

FAST POLYNOMIAL TRANSFORMS BASED ON TOEPLITZ AND HANKEL MATRICES

ALEX TOWNSEND, MARCUS WEBB, AND SHEEHAN OLVER

ABSTRACT. Many standard conversion matrices between coefficients in classical orthogonal polynomial expansions can be decomposed using diagonally-scaled Hadamard products involving Toeplitz and Hankel matrices. This allows us to derive algorithms with an observed complexity of $\mathcal{O}(N \log^2 N)$, based on the fast Fourier transform, for converting coefficients of a degree N polynomial in one polynomial basis to coefficients in another. Numerical results show that this approach is competitive with state-of-the-art techniques, requires no precomputational cost, can be implemented in a handful of lines of code, and is easily adapted to extended precision arithmetic.

1. INTRODUCTION

Expansions of polynomials as finite series in orthogonal polynomial bases have applications throughout scientific computing, engineering, and physics [3, 10, 37]. The most popular expansions are in the Chebyshev and Legendre basis,

$$(1) \quad p_N(x) = \sum_{k=0}^N c_k^{cheb} T_k(x) = \sum_{k=0}^N c_k^{leg} P_k(x), \quad x \in [-1, 1],$$

where p_N is a degree N polynomial and T_k and P_k are the degree k Chebyshev and Legendre polynomials, respectively. Chebyshev expansions are used because of their near-optimal approximation properties and associated fast transforms [14, 26, 43] and Legendre expansions for their L^2 orthogonality [31, Table 18.3.1] as well as other recurrence relations that they satisfy [17]. A useful working paradigm is to represent a smooth function on a finite interval using the Chebyshev basis and to convert it to a different polynomial basis, such as Legendre, whenever it is algorithmically convenient to do so [17]. It is therefore important to have fast transforms for converting a vector of coefficients in one polynomial basis to another.

Given two sequences of orthogonal polynomials $\{\phi_0, \phi_1, \dots\}$ and $\{\psi_0, \psi_1, \dots\}$, there is an $(N+1) \times (N+1)$ upper-triangular *conversion matrix* (sometimes called

Received by the editor May 10, 2016, and, in revised form, November 20, 2016, and March 9, 2017.

2010 *Mathematics Subject Classification*. Primary 65T50, 65D05, 15B05.

Key words and phrases. Conversion matrix, Toeplitz, Hankel, Hadamard product.

The work of the first author was supported by the National Science Foundation grant No. 1522577.

The work of the second author was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/H023348/1 for the University of Cambridge Centre for Doctoral Training, the Cambridge Centre for Analysis, and the London Mathematical Society Cecil King Travel Scholarship 2015.

the connection coefficients matrix), A , such that

$$\underline{\beta} = A\underline{\alpha}, \quad p_N(x) = \sum_{k=0}^N \alpha_k \phi_k(x) = \sum_{k=0}^N \beta_k \psi_k(x),$$

where $\underline{\alpha} = (\alpha_0, \dots, \alpha_N)^T$ and $\underline{\beta} = (\beta_0, \dots, \beta_N)^T$. In this paper we describe how to compute the matrix-vector product $A\underline{\alpha}$ in an observed complexity of $\mathcal{O}(N \log^2 N)$ when $\{\phi_0, \phi_1, \dots\}$ and $\{\psi_0, \psi_1, \dots\}$ are sets of classical orthogonal polynomials such as Chebyshev, Legendre (see Sections 3 and 5.1), ultraspherical (see Section 5.2), Jacobi (see Section 5.3), and Laguerre (see Section 5.4).

There are many existing fast algorithms for computing $A\underline{\alpha}$ that exploit a variety of structures, including: (1) a hierarchical off-diagonal low rank structure [1], (2) an eigenvalue decomposition involving a semiseparable matrix [23, 24, 35], and (3) trigonometric-like behavior of orthogonal polynomials via asymptotic expansions [16, 28, 33, 39]. In this paper, we exploit a new observation that all standard conversion matrices can be decomposed using diagonally-scaled Hadamard products involving Toeplitz and Hankel matrices.¹ This structure allows us to derive FFT-based algorithms for computing $A\underline{\alpha}$ that can be implemented in many programming languages using just a handful of lines of code. We believe this algorithm is conceptually simpler than previous approaches, while being competitive in terms of computational speed (see Sections 4 and 5).

