbf{S}', \mathbf{S}')^\dagger \mathbf{A}^{1/2}(\mathbf{S}', :)$$

is the orthogonal projection onto the range of $\mathbf{A}^{1/2}(:, \mathbf{S}')$. In contrast, accelerated RPCHOLESKY accepts a subset of the proposals $\mathbf{S} \subseteq \mathbf{S}'$, leading to a residual matrix

$$\mathbf{A}^{(1,b)} = \mathbf{A}^{1/2}(\mathbf{I} - \Pi_{\mathbf{A}^{1/2}(:, \mathbf{S})}) \mathbf{A}^{1/2}.$$

The range of $\mathbf{A}^{1/2}(:, \mathbf{S}')$ is no smaller than the range of $\mathbf{A}^{1/2}(:, \mathbf{S})$, so

$$(4.1) \quad \mathbf{I} - \Pi_{\mathbf{A}^{1/2}(:, \mathbf{S}')} \preceq \mathbf{I} - \Pi_{\mathbf{A}^{1/2}(:, \mathbf{S})}.$$

Conjugate both sides of (4.1) with $\mathbf{A}^{1/2}$ to yield

$$\tilde{\mathbf{A}}^{(1,b)} = \mathbf{A}^{1/2}(\mathbf{I} - \Pi_{\mathbf{A}^{1/2}(:, \mathbf{S}')}) \mathbf{A}^{1/2} \preceq \mathbf{A}^{1/2}(\mathbf{I} - \Pi_{\mathbf{A}^{1/2}(:, \mathbf{S})}) \mathbf{A}^{1/2} = \mathbf{A}^{(1,b)}.$$

Taking expectations, it follows that $\mathbb{E}[\tilde{\mathbf{A}}^{(1,b)}] \preceq \mathbb{E}[\mathbf{A}^{(1,b)}]$ and, by Weyl’s inequality,

$$\lambda_i(\mathbb{E}[\tilde{\mathbf{A}}^{(1,b)}]) \leq \lambda_i(\mathbb{E}[\mathbf{A}^{(1,b)}]), \quad \text{for each } i = 1, \dots, N.$$

This completes the comparison of block and accelerated RPCHOLESKY.

Step 2: Bounding $\mathbf{v}^* \mathbb{E}[\mathbf{A}^{(1,b)}] \mathbf{v}$. Now it suffices to analyze the accelerated RPCHOLESKY error. This proof will do so by (a) deriving an upper bound for the quadratic form $\mathbf{v}^* \mathbb{E}[\mathbf{A}^{(1,b)}] \mathbf{v}$, (b) constructing a $(N - i + 1)$ -dimensional subspace where the Rayleigh quotients are no larger than $\lambda_i(\Phi_b(\mathbb{E}[\mathbf{A}]))$, and (c) invoking the min-max principle.

The residual for accelerated RPCHOLESKY after the first proposal has expectation

$$\mathbb{E}[\mathbf{A}^{(1,1)}] = \mathbb{E}[\mathbb{E}[\mathbf{A}^{(1,1)} | \mathbf{A}]] = \mathbb{E} \left[\mathbf{A} - \frac{1}{\text{tr}(\mathbf{A})} \mathbf{A}^2 \right] = \mathbb{E}[\Phi_1(\mathbf{A})].$$

Since Φ_1 is concave over psd matrices (Lemma 4.7c), Jensen's inequality implies

$$\mathbb{E}[\mathbf{A}^{(1,1)}] = \mathbb{E}[\Phi_1(\mathbf{A})] \preceq \Phi_1(\mathbb{E}[\mathbf{A}]).$$

Next fix a unit-length vector $\mathbf{v} \in \mathbb{C}^N$ and conjugate both sides of this display with \mathbf{v} :

$$\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,1)} \mathbf{v}] \leq \mathbf{v}^* \Phi_1(\mathbb{E}[\mathbf{A}]) \mathbf{v}.$$

By boundedness of $\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}$ and monotonicity of $\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}$ (Lemma 4.6a-b), it follows

$$(4.2) \quad \phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,1)} \mathbf{v}]) \leq \phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\mathbf{v}^* \Phi_1(\mathbb{E}[\mathbf{A}]) \mathbf{v}).$$

The residual matrix after the $(\ell + 1)$ st proposal has conditional expectation

$$\mathbb{E}[\mathbf{A}^{(1,\ell+1)} | \mathbf{A}^{(1,\ell)}, \mathbf{A}] = \mathbf{A}^{(1,\ell)} - \frac{1}{\text{tr}(\mathbf{A})} (\mathbf{A}^{(1,\ell)})^2.$$

Conjugate both sides of this display with \mathbf{v} to obtain:

$$\begin{aligned} \mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,\ell+1)} \mathbf{v} | \mathbf{A}^{(1,\ell)}, \mathbf{A}] &= \mathbf{v}^* \mathbf{A}^{(1,\ell)} \mathbf{v} - \frac{1}{\text{tr}(\mathbf{A})} \mathbf{v}^* (\mathbf{A}^{(1,\ell)})^2 \mathbf{v} \\ &\leq \mathbf{v}^* \mathbf{A}^{(1,\ell)} \mathbf{v} - \frac{1}{\text{tr}(\mathbf{A})} (\mathbf{v}^* \mathbf{A}^{(1,\ell)} \mathbf{v})^2 = \phi_{\text{tr}(\mathbf{A})}(\mathbf{v}^* \mathbf{A}^{(1,\ell)} \mathbf{v}). \end{aligned}$$

The inequality follows because $\mathbf{v} \mathbf{v}^*$ is a contraction. Take expectations and apply Jensen's inequality to the concave function $x \mapsto \phi_{\text{tr}(\mathbf{A})}(x)$ (Lemma 4.6d):

$$\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,\ell+1)} \mathbf{v} | \mathbf{A}] \leq \mathbb{E}[\phi_{\text{tr}(\mathbf{A})}(\mathbf{v}^* \mathbf{A}^{(1,\ell)} \mathbf{v}) | \mathbf{A}] \leq \phi_{\text{tr}(\mathbf{A})}(\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,\ell)} \mathbf{v} | \mathbf{A}]).$$

Use the boundedness of $\phi_{\text{tr}(\mathbf{A})}$ and monotonicity of $\phi_{\text{tr}(\mathbf{A})}^{(b-\ell-1)}$ (Lemma 4.6a-b) to show

$$\phi_{\text{tr}(\mathbf{A})}^{(b-\ell-1)}(\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,\ell+1)} \mathbf{v} | \mathbf{A}]) \leq \phi_{\text{tr}(\mathbf{A})}^{(b-\ell)}(\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,\ell)} \mathbf{v} | \mathbf{A}]).$$

Chaining the above inequalities for $\ell = 1, 2, \dots, b - 1$ yields

$$\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,b)} \mathbf{v} | \mathbf{A}] \leq \phi_{\text{tr}(\mathbf{A})}^{(b-1)}(\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,1)} \mathbf{v} | \mathbf{A}]).$$

Taking expectations yields

$$(4.3) \quad \begin{aligned} \mathbf{v}^* \mathbb{E}[\mathbf{A}^{(1,b)}] \mathbf{v} &\leq \mathbb{E}[\phi_{\text{tr}(\mathbf{A})}^{(b-1)}(\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,1)} \mathbf{v} | \mathbf{A}])] \\ &\leq \phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\mathbb{E}[\mathbf{v}^* \mathbf{A}^{(1,1)} \mathbf{v}]) \leq \phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\mathbf{v}^* \Phi_1(\mathbb{E}[\mathbf{A}]) \mathbf{v}). \end{aligned}$$

The second inequality is Jensen's inequality applied to the jointly concave function $(x, \alpha) \mapsto \phi_\alpha^{(b-1)}(x)$ (Lemma 4.6d), and the third inequality is (4.2). This establishes a convenient upper bound for the quadratic form $\mathbf{v}^* \mathbb{E}[\mathbf{A}^{(1,b)}] \mathbf{v}$.

Step 3: Apply the min-max principle. Let $\mathbf{v}_1, \dots, \mathbf{v}_N$ be eigenvectors of $\Phi_1(\mathbb{E}[\mathbf{A}])$, associated with its decreasingly ordered eigenvalues. For any unit-length vector $\mathbf{v} \in \text{span}\{\mathbf{v}_i, \dots, \mathbf{v}_N\}$, the min-max variational principle implies

$$\mathbf{v}^* \Phi_1(\mathbb{E}[\mathbf{A}]) \mathbf{v} \leq \lambda_i(\Phi_1(\mathbb{E}[\mathbf{A}])).$$

By (4.3), boundedness of $\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}$, and monotonicity of $\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}$ (Lemma 4.6a-b), it follows

$$\mathbf{v}^* \mathbb{E}[\mathbf{A}^{(1,b)}] \mathbf{v} \leq \phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\mathbf{v}^* \Phi_1(\mathbb{E}[\mathbf{A}]) \mathbf{v}) \leq \phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\lambda_i(\Phi_1(\mathbb{E}[\mathbf{A}]))).$$

By boundedness of $\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}$ and the monotonicity of $\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}$ (Lemma 4.6a-b), the matrices $\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\Phi_1(\mathbb{E}[\mathbf{A}]))$ and $\Phi_1(\mathbb{E}[\mathbf{A}])$ have the same eigenvectors with the same ordering of the eigenvalues. Hence,

$$\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\lambda_i(\Phi_1(\mathbb{E}[\mathbf{A}]))) = \lambda_i(\phi_{\mathbb{E}[\text{tr}(\mathbf{A})]}^{(b-1)}(\Phi_1(\mathbb{E}[\mathbf{A}]))) = \lambda_i(\Phi_b(\mathbb{E}[\mathbf{A}])).$$

Combining the two previous displays yields

$$\mathbf{v}^* \mathbb{E}[\mathbf{A}^{(1,b)}] \mathbf{v} \leq \lambda_i(\Phi_b(\mathbb{E}[\mathbf{A}])) \quad \text{for any unit-length } \mathbf{v} \in \text{span}\{\mathbf{v}_i, \dots, \mathbf{v}_N\}.$$

It follows from the min-max variational principle that $\lambda_i(\mathbb{E}[\mathbf{A}^{(1,b)}]) \leq \lambda_i(\Phi_b(\mathbb{E}[\mathbf{A}]))$, completing the proof. \square

The central result of this section is proved by iteratively applying Proposition 4.8.

Proof of Lemma 4.4. The $(s+1)$ st round of accelerated RPCHOLESKY or block RPCHOLESKY leads to the same residual matrix as a single round of accelerated RPCHOLESKY or block RPCHOLESKY applied to the psd matrix $\mathbf{A}^{(s)}$. Thus, Proposition 4.8 shows that

$$(4.4) \quad \lambda_i(\mathbb{E}[\mathbf{A}^{(s+1)}]) \leq \lambda_i(\Phi_b(\mathbb{E}[\mathbf{A}^{(s)}])) \quad \text{for each } i = 1, \dots, N.$$

Next take eigendecompositions

$$\mathbb{E}[\mathbf{A}^{(s+1)}] = \mathbf{Q}_1 \Lambda_1 \mathbf{Q}_1^* \quad \text{and} \quad \Phi_b(\mathbb{E}[\mathbf{A}^{(s)}]) = \mathbf{Q}_2 \Lambda_2 \mathbf{Q}_2^*,$$

where the eigenvalues are listed in weakly decreasing order. Equation (4.4) implies $\Lambda_1 \preceq \Lambda_2$, so the monotonicity of $\Phi_b^{(t-s-1)}$ (Lemma 4.7a) further implies

$$(4.5) \quad \Phi_b^{(t-s-1)}(\Lambda_1) \preceq \Phi_b^{(t-s-1)}(\Lambda_2).$$

The function Φ_b is covariant under unitary basis transformations:

$$\Phi_b(\mathbf{Q} \Lambda \mathbf{Q}^*) = \mathbf{Q} \Phi_b(\Lambda) \mathbf{Q}^* \quad \text{for unitary } \mathbf{Q} \in \mathbb{C}^{N \times N}.$$

Therefore, the eigenvalue ordering (4.5) implies

$$\begin{aligned} \mathbf{Q}_1^* \Phi_b^{(t-s-1)}(\mathbb{E}[\mathbf{A}^{(s+1)}]) \mathbf{Q}_1 &= \Phi_b^{(t-s-1)}(\Lambda_1) \\ &\preceq \Phi_b^{(t-s-1)}(\Lambda_2) = \mathbf{Q}_2^* \Phi_b^{(t-s)}(\mathbb{E}[\mathbf{A}^{(s)}]) \mathbf{Q}_2. \end{aligned}$$

Since the eigenvalues are invariant under unitary basis transformations, Weyl's inequality implies

$$(4.6) \quad \begin{aligned} \lambda_i(\Phi_b^{(t-s-1)}(\mathbb{E}[\mathbf{A}^{(s+1)}])) &= \lambda_i(\mathbf{Q}_1^* \Phi_b^{(t-s-1)}(\mathbb{E}[\mathbf{A}^{(s+1)}]) \mathbf{Q}_1) \\ &\leq \lambda_i(\mathbf{Q}_2^* \Phi_b^{(t-s)}(\mathbb{E}[\mathbf{A}^{(s)}]) \mathbf{Q}_2) = \lambda_i(\Phi_b^{(t-s)}(\mathbb{E}[\mathbf{A}^{(s)}])). \end{aligned}$$

Apply (4.6) for $s = 0, 1, \dots, t-1$ to complete the proof. \square

4.5. Part III: Permutation averaging lemma. The permutation averaging lemma has the following short and simple proof.

Proof of Lemma 4.5. Since the trace is invariant under permutations and the function Φ_b is covariant under permutations,

$$\text{tr}(\Phi_b^{(t)}(\mathbf{\Lambda})) = \text{tr}(\mathbf{P} \Phi_b^{(t)}(\mathbf{\Lambda}) \mathbf{P}^*) = \text{tr}(\Phi_b^{(t)}(\mathbf{P} \mathbf{\Lambda} \mathbf{P}^*)).$$

Now apply Jensen's inequality to the concave function $\Phi_b^{(t)}$ (Lemma 4.7b):

$$\mathbb{E}[\Phi_b^{(t)}(\mathbf{P} \mathbf{\Lambda} \mathbf{P}^*)] \preceq \Phi_b^{(t)}(\mathbb{E}[\mathbf{P} \mathbf{\Lambda} \mathbf{P}^*]).$$

Since the trace respects the psd ordering, it follows

$$\text{tr}(\Phi_b^{(t)}(\mathbf{\Lambda})) = \text{tr}(\mathbb{E}[\Phi_b^{(t)}(\mathbf{P} \mathbf{\Lambda} \mathbf{P}^*)]) \leq \text{tr}(\Phi_b^{(t)}(\mathbb{E}[\mathbf{P} \mathbf{\Lambda} \mathbf{P}^*])).$$

This completes the proof. \square

4.6. Part IV: Main error bound. At this point, the main error bound for accelerated RPCHOLESKY is nearly in hand. The last step is an analysis argument that explicitly bounds the error for a worst-case instance of accelerated RPCHOLESKY.