It is easy to see the structure that we exploit by considering an example. Consider the $(N+1) \times (N+1)$ Legendre-to-Chebyshev conversion matrix, which converts a vector of Legendre coefficients to Chebyshev coefficients for a degree N polynomial. Following [1, eqn. (2.18)] (see also [40, (4.9.4)]), it is given by

$$(2) \quad \underline{c}^{cheb} = M\underline{c}^{leg}, \quad M_{jk} = \begin{cases} \frac{1}{\pi} \Lambda\left(\frac{k}{2}\right)^2, & 0 = j \leq k \leq N, \text{ } k \text{ even,} \\ \frac{2}{\pi} \Lambda\left(\frac{k-j}{2}\right) \Lambda\left(\frac{k+j}{2}\right), & 0 < j \leq k \leq N, \text{ } k-j \text{ even,} \\ 0, & \text{otherwise,} \end{cases}$$

where² $\Lambda(z) = \Gamma(z+1/2)/\Gamma(z+1)$, $\Gamma(z)$ is the gamma function, and $\underline{c}^{cheb} = (c_0^{cheb}, \dots, c_N^{cheb})^T$ and $\underline{c}^{leg} = (c_0^{leg}, \dots, c_N^{leg})^T$ are the Chebyshev and Legendre coefficients of $p_N(x)$ in (1), respectively. The explicit formula in (2) reveals that after a trivial diagonal scaling, the matrix M can be written as a Hadamard product between an upper-triangular Toeplitz matrix, from the $\Lambda(\frac{k-j}{2})$ term, and a Hankel matrix, from the term $\Lambda(\frac{j+i}{2})$. Thus, for the matrix M in (2) we can write

$$(3) \quad M = D_1(T \circ H)D_2,$$

where D_1 and D_2 are diagonal matrices, T is an upper-triangular Toeplitz matrix, H is a Hankel matrix, and ‘ \circ ’ is the Hadamard matrix product, i.e., entrywise multiplication between two matrices (see (7) for explicit formulas for D_1 , D_2 , T and H). We find that the structure in (3) holds for many of the standard conversion matrices (see Section 5).

At first it is not obvious why the decomposition in (3) is useful for deriving a fast matrix-vector product because for general Toeplitz and Hankel matrices we are not

¹A Toeplitz matrix, T , is a matrix that is constant along each diagonal, i.e., $T_{jk} = t_{j-k}$. A Hankel matrix, H , is a matrix that is constant along each anti-diagonal, i.e., $H_{jk} = h_{j+k}$.

²We note that $\Lambda(z) = B(1/2, z+1/2)/\sqrt{\pi}$, where $B(\alpha, \beta)$ is the Beta function [31, (5.12.1)].

Outline for fast polynomial basis conversion: Computing $\underline{\beta} = A\underline{\alpha}$.	
	Cost
1. Decompose A into $A = D_1(T \circ H)D_2$ (see Section 5)	-
2. Calculate $H \approx \sum_{r=1}^K a_r \underline{\ell}_r \underline{\ell}_r^T$ (see Section 2.1)	$\mathcal{O}(N \log^2 N)$
3. Compute $\underline{v} = D_2 \underline{\alpha}$	$\mathcal{O}(N)$
4. Compute $\underline{w} = (T \circ H)\underline{v}$ using (5) and the FFT [29]	$\mathcal{O}(N \log^2 N)$
5. Compute $\underline{\beta} = D_1 \underline{w}$	$\mathcal{O}(N)$
	<hr/>
	$\mathcal{O}(N \log^2 N)$

FIGURE 1. A summary of the $\mathcal{O}(N \log^2 N)$ algorithm for converting between many classical orthogonal polynomial bases, including Chebyshev, Legendre, ultraspherical, Jacobi, and Laguerre polynomial bases. The algorithmic complexity of the second step is, more precisely, expected to be $\mathcal{O}(N \log^2 N \log(1/\epsilon))$ and numerical evidence is provided (see Figure 4) but a full theoretical justification is elusive. For the algorithmic complexity of the fourth step see Section 3.2.

aware of a fast algorithm for computing $(T \circ H)\underline{v}$. However, for conversion matrices the Hankel matrix in (3) is often real and positive semidefinite and hence, is severely ill-conditioned [4]. Moreover, Theorem 3 shows that all positive semidefinite Hankel matrices in (3) can be approximated, up to an accuracy of $0 < \epsilon < 1$, by a rank $K = \mathcal{O}(\log N \log(1/\epsilon))$ matrix. To give an example value for K in practice, if we take the H from the Legendre-to-Chebyshev matrix given in (2), with $N + 1 = 5000$ and $\epsilon = 10^{-14}$, then there exists a rank $K = 30$ matrix \hat{H} such that $\|H - \hat{H}\|_2 < \epsilon \|H\|_2$.

For an integer K , we construct a rank K approximation of H in $\mathcal{O}(K^2 N)$ operations using the pivoted Cholesky algorithm (see Section 2.1) to obtain the approximation,

$$(4) \quad H \approx \sum_{r=1}^K a_r \underline{\ell}_r \underline{\ell}_r^T.$$

Since K needs to be no larger than $\mathcal{O}(\log N \log(1/\epsilon))$ we can compute (4) in a total of $\mathcal{O}(N \log^2 N \log(1/\epsilon)^2)$ operations. Moreover, using $(A \circ \underline{\ell} \underline{\ell}^T)\underline{v} = (D_{\underline{\ell}} A D_{\underline{\ell}})\underline{v}$, where $D_{\underline{\ell}} = \text{diag}(\underline{\ell})$, we can write

$$(5) \quad (T \circ H)\underline{v} \approx \sum_{r=1}^K a_r \left(D_{\underline{\ell}_r} T D_{\underline{\ell}_r} \right) \underline{v},$$

where each term is a diagonally-scaled Toeplitz matrix whose matrix-vector product can be computed in $\mathcal{O}(N \log N)$ operations via the fast Fourier transform (FFT) [29, Sec. 3.4]. There are K terms in (5), so the matrix-vector product $(T \circ H)\underline{v}$ can be computed, up to an expected error of $\mathcal{O}(\epsilon)$, in $\mathcal{O}(KN \log N + K^2 N)$ operations. For the majority of this paper we write algorithmic complexities without the explicit dependency on ϵ . A more careful analysis of the error incurred in the approximation in (5) can be found in Section 3.3.

The approach for fast polynomial basis conversion (see Figure 1) requires no hierarchical data structures, no precomputational setup cost, and no matrix partitioning. The fundamental step is a fast Toeplitz matrix-vector product and the cost of the matrix-vector product is precisely $\mathcal{O}(K)$ FFTs of size $2N - 1$ [29, Sec. 3.4]. The fact that our algorithm relies on the FFT means that the implementation is automatically tuned to personal computer architectures, thanks to FFTW [12]. Our algorithm is now the default algorithm for polynomial basis conversion in the MATLAB software system called Chebfun [11] (see the commands `leg2cheb`, `cheb2leg`, `ultra2ultra`, and `jac2jac`) and the Julia package called ApproxFun.jl [30], via the package FastTransforms.jl [41] (see commands `leg2cheb` and `cheb2leg`).

More generally, matrix decompositions of the form $A \circ B$, where A has a known fast matrix-vector product and B can be approximated by a low rank matrix, are often useful for deriving fast matrix-vector products. The algorithmic challenge is to construct a low rank approximation to B in quasilinear complexity. Decompositions of the form $A \circ B$ have been exploited to derive a simple nonuniform FFT [2], a discrete Legendre transform [18], and a fast evaluation of exponential sums [21]. We hope that this list continues to grow in the future.

This paper is structured as follows. In Section 2 we describe a pivoted Cholesky algorithm for constructing low rank approximations of real, symmetric, and positive semidefinite matrices and use it to compute $(T \circ H)\underline{v}$ in $\mathcal{O}(N \log^2 N)$ operations, where T is an $(N + 1) \times (N + 1)$ Toeplitz matrix and H is a real, symmetric, and positive semidefinite Hankel matrix. In Section 3 we write $M = D_1(T \circ H)D_2$ as in (3) and show that the Hankel part, H , can be approximated, up to an error of $0 < \epsilon < 1$, by a rank $K = \mathcal{O}(\log N \log(1/\epsilon))$ matrix. In Section 4 we compare various algorithms for converting Legendre-to-Chebyshev basis conversion. In Section 5 we show that our algorithm allows for fast matrix-vector products involving many of the standard conversion matrices and we give numerical results throughout that section.