Proof of Theorem 4.2. Take an eigendecomposition $\mathbf{A} = \mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^*$. By Lemma 4.4,

$$\mathbb{E}[\text{tr}(\mathbf{A}^{(t)})] \leq \text{tr}(\Phi_b^{(t)}(\mathbf{A})) = \text{tr}(\Phi_b^{(t)}(\mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^*)).$$

Since the trace is invariant under unitary basis transformations and the function Φ_b is covariant under unitary basis transformations,

$$\text{tr}(\Phi_b^{(t)}(\mathbf{Q} \mathbf{\Lambda} \mathbf{Q}^*)) = \text{tr}(\mathbf{Q} \Phi_b^{(t)}(\mathbf{\Lambda}) \mathbf{Q}^*) = \text{tr}(\Phi_b^{(t)}(\mathbf{\Lambda})).$$

Therefore, the upper bound only depends on the eigenvalue matrix $\mathbf{\Lambda} \in \mathbb{R}^{N \times N}$.

Next, take a random permutation \mathbf{P} that permutes the first r coordinates uniformly at random and permutes the trailing $N-r$ coordinates uniformly at random. By Lemma 4.5,

$$\text{tr}(\Phi_b^{(t)}(\mathbf{\Lambda})) \leq \text{tr}(\Phi_b^{(t)}(\widehat{\mathbf{\Lambda}})), \quad \text{where } \widehat{\mathbf{\Lambda}} = \mathbb{E}[\mathbf{P} \mathbf{\Lambda} \mathbf{P}^*].$$

The matrix $\widehat{\mathbf{\Lambda}}$ has the explicit form

$$\widehat{\mathbf{\Lambda}} = \text{diag} \left\{ \underbrace{\frac{\alpha}{r}, \dots, \frac{\alpha}{r}}_{r \text{ times}}, \underbrace{\frac{\beta}{N-r}, \dots, \frac{\beta}{N-r}}_{N-r \text{ times}} \right\}, \quad \text{where} \quad \begin{cases} \alpha = \sum_{i=1}^r \lambda_i(\mathbf{A}), \\ \beta = \sum_{i=r+1}^N \lambda_i(\mathbf{A}). \end{cases}$$

It is a two-cluster matrix that leads to the worst possible value of $\text{tr}(\Phi_b^{(t)}(\mathbf{\Lambda}))$

For the next step, argue by induction that

$$(4.7) \quad \Phi_b^{(t)}(\widehat{\Lambda}) \preceq \widehat{\Lambda}^{(t)}, \quad \text{where } \widehat{\Lambda}^{(t)} = \text{diag} \left\{ \underbrace{\frac{\alpha^{(t)}}{r}, \dots, \frac{\alpha^{(t)}}{r}}_{r \text{ times}}, \underbrace{\frac{\beta}{N-r}, \dots, \frac{\beta}{N-r}}_{N-r \text{ times}} \right\},$$

and the entries $\alpha^{(t)}$ are given by the recursion

$$(4.8) \quad \alpha^{(t+1)} = \alpha^{(t)} - \frac{b(\alpha^{(t)})^2}{(b+r)\alpha^{(t)} + r\beta}, \quad \alpha^{(0)} = \alpha.$$

The comparison (4.7) is valid with equality at time $t = 0$. Assume the comparison is valid at a time $t \geq 0$. Using the monotonicity of Φ_b (Lemma 4.7a), it follows that

$$\Phi_b^{(t+1)}(\widehat{\Lambda}) \preceq \Phi_b(\widehat{\Lambda}^{(t)}) = \phi_{\alpha^{(t)}+\beta}^{(b)}(\widehat{\Lambda}^{(t)}).$$

By Lemma 4.6e, the trailing $N - r$ diagonal entries of $\phi_{\alpha^{(t)}+\beta}^{(b)}(\widehat{\Lambda}^{(t)})$ are bounded by $\frac{\beta}{N-r}$ and the leading r diagonal entries are bounded by

$$\frac{1}{\frac{b}{\alpha^{(t)}+\beta} + \frac{r}{\alpha^{(t)}}} = \frac{1}{r} \left[\alpha^{(t)} - \frac{b(\alpha^{(t)})^2}{(b+r)\alpha^{(t)} + r\beta} \right].$$

Therefore, (4.7) holds for all $t \geq 0$ by induction.

At each instant $t = 0, 1, \dots$, the discrete-time dynamical system $\alpha^{(t)}$ (4.8) is bounded from above by a continuous-time dynamical system $\alpha(t)$ that satisfies

$$(4.9) \quad \dot{\alpha}(t) = -\frac{b(\alpha(t))^2}{(b+r)\alpha(t) + r\beta}, \quad \alpha(0) = \alpha.$$

The comparison holds because $x \mapsto bx^2/((b+r)x+r\beta)$ is a non-negative, nondecreasing function.

Last, assume that $\alpha \leq \varepsilon\beta$. (Otherwise, the theorem holds trivially at all times $t \geq 0$.) By separation of variables, the continuous-time dynamical system (4.9) satisfies $\alpha(t) \leq \varepsilon\beta$ as soon as

$$t = \left[-\frac{b+r}{b} \log(\alpha(t)) + \frac{r\beta}{b} \alpha(t)^{-1} \right]_{\alpha(0)=\alpha}^{\alpha(t)=\varepsilon\beta} = \frac{b+r}{b} \log\left(\frac{1}{\varepsilon\eta}\right) + \frac{r}{b\varepsilon} - \frac{r}{b}\eta.$$

This establishes a slightly stronger version of the theorem, completing the proof. \square

5. Extension: Accelerated randomly pivoted QR. The pivoted partial Cholesky decomposition of a psd matrix \mathbf{A} is closely related to the pivoted partial QR decomposition of a general, possibly rectangular, matrix $\mathbf{B} \in \mathbb{C}^{M \times N}$. This section reviews the connection between algorithms and introduces a new *accelerated randomly pivoted QR algorithm* for general low-rank matrix approximation.

A low-rank approximation for a general matrix can be computed by column-pivoted QR as follows: Begin by initializing the approximation $\widehat{\mathbf{B}}^{(0)} := \mathbf{0}$ and the residual $\mathbf{B}^{(0)} := \mathbf{B}$. For each $i = 0, \dots, k - 1$, perform the following steps:

1. **Select a pivot column.** Choose $s_{i+1} \in \{1, \dots, N\}$.
2. **Update.** Evaluate the column $\mathbf{B}^{(i)}(:, s_{i+1})$ indexed by the pivot index s_{i+1} . Update the approximation and the residual:

$$\begin{aligned}\widehat{\mathbf{B}}^{(i+1)} &:= \widehat{\mathbf{B}}^{(i)} + \frac{\mathbf{B}^{(i)}(:, s_{i+1})\mathbf{B}^{(i)}(:, s_{i+1})^* \mathbf{B}^{(i)}}{\|\mathbf{B}^{(i)}(:, s_{i+1})\|^2}, \\ \mathbf{B}^{(i+1)} &:= \mathbf{B}^{(i)} - \frac{\mathbf{B}^{(i)}(:, s_{i+1})\mathbf{B}^{(i)}(:, s_{i+1})^* \mathbf{B}^{(i)}}{\|\mathbf{B}^{(i)}(:, s_{i+1})\|^2}.\end{aligned}$$

Each update has the effect of orthogonalizing the residual $\mathbf{B}^{(i)}$ against the column indexed by the selected pivot and adjusting the approximation $\widehat{\mathbf{B}}^{(i)}$ accordingly. This version of column-pivoted QR has issues with numerical stability and is provided for conceptual understanding only. Numerically stable implementation of QR decomposition methods is discussed in numerical linear algebra textbooks, e.g., [20, ch. 5].

Pivoted QR and Cholesky approximations are closely related (see, e.g., [21]):

PROPOSITION 5.1 (QR and Cholesky). *Suppose that column-pivoted QR is executed on a matrix \mathbf{B} with pivots s_1, \dots, s_k producing an approximation $\widehat{\mathbf{B}} \approx \mathbf{B}$. Similarly, suppose that partial Cholesky is performed on the Gram matrix $\mathbf{A} := \mathbf{B}^* \mathbf{B}$ with the same pivots s_1, \dots, s_k , producing an approximation $\widehat{\mathbf{A}} \approx \mathbf{A}$. Then $\widehat{\mathbf{B}}^* \widehat{\mathbf{B}} = \widehat{\mathbf{A}}$.*

This result establishes the connection between the low-rank approximation of a psd matrix based on the pivoted Cholesky decomposition and the low-rank approximation of a general, rectangular matrix by the column-pivoted QR decomposition. In particular, any pivot selection strategy for Cholesky decomposition (i.e., random pivoting) has an analogue for QR decomposition and visa versa.

The analogue of RPCHOLESKY is the randomly pivoted QR algorithm, which was introduced by Deshpande, Vempala, Rademacher, and Wang [9, 10] under the name adaptive sampling. See Algorithm 5.1 for randomly pivoted QR pseudocode. As with RPCHOLESKY, randomly pivoted QR can be extended using block and accelerated variants; see [4] for block randomly pivoted QR pseudocode and see Algorithm 5.2 for accelerated randomly pivoted QR pseudocode. Note that Algorithm 5.2 uses a block Gram–Schmidt procedure to perform orthogonalization, which employs two orthogonalization steps for improved stability [19]. A different, even more numerically stable approach uses Householder reflectors instead [20, Ch. 5].

Due to the connection between QR and Cholesky (Proposition 5.1), theoretical results for block and accelerated RPCHOLESKY immediately extend to their QR analogues.

Algorithm 5.1. Randomly pivoted QR.

Input: Matrix $\mathbf{B} \in \mathbb{C}^{M \times N}$, approximation rank k

Output: Matrices $\mathbf{Q} \in \mathbb{C}^{M \times k}$, $\mathbf{F} \in \mathbb{C}^{N \times k}$ defining $\widehat{\mathbf{B}} = \mathbf{Q}\mathbf{F}^*$; pivot set \mathbf{S}

Initialize $\mathbf{Q} \leftarrow \mathbf{0}_{M \times k}$, $\mathbf{F} \leftarrow \mathbf{0}_{N \times k}$, $\mathbf{S} \leftarrow \emptyset$, and $\mathbf{d} \leftarrow \text{SQUARED COLUMN NORMS } \mathbf{B}$

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

Sample pivot $s \sim \mathbf{d} / \sum_j d_j$

Induct new pivot $\mathbf{S} \leftarrow \mathbf{S} \cup \{s\}$

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

$\mathbf{Q}(:, i) \leftarrow \mathbf{g} / \|\mathbf{g}\|$

$\mathbf{F}(:, i) \leftarrow \mathbf{B}^* \mathbf{Q}(:, i)$

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

end for

Algorithm 5.2. Accelerated randomly pivoted QR.

Input: Matrix $\mathbf{A} \in \mathbb{C}^{M \times N}$; block size b ; number of rounds t

Output: Matrices $\mathbf{Q} \in \mathbb{C}^{M \times k}$, $\mathbf{F} \in \mathbb{C}^{N \times k}$ defining $\widehat{\mathbf{B}} = \mathbf{Q}\mathbf{F}^*$; pivot set \mathbf{S}

Initialize $\mathbf{Q} \leftarrow \mathbf{0}_{M \times 0}$, $\mathbf{F} \leftarrow \mathbf{0}_{N \times 0}$, $\mathbf{S} \leftarrow \emptyset$, and $\mathbf{u} \leftarrow \text{SQUARED COLUMN NORMS } \mathbf{B}$

for $i = 0$ to $t - 1$ **do**

Step 1: Propose a block of pivots

Sample $s'_{ib+1}, \dots, s'_{(i+1)b} \stackrel{\text{iid}}{\sim} \mathbf{u}$ and set $\mathbf{S}'_i \leftarrow \{s'_{ib+1}, \dots, s'_{(i+1)b}\}$

Step 2: Rejection sampling

$\mathbf{C} \leftarrow \mathbf{B}(:, \mathbf{S}'_i) - \mathbf{Q}\mathbf{F}(\mathbf{S}'_i, :)^*$

$\mathbf{S}_i \leftarrow \text{REJECTION SAMPLE SUBMATRIX } \mathbf{S}'_i, \mathbf{C}^* \mathbf{C} \quad \triangleright \text{Discard second output}$

$\mathbf{S} \leftarrow \mathbf{S} \cup \mathbf{S}_i \quad \triangleright \text{Update pivots}$

Step 3: Update low-rank approximation and proposal distribution

$\mathbf{Q}_\perp \leftarrow \mathbf{B}(:, \mathbf{S}_i) - \mathbf{Q}\mathbf{F}(\mathbf{S}_i, :)^*$

$\mathbf{Q}_\perp \leftarrow \mathbf{Q}_\perp - \mathbf{Q}(\mathbf{Q}^* \mathbf{Q}_\perp) \quad \triangleright \text{Extra step of Gram-Schmidt, for stability}$

$\mathbf{Q}_\perp \leftarrow \text{ORTH } \mathbf{Q}_\perp \quad \triangleright \text{Orthonormalize columns}$

$\mathbf{Q} \leftarrow [\mathbf{Q} \quad \mathbf{Q}_\perp]$

$\mathbf{G} \leftarrow \mathbf{B}^* \mathbf{Q}_\perp$

$\mathbf{F} \leftarrow [\mathbf{F} \quad \mathbf{G}]$

$\mathbf{u} \leftarrow \mathbf{u} - \text{SQUARED ROW NORMS } \mathbf{G} \quad \triangleright \text{Update diagonal}$

$\mathbf{u} \leftarrow \max\{\mathbf{u}, \mathbf{0}\} \quad \triangleright \text{Helpful in floating point arithmetic}$

end for

COROLLARY 5.2 (sufficient iterations for accelerated and block randomly pivoted QR). Consider a rectangular matrix $\mathbf{B} \in \mathbb{C}^{M \times N}$, and let $\mathbf{B}^{(t)}$ denote the random residual after applying t rounds of accelerated randomly pivoted QR or block randomly pivoted QR with a block size $b \geq 1$. Then,

$$\mathbb{E} \|\mathbf{B}^{(t)}\|_{\text{F}}^2 \leq (1 + \varepsilon) \cdot \sum_{i=r+1}^N \sigma_i(\mathbf{B})^2$$

as soon as

$$tb \geq \frac{r}{\varepsilon} + (r + b) \log \left(\frac{1}{\varepsilon \eta} \right), \quad \text{where } \eta = \frac{\sum_{i=r+1}^N \sigma_i(\mathbf{B})^2}{\sum_{i=1}^r \sigma_i(\mathbf{B})^2}.$$

Accelerated randomly pivoted QR is a new algorithm, while block randomly pivoted QR was introduced and analyzed in [9, Thm. 1.2]. This earlier analysis has the large block size limitation that is discussed in section 4.