2. A FAST MATRIX-VECTOR PRODUCT FOR CERTAIN TOEPLITZ-DOT-HANKEL MATRICES

We say that a matrix A is a *Toeplitz-dot-Hankel* matrix if A can be written as a Hadamard product of a Toeplitz and Hankel matrix, i.e., $A = T \circ H$, where T is a Toeplitz matrix, H is a Hankel matrix, and ‘ \circ ’ denotes the Hadamard product. In this section, we suppose that H is a real, symmetric, and positive semidefinite Hankel matrix and that it is approximated, up to an error of $0 < \epsilon < 1$, by a rank K matrix. Later, in Section 3 we show that this holds for the Hankel part of the Legendre-to-Chebyshev conversion matrix in (2) when $K = \mathcal{O}(\log N \log(1/\epsilon))$.

Once we have constructed the low rank approximation in (4), costing $\mathcal{O}(K^2 N)$ operations, a fast matrix-vector product is immediate as (5) shows that $(T \circ H)\underline{v}$ can be computed as a sum of matrix-vector products involving a Toeplitz matrix.

2.1. Low rank approximation of a real positive semidefinite matrices using the pivoted Cholesky algorithm. Suppose that H is approximated up to an error of ϵ by a rank K matrix. We would like to compute a rank K approximation to H , i.e., a_1, \dots, a_K and $\underline{\ell}_1, \dots, \underline{\ell}_K$ in (4).

In principle, one could construct a best rank K approximation of H by computing the singular value decomposition (SVD) of H and taking the first K left and right singular vectors as well as the first K singular values. Naively, this costs a

prohibitive $\mathcal{O}(N^3)$ operations, which can be reduced to $\mathcal{O}(N^2 \log N)$ operations if the Hankel structure is exploited [46]. Instead, we describe an algorithm that costs only $\mathcal{O}(K^2 N)$ operations based on a pivoted Cholesky algorithm. It can be applied to any real, symmetric, and positive semidefinite matrix and does not exploit the Hankel structure of H . In practice, if the singular values of H decay rapidly, then the rank K approximation constructed by our pivoted Cholesky algorithm is observed to be near-best (see Figure 4). Our algorithm is very similar to the pivoted Cholesky algorithm described in [19], except we avoid square roots and have a different stopping criterion. One can trace back the origin of our algorithm to the rank revealing literature [15, 20].

Set $H^{(0)} = H$ and assume that H is a nonzero, real, symmetric, and positive semidefinite matrix. First, the maximum on the diagonal of H is selected, say $H_{p_1 p_1}$, which is also the absolute global maximum entry of H [13, Thm. 4.2.8]. Then, one step of the Cholesky algorithm is performed with the entry (p_1, p_1) as the pivot, i.e.,

$$(6) \quad H^{(1)} = H^{(0)} - H_{:,p_1}^{(0)} H_{p_1,:}^{(0)} / H_{p_1 p_1}^{(0)},$$

where $H_{:,p_1}^{(0)}$ and $H_{p_1,:}^{(0)}$ denotes the p_1 th column and p_1 th row of $H^{(0)}$, respectively. The matrix $H^{(1)}$ has a zero p_1 th column and p_1 th row. The Cholesky step in (6) is closely related to the Schur complement of the (p_1, p_1) entry in $H^{(0)}$ and the matrix $H^{(1)}$ is guaranteed to be real, symmetric, and positive semidefinite. Furthermore, the rank of $H^{(1)}$ is exactly one less than $H^{(0)}$ [27, Cor. 19.2].

Next, the maximum on the diagonal of $H^{(1)}$ is found, say $H_{p_2 p_2}^{(1)}$. If $H_{p_2 p_2}^{(1)} = 0$ or is sufficiently small, then the process is terminated; otherwise, another Cholesky step is taken with the entry (p_2, p_2) as the pivot, i.e.,

$$H^{(2)} = H^{(1)} - H_{:,p_2}^{(1)} H_{p_2,:}^{(1)} / H_{p_2 p_2}^{(1)}.$$

Again, the matrix $H^{(2)}$ is guaranteed to be real, symmetric, and positive semidefinite and has a rank that is exactly one less than $H^{(1)}$. The pivoted Cholesky algorithm continues until the maximum on diagonal is either zero or sufficiently small. Since the rank of H is at most N and the rank decreases by precisely one after each Cholesky step, the algorithm terminates in at most N steps. For the algorithm to be computationally more efficient than the SVD one hopes to need just $K \ll N$ steps.