Appendix A. Why blocking? Figures 1 to 3 demonstrate the empirical speed-ups of algorithms that simultaneously process large blocks of columns from kernel matrices. Subsection 1.2 has discussed one reason for this speed-up, due to data movement. This appendix describes two additional reasons why blocked (kernel) matrix algorithms can be faster: the fast Euclidean distances trick (subsection A.1) and matrix-matrix operations (subsection A.2).

A.1. The fast Euclidean distances trick. The fast Euclidean distances trick is a method for evaluating kernel matrix entries that take the special form

$$\kappa(\mathbf{x}_i, \mathbf{x}_j) = \phi(\|\mathbf{x}_i - \mathbf{x}_j\|) \quad \text{for } i \in \mathbf{S}_1, j \in \mathbf{S}_2,$$

where $\phi: \mathbb{R} \rightarrow \mathbb{R}$ is a univariate function. The trick relies on decomposing the square Euclidean distance into three summands:

$$(A.1) \quad \|\mathbf{x}_i - \mathbf{x}_j\|^2 = \|\mathbf{x}_i\|^2 - 2\operatorname{Re}\{\mathbf{x}_i^* \mathbf{x}_j\} + \|\mathbf{x}_j\|^2.$$

It takes just $\mathcal{O}(d|\mathcal{S}_1| + d|\mathcal{S}_2|)$ operations to compute the square norms $\|\mathbf{x}_i\|^2$ and $\|\mathbf{x}_j\|^2$ for each $i \in \mathcal{S}_1$ and $j \in \mathcal{S}_2$. Meanwhile, the inner products $\mathbf{x}_i^* \mathbf{x}_j$ can be obtained from the entries of the inner product matrix

$$(A.2) \quad \mathbf{X}(:, \mathcal{S}_1)^* \mathbf{X}(:, \mathcal{S}_2) \quad \text{where} \quad \mathbf{X} = [\mathbf{x}_1 \quad \cdots \quad \mathbf{x}_N].$$

Computing the inner product matrix requires $\mathcal{O}(d|\mathcal{S}_1||\mathcal{S}_2|)$ operations using a single matrix–matrix multiplication, a highly optimized operation on modern computers [13, Ch. 20.1]. The fast Euclidean distances trick explains why generating columns can be $10\times$ faster for a high-dimensional Gaussian kernel (Figure 1, left) than an ℓ_1 Laplace kernel (Figure 1, right).

As a potential concern, the fast Euclidean distances trick can lead to *catastrophic cancellations* [22, sec. 1.7] in floating point arithmetic when two points \mathbf{x}_i and \mathbf{x}_j are close. Catastrophic cancellations sometimes produce inaccurate results for block RPCHOLESKY, even after adding a multiple of the machine precision to the diagonal of the target matrix to promote stability. In contrast, accelerated RPCHOLESKY avoids catastrophic cancellations by rejecting the small diagonal values as pivots. The opportunity to use the fast Euclidean distances trick without catastrophic cancellations is a major advantage of accelerated RPCHOLESKY. Nonetheless, to promote a simple comparison, all the experiments involving block RPCHOLESKY avoid the fast Euclidean distances trick; the code computes distances by subtracting $\mathbf{x}_i - \mathbf{x}_j$ and directly calculating the norm.

A.2. Matrix–matrix operations. The cost of RPCHOLESKY-type methods can be decomposed into the cost of generating columns of the kernel matrix and the cost of processing the columns to execute the algorithm. Blocking leads to speed-ups in both steps of the algorithm. For the second step, blocking organizes the operations into efficient matrix–matrix operations. Accelerated RPCHOLESKY takes advantage of these optimized operations for an improved runtime. Simple RPCHOLESKY, by contrast, uses only vector–vector and matrix–vector operations, which are slower.

Appendix B. Deferred proofs from Appendix 4. This section contains deferred proofs from Appendix 4.

Proof of Lemma 4.6. The proof is based on a series of direct calculations.

Parts (a)–(b). Observe that

$$(B.1) \quad \phi_\alpha(x) = \frac{\alpha}{4} - \frac{1}{\alpha} \left(x - \frac{\alpha}{2}\right)^2$$

is a concave quadratic which is uniformly bounded by $\frac{\alpha}{4}$. This immediately establishes part (a) and part (b) in the case $b = 1$. Next assume part (b) is valid for some $b \geq 1$ and assume $x \leq x' \leq \alpha/2$. By the boundedness property (a), $\phi_\alpha^{(b)}(x) \leq \phi_\alpha^{(b)}(x') \leq \alpha/4$. Since $x \mapsto \phi_\alpha(x)$ is nondecreasing over the range $[0, \frac{\alpha}{2}]$, it follows:

$$\phi_\alpha^{(b+1)}(x) = \phi_\alpha(\phi_\alpha^{(b)}(x)) \leq \phi_\alpha(\phi_\alpha^{(b)}(x')) = \phi_\alpha^{(b+1)}(x'),$$

which establishes part (b) by induction.

Part (c). Calculate

$$\phi_{\beta \cdot \alpha}(\beta \cdot x) = \beta x - \frac{\beta}{\alpha} x^2 = \beta \left[x - \frac{x^2}{\alpha} \right] = \beta \cdot \phi_{\alpha}(x),$$

which confirms part (c) when $b = 1$. Homogeneity for $b \geq 1$ follows by induction.

Part (d). The quadratic representation (B.1) shows that ϕ_1 is concave. Assume that $\phi_1^{(b)}$ is concave for some $b \geq 1$. Then, calculate for $\theta \in (0, 1)$:

$$\begin{aligned} \theta \phi_1^{(b+1)}(x) + (1 - \theta) \phi_1^{(b+1)}(y) &= \theta \phi_1(\phi_1^{(b)}(x)) + (1 - \theta) \phi_1(\phi_1^{(b)}(y)) \\ &\leq \phi_1(\theta \phi_1^{(b)}(x) + (1 - \theta) \phi_1^{(b)}(y)) \\ &\leq \phi_1(\phi_1^{(b)}(\theta x + (1 - \theta)y)) = \phi_1^{(b+1)}(\theta x + (1 - \theta)y). \end{aligned}$$

The first inequality is concavity of ϕ_1 , and the second inequality is concavity of $\phi_1^{(b)}$ and parts (a)–(b). Conclude by induction that $\phi_1^{(b)}$ is concave for each $b \geq 1$.

The next step is to establish joint concavity of $(x, \alpha) \mapsto \phi_{\alpha}^{(b)}(x)$. For any $\theta \in (0, 1)$:

$$\begin{aligned} \theta \phi_{\alpha_1}^{(b)}(x_1) + (1 - \theta) \phi_{\alpha_2}^{(b)}(x_2) &= \theta \alpha_1 \phi_1^{(b)}\left(\frac{x_1}{\alpha_1}\right) + (1 - \theta) \alpha_2 \phi_1^{(b)}\left(\frac{x_2}{\alpha_2}\right) \\ &\leq [\theta \alpha_1 + (1 - \theta) \alpha_2] \phi_1^{(b)}\left(\frac{\theta x_1 + (1 - \theta)x_2}{\theta \alpha_1 + (1 - \theta)\alpha_2}\right) \\ &= \phi_{\theta \alpha_1 + (1 - \theta)\alpha_2}^{(b)}(\theta x_1 + (1 - \theta)x_2). \end{aligned}$$

The first equality is part (c), and the inequality is the concavity of $\phi_1^{(b)}$. Part (d) is proven.

Part (e). Calculate

$$\phi_{\alpha}(x) = x - \frac{x^2}{\alpha} \leq x - \frac{x^2}{x + \alpha} = \frac{\alpha x}{x + \alpha} = \frac{1}{\alpha^{-1} + x^{-1}},$$

which confirms part (e) in the case $b = 1$. Next assume part (e) holds for some value $b \geq 1$. Then using the monotonicity property (b),

$$\begin{aligned} \phi_{\alpha}^{(b+1)}(x) &= \phi_{\alpha}(\phi_{\alpha}^{(b)}(x)) \\ &\leq \phi_{\alpha}\left(\frac{1}{b\alpha^{-1} + x^{-1}}\right) = \frac{1}{b\alpha^{-1} + x^{-1}} - \frac{1}{(b\alpha^{-1} + x^{-1})(b + \alpha x^{-1})} \\ &\leq \frac{1}{b\alpha^{-1} + x^{-1}} - \frac{1}{(b\alpha^{-1} + x^{-1})(b + 1 + \alpha x^{-1})} = \frac{1}{(b + 1)\alpha^{-1} + x^{-1}}. \end{aligned}$$

Therefore, part (e) holds by induction. \square

Proof of Lemma 4.7. The proof is based on a series of direct calculations.

Part (a). Assume $\mathbf{\Lambda}_1 \preceq \mathbf{\Lambda}_2$. Use the positive homogeneity and concavity of $(x, \alpha) \mapsto \phi_{\alpha}^{(b)}(x)$ (Lemma 4.6c–d) to calculate

$$\begin{aligned} \frac{1}{2} \phi_{\text{tr}(\mathbf{\Lambda}_1)}^{(b)}(\mathbf{\Lambda}_1) &\preceq \frac{1}{2} \phi_{\text{tr}(\mathbf{\Lambda}_1)}^{(b)}(\mathbf{\Lambda}_1) + \frac{1}{2} \phi_{\text{tr}(\mathbf{\Lambda}_2 - \mathbf{\Lambda}_1)}^{(b)}(\mathbf{\Lambda}_2 - \mathbf{\Lambda}_1) \\ &\preceq \phi_{\frac{1}{2} \text{tr}(\mathbf{\Lambda}_2)}^{(b)}\left(\frac{1}{2} \mathbf{\Lambda}_2\right) = \frac{1}{2} \phi_{\text{tr}(\mathbf{\Lambda}_2)}^{(b)}(\mathbf{\Lambda}_2). \end{aligned}$$

Equivalently, $\Phi_b(\mathbf{\Lambda}_1) \preceq \Phi_b(\mathbf{\Lambda}_2)$. Iterating this property establishes part (a).

Part (b). Use the concavity of $(x, \alpha) \mapsto \phi_\alpha^{(b)}(x)$ (Lemma 4.6d) to conclude:

$$\theta \Phi_b(\Lambda_1) + (1 - \theta) \Phi_b(\Lambda_2) \preceq \Phi_b(\theta \Lambda_1 + (1 - \theta) \Lambda_2) \quad \text{for } \theta \in (0, 1).$$

This confirms part (b) in the case $t = 1$. Next assume part (b) holds for $t \geq 1$. Then

$$\begin{aligned} \theta \Phi_b^{(t+1)}(\Lambda_1) + (1 - \theta) \Phi_b^{(t+1)}(\Lambda_2) &= \theta \Phi_b(\Phi_b^{(t)}(\Lambda_1)) + (1 - \theta) \Phi_b(\Phi_b^{(t)}(\Lambda_2)) \\ &\preceq \Phi_b(\theta \Phi_b^{(t)}(\Lambda_1) + (1 - \theta) \Phi_b^{(t)}(\Lambda_2)) \\ &\preceq \Phi_b(\Phi_b^{(t)}(\theta \Lambda_1 + (1 - \theta) \Lambda_2)) \\ &= \Phi_b^{(t+1)}(\theta \Lambda_1 + (1 - \theta) \Lambda_2). \end{aligned}$$

The first inequality is the concavity of Φ_b , and the second inequality is the concavity of $\Phi_b^{(t)}$ and the monotonicity property (a). Therefore, by induction, part (b) holds for $t \geq 1$.

Part (c). Calculate

$$\begin{aligned} &\Phi_1(\theta \mathbf{A}_1 + (1 - \theta) \mathbf{A}_2) - \theta \Phi_1(\mathbf{A}_1) - (1 - \theta) \Phi_1(\mathbf{A}_2) \\ &= \frac{\theta \mathbf{A}_1^2}{\text{tr}(\mathbf{A}_1)} + \frac{(1 - \theta) \mathbf{A}_2^2}{\text{tr}(\mathbf{A}_2)} - \frac{[\theta \mathbf{A}_1 + (1 - \theta) \mathbf{A}_2]^2}{\theta \text{tr}(\mathbf{A}_1) + (1 - \theta) \text{tr}(\mathbf{A}_2)} \\ &= \frac{\theta(1 - \theta) \text{tr}(\mathbf{A}_1) \text{tr}(\mathbf{A}_2)}{\theta \text{tr}(\mathbf{A}_1) + (1 - \theta) \text{tr}(\mathbf{A}_2)} \left[\frac{\mathbf{A}_1}{\text{tr}(\mathbf{A}_1)} - \frac{\mathbf{A}_2}{\text{tr}(\mathbf{A}_2)} \right]^2 \succeq \mathbf{0}. \end{aligned}$$

This establishes the concavity of Φ_1 over all psd matrices. \square

Appendix C. Robust blockwise random pivoting. The robust blockwise random pivoting (RBRP) method was introduced by Dong, Chen, Martinsson, and Pearce to improve the performance of block randomly pivoted QR [12]. Similar to accelerated randomly pivoted QR, each round of RBRP begins by generating a block of proposed pivots $S'_i = \{s'_{ib+1}, \dots, s'_{(i+1)b}\}$. Then RBRP chooses a subset of the pivots by a deterministic, greedy method. In detail, let \mathbf{G} denote the residual submatrix with columns indexed by S'_i :

$$\mathbf{G} = \mathbf{B}(:, S'_i) - \widehat{\mathbf{B}}(:, S'_i).$$

RBRP applies a partial QR decomposition to \mathbf{G} using greedy column pivoting, stopping when the squared Frobenius norm of \mathbf{G} has been reduced by a factor b . RBRP then accepts the pivots chosen before the tolerance is met and rejects all others.

RBRP can be extended to psd low-rank approximation, resulting in a *robust block randomly pivoted Cholesky* method. A comparison of RBRP Cholesky, block RPCHOLESKY, and accelerated RPCHOLESKY appears in Table 3. Both accelerated

TABLE 3

Runtime and relative error for rank-1,000 approximation of the smile example from Figure 2 with block size $b = 120$. Here the mean and standard deviation are computed empirically over 100 independent trials.

	Relative trace error	Runtime (sec)
Block RPCHOLESKY	$3.89\text{e-}04 \pm 1.40\text{e-}04$	7.00 ± 0.63
Accelerated RPCHOLESKY	$4.85\text{e-}07 \pm 3.60\text{e-}08$	7.43 ± 1.06
RBRP Cholesky	$6.62\text{e-}07 \pm 1.35\text{e-}07$	7.53 ± 0.80

Algorithm D.1. RPCHOLESKY.

Input: Psd matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$; approximation rank k
Output: Matrix $\mathbf{F} \in \mathbb{C}^{N \times k}$ defining low-rank approximation $\widehat{\mathbf{A}} = \mathbf{F}\mathbf{F}^*$; pivot set \mathbf{S}
Initialize $\mathbf{F} \leftarrow \mathbf{0}_{N \times k}$ and $\mathbf{d} \leftarrow \text{diag } \mathbf{A}$
for $i = 1$ to k **do**
 Sample pivot $s \sim \mathbf{d} / \sum_j d_j$
 Induct new pivot $\mathbf{S} \leftarrow \mathbf{S} \cup \{s\}$
 $\mathbf{g} \leftarrow \mathbf{A}(:, s) - \mathbf{F}(:, 1:i-1)\mathbf{F}(s, 1:i-1)^*$
 $\mathbf{F}(:, i) \leftarrow \mathbf{g} / \sqrt{g_s}$
 $\mathbf{d} \leftarrow \mathbf{d} - |\mathbf{F}(:, i)|^2$
end for

Algorithm D.2. Block RPCHOLESKY.

Input: Psd matrix $\mathbf{A} \in \mathbb{C}^{N \times N}$; block size b ; number of rounds t
Output: Matrix $\mathbf{F} \in \mathbb{C}^{N \times bt}$ defining low-rank approximation $\widehat{\mathbf{A}} = \mathbf{F}\mathbf{F}^*$; pivot set \mathbf{S}
Initialize $\mathbf{F} \leftarrow \mathbf{0}_{N \times bk}$, $\mathbf{S} \leftarrow \emptyset$, and $\mathbf{d} \leftarrow \text{diag } \mathbf{A}$
for $i = 0$ to $t - 1$ **do**
 Sample $s_{ib+1}, \dots, s_{ib+b} \stackrel{\text{iid}}{\sim} \mathbf{d} / \sum_j d_j$
 $\mathbf{S}' \leftarrow \text{UNIQUE}\{s_{ib+1}, \dots, s_{ib+b}\}$
 $\mathbf{S} \leftarrow \mathbf{S} \cup \mathbf{S}'$
 $\mathbf{G} \leftarrow \mathbf{A}(:, \mathbf{S}') - \mathbf{F}(:, 1:iT)\mathbf{F}(\mathbf{S}', 1:iT)^*$
 $\mathbf{R} \leftarrow \text{CHOLG}(\mathbf{S}', :)$
 $\mathbf{F}(:, ib+1:ib+|\mathbf{S}'|) \leftarrow \mathbf{G}\mathbf{R}^{-1}$
 $\mathbf{d} \leftarrow \mathbf{d} - \text{SQUAREDROWNORMS}\mathbf{F}(:, ib+1:ib+b)$
end for
Remove zero columns from \mathbf{F}

RPCHOLESKY and RBRP Cholesky achieve dramatically lower trace error than standard block RPCHOLESKY on this example, demonstrating that both methods are effective at filtering redundant columns. This example does suggest a few preliminary reasons to (slightly) prefer accelerated RPCHOLESKY over RBRP Cholesky, though. First, RBRP Cholesky has $1.4\times$ higher trace error than accelerated RPCHOLESKY on this example, and the error of the output is more variable for RBRP Cholesky. Second, RBRP Cholesky is slightly slower than accelerated RPCHOLESKY in our implementation. Finally, rigorous error bounds are known for accelerated RPCHOLESKY (and QR) but not known for RBRP methods. Notwithstanding these difference, RBRP methods and accelerated RPCHOLESKY/QR both appear to be fast and accurate in practice, and both methods significantly outperform standard block RPCHOLESKY/QR.

Appendix D. Existing algorithms. This section provides the pseudocode for RPCHOLESKY (Algorithm D.1) and block RPCHOLESKY (Algorithm D.2).

Reproducibility of computational results. This paper has been awarded the “SIAM Reproducibility Badge: Code and data available” as a recognition that the authors have followed reproducibility principles valued by SIMAX and the scientific computing community. Code and data that allow readers to reproduce the results in

this paper are available at <https://github.com/epperly/Randomly-Pivoted-Cholesky> and in the supplementary materials (github-repo.zip [local/web 82.9KB]).

Acknowledgments. We thank Vivek Bharadwaj, Chao Chen, and Yijun Dong for helpful discussions.

REFERENCES

- [1] A. ALAOU AND M. W. MAHONEY, *Fast randomized kernel ridge regression with statistical guarantees*, in Proceedings of the 28th International Conference on Neural Information Processing Systems, 2015, <https://dl.acm.org/doi/10.5555/2969239.2969326>.
- [2] D. ARISTOFF, M. JOHNSON, G. SIMPSON, AND R. J. WEBBER, *The fast committor machine: Interpretable prediction with kernels*, *J. Chem. Phys.*, 161 (2024), 084113, <https://doi.org/10.1063/5.0222798>.
- [3] H. AVRON, K. L. CLARKSON, AND D. P. WOODRUFF, *Faster kernel ridge regression using sketching and preconditioning*, *SIAM J. Matrix Anal. Appl.*, 38 (2017), pp. 1116–1138, <https://doi.org/10.1137/16M1105396>.
- [4] Y. CHEN, E. N. EPPERLY, J. A. TROPP, AND R. J. WEBBER, *Randomly pivoted Cholesky: Practical approximation of a kernel matrix with few entry evaluations*, *Commun. Pure Appl. Math.*, 78 (2025), pp. 995–1041, <https://doi.org/10.1002/cpa.22234>.
- [5] S. CHMIELA, H. E. SAUCEDA, K.-R. MÜLLER, AND A. TKATCHENKO, *Towards exact molecular dynamics simulations with machine-learned force fields*, *Nat. Commun.*, 9 (2018), 3887, <https://doi.org/10.1038/s41467-018-06169-2>.
- [6] S. CHMIELA, A. TKATCHENKO, H. E. SAUCEDA, I. POLTAVSKY, K. T. SCHÜTT, AND K.-R. MÜLLER, *Machine learning of accurate energy-conserving molecular force fields*, *Sci. Adv.*, 3 (2017), e1603015, <https://doi.org/10.1126/sciadv.1603015>.
- [7] S. CHMIELA, V. VASSILEV-GALINDO, O. T. UNKE, A. KABYLDA, H. E. SAUCEDA, A. TKATCHENKO, AND K.-R. MÜLLER, *Accurate global machine learning force fields for molecules with hundreds of atoms*, *Sci. Adv.*, 9 (2023), eadf0873, <https://doi.org/10.1126/sciadv.adf0873>.
- [8] K. CUTAJAR, M. OSBORNE, J. CUNNINGHAM, AND M. FILIPPONE, *Preconditioning kernel matrices*, in Proceedings of the 33rd International Conference on Machine Learning, 2016, <https://proceedings.mlr.press/v48/cutajar16.html>.
- [9] A. DESHPANDE, L. RADEMACHER, S. S. VEMPALA, AND G. WANG, *Matrix approximation and projective clustering via volume sampling*, *Theory Comput.*, 2 (2006), pp. 225–247, <https://doi.org/10.4086/toc.2006.v002a012>.
- [10] A. DESHPANDE AND S. VEMPALA, *Adaptive sampling and fast low-rank matrix approximation*, in Proceedings of the 9th International Conference on Approximation Algorithms for Combinatorial Optimization Problems, and 10th International Conference on Randomization and Computation, 2006, pp. 292–303, https://doi.org/10.1007/11830924_28.
- [11] M. DÍAZ, E. N. EPPERLY, Z. FRANGELLA, J. A. TROPP, AND R. J. WEBBER, *Robust, Randomized Preconditioning for Kernel Ridge Regression*, preprint, 2024, <https://arxiv.org/abs/2304.12465v4>.
- [12] Y. DONG, C. CHEN, P.-G. MARTINSSON, AND K. PEARCE, *Robust blockwise random pivoting: Fast and accurate adaptive interpolative decomposition*, *SIAM J. Matrix Anal. Appl.*, 46 (2025), pp. 1791–1815, <https://doi.org/10.1137/24M1678027>.
- [13] J. DONGARRA, I. FOSTER, G. FOX, W. GROPP, K. KENNEDY, L. TORCZON, AND A. WHITE, eds., *Sourcebook of Parallel Computing*, Morgan Kaufmann Publishers, San Francisco, CA, 2003.
- [14] P. DRINEAS AND M. W. MAHONEY, *On the Nyström method for approximating a Gram matrix for improved kernel-based learning*, *J. Mach. Learn. Res.*, 6 (2005), pp. 2153–2175, <http://jmlr.org/papers/v6/drineas05a.html>.
- [15] E. N. EPPERLY AND E. MORENO, *Kernel quadrature with randomly pivoted Cholesky*, in Proceedings of the 37th International Conference on Neural Information Processing Systems, 2023, <https://dl.acm.org/doi/10.5555/3666122.3668997>.
- [16] S. FINE AND K. SCHEINBERG, *Efficient SVM training using low-rank kernel representations*, *J. Mach. Learn. Res.*, 2 (2002), pp. 243–264, <https://dl.acm.org/doi/10.5555/944790.944812>.
- [17] Z. FRANGELLA, J. A. TROPP, AND M. UDELL, *Randomized Nyström preconditioning*, *SIAM J. Matrix Anal. Appl.*, 44 (2023), pp. 718–752, <https://doi.org/10.1137/21M1466244>.

- [18] J. GARDNER, G. PLEISS, K. Q. WEINBERGER, D. BINDEL, AND A. G. WILSON, *GPyTorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration*, in Proceedings of the 32nd International Conference on Neural Information Processing Systems, 2018, <https://dl.acm.org/doi/10.5555/3327757.3327857>.
- [19] L. GIRAUD, J. LANGOU, M. ROZLOŽNÍK, AND J. V. D. ESHOF, *Rounding error analysis of the classical Gram–Schmidt orthogonalization process*, Numer. Math., 101 (2005), pp. 87–100, <https://doi.org/10.1007/s00211-005-0615-4>.
- [20] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 4th ed., Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins Press, Baltimore, 2013.
- [21] N. J. HIGHAM, *Analysis of the Cholesky decomposition of a semi-definite matrix*, in Reliable Numerical Computation, M. G. Cox and S. Hammarling, eds., Oxford University Press, 1990, <https://doi.org/10.1093/oso/9780198535645.003.0010>.
- [22] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., Society for Industrial and Applied Mathematics, Philadelphia, 2002.
- [23] M. KANAGAWA, P. HENNIG, D. SEJDINOVIC, AND B. K. SRIPERUMBUDUR, *Gaussian Processes and Kernel Methods: A Review on Connections and Equivalences*, preprint, 2018, <https://arxiv.org/abs/1807.02582v1>.
- [24] J. S. LIU, *Monte Carlo Strategies in Scientific Computing*, Springer New York, 2004, <https://doi.org/10.1007/978-0-387-76371-2>.
- [25] G. MEANTI, L. CARRATINO, L. ROSASCO, AND A. RUDI, *Kernel methods through the roof: Handling billions of points efficiently*, in Proceedings of the 34th International Conference on Neural Information Processing Systems, 2020, <https://dl.acm.org/doi/abs/10.5555/3495724.3496932>.
- [26] R. MURRAY, J. DEMMEL, M. W. MAHONEY, N. B. ERICHSOHN, M. MELNICHENKO, O. A. MALIK, L. GRIGORI, P. LUSZCZEK, M. DEREZIŃSKI, M. E. LOPES, T. LIANG, H. LUO, AND J. DONGARRA, *Randomized Numerical Linear Algebra: A Perspective on the Field with an Eye to Software*, preprint, 2023, <https://arxiv.org/abs/2302.11474v2>.
- [27] C. MUSCO AND C. MUSCO, *Recursive sampling for the Nyström method*, in Proceedings of the 31st International Conference on Neural Information Processing Systems, 2017, <https://dl.acm.org/doi/10.5555/3294996.3295140>.
- [28] 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, 2017, pp. 672–683, <https://doi.org/10.1109/FOCS.2017.68>.
- [29] A. REZAEI AND S. O. GHARAN, *A polynomial time MCMC method for sampling from continuous determinantal point processes*, in Proceedings of the 36th International Conference on Machine Learning, 2019, <https://proceedings.mlr.press/v97/rezaei19a.html>.
- [30] A. RUDI, D. CALANDRIELLO, L. CARRATINO, AND L. ROSASCO, *On fast leverage score sampling and optimal learning*, in Proceedings of the 32nd International Conference on Neural Information Processing Systems, Vol. 31, 2018, <https://dl.acm.org/doi/10.5555/3327345.3327470>.
- [31] B. SCHÖLKOPF AND A. J. SMOLA, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press, Cambridge, MA, 2018.
- [32] J. A. TROPP AND R. J. WEBBER, *Randomized Algorithms for Low-Rank Matrix Approximation: Design, Analysis, and Applications*, preprint, 2023, <https://arxiv.org/abs/2306.12418v3>.
- [33] O. T. UNKE, S. CHMIELA, H. E. SAUCEDA, M. GASTEGGER, I. POLTAVSKY, K. T. SCHÜTT, A. TKATCHENKO, AND K.-R. MÜLLER, *Machine learning force fields*, Chem. Rev., 121 (2021), pp. 10142–10186, <https://doi.org/10.1021/acs.chemrev.0c01111>.
- [34] K. WANG, G. PLEISS, J. GARDNER, S. TYREE, K. Q. WEINBERGER, AND A. G. WILSON, *Exact Gaussian processes on a million data points*, in Proceedings of the 33rd International Conference on Neural Information Processing Systems, 2019, <https://dl.acm.org/doi/10.5555/3454287.3455599>.
- [35] C. K. I. WILLIAMS AND M. SEEGER, *Using the Nyström method to speed up kernel machines*, in Proceedings of the 13th International Conference on Neural Information Processing Systems, 2000, <https://dl.acm.org/doi/10.5555/3008751.3008847>.
- [36] D. P. WOODRUFF, *Sketching as a tool for numerical linear algebra*, Found. Trends Theoret. Comput. Sci., 10 (2014), pp. 1–157, <https://doi.org/10.1561/04000000060>.
- [37] F. ZHANG, *Matrix Theory: Basic Results and Techniques*, 2nd ed., Springer, New York, 2011.