Suppose that the pivoted Cholesky algorithm takes K steps before terminating. Since $H^{(0)}, \dots, H^{(K-1)}$ are symmetric matrices, we can write the r th Cholesky step as follows:

$$H^{(r)} = H^{(r-1)} - a_r \underline{\ell}_r \underline{\ell}_r^T, \quad \underline{\ell}_r = H_{:,p_r}^{(r-1)}, \quad a_r = (H_{p_r p_r}^{(r-1)})^{-1}.$$

Therefore, we use the r th pivoted Cholesky step to construct the r th term in (4), where after K steps the rank K approximation to H is constructed.

The pivoted Cholesky algorithm described so far requires a total of $\mathcal{O}(KN^2)$ operations because at each step an $(N + 1) \times (N + 1)$ matrix is updated. Now we will describe how to construct the same rank K approximant to H in $\mathcal{O}(K^2 N)$ operations, which is a significant saving when $K \ll N$.

The main idea to reduce the computational cost is to note that it is not necessary to update the whole matrix at each Cholesky step. For example, consider the Hankel




FIGURE 2. The pivoted Cholesky algorithm is applied to the Hankel part of the Legendre-to-Chebyshev conversion matrix in (2) when $N = 300$. The Cholesky algorithm takes 27 pivots (blue dots) and evaluates the matrix at the 27 corresponding columns (black lines). One can see that the majority of the entries in the Hankel part of M are not required during the pivoted Cholesky algorithm, allowing for the computational cost to be reduced from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$ operations.

part, H , of the matrix M in (2) for $N = 300$. Figure 2 shows the pivot locations selected by the pivoted Cholesky algorithm, where a total of 27 steps was required to construct a low rank approximant of H that is accurate up to double precision. This means that only 27 columns from H are required to compute (4) (see black vertical lines on Figure 2). Therefore, we rewrite the pivoted Cholesky algorithm so that it only updates the diagonal entries (required to determine the pivot locations) and those 27 columns of H . This allows for a significant computational saving when $K \ll N$.

Let $\underline{d} = (H_{00}, \dots, H_{NN})^T$ be the diagonal entries of H . In the first step, instead of (6), we only update the diagonal as follows:

$$\underline{d} \leftarrow \underline{d} - (H_{:,p_1}^{(0)})^2 / H_{p_1 p_1}^{(0)},$$

where we have used the fact that $H^T = H$ and $(H_{:,p_1}^{(0)})^2$ is the vector $H_{:,p_1}^{(0)}$ squared entry-by-entry. The diagonal can then be used to determine the location of the second pivot. For the second step, we again only update the diagonal,

$$\underline{d} \leftarrow \underline{d} - (H_{:,p_2}^{(1)})^2 / H_{p_2 p_2}^{(1)},$$

where $(H_{:,p_2}^{(1)})^2$ means that the vector $H_{:,p_2}^{(1)}$ is squared entry-by-entry. Since the pivot locations at each step are determined by the diagonal entries, one can select the pivoting entries by only updating the diagonal vector \underline{d} . At the r th Cholesky step the column $H_{:,p_r}^{(r-1)}$ is required, which is not directly available from the matrix H .

<p>Pivoted Cholesky algorithm for real, symmetric, and positive semi-definite matrices</p> <p>Input: A real, symmetric, and positive semidefinite matrix, H, and an accuracy goal $0 < \epsilon < 1$.</p> <p>Output: Vectors $\underline{\ell}_1, \dots, \underline{\ell}_K$ and a_1, \dots, a_K such that $H \approx \sum_{r=1}^K a_r \underline{\ell}_r \underline{\ell}_r^T$.</p> <p>Set $\underline{d} = (H_{00}, \dots, H_{NN})^T$, the diagonal of H</p> <p>for $r = 1, 2, \dots, N$</p> <p style="padding-left: 2em;">$p_r = \arg \max_{0 \leq j \leq N} d_j$</p> <p style="padding-left: 2em;">Set $\underline{\ell}_r = H_{:,p_r}$</p> <p style="padding-left: 2em;">for $j = 1, \dots, r - 1$</p> <p style="padding-left: 4em;">$\underline{\ell}_r \leftarrow \underline{\ell}_r - \underline{\ell}_j (\underline{\ell}_j)_{p_r} / (\underline{\ell}_j)_{p_j}$</p> <p style="padding-left: 2em;">end</p> <p style="padding-left: 2em;">Set $a_r = 1 / (\underline{\ell}_r)_{p_r}$</p> <p style="padding-left: 2em;">$\underline{d} \leftarrow \underline{d} - a_r (\underline{\ell}_r)^2$</p> <p style="padding-left: 2em;">if $\max_{0 \leq j \leq N} d_j \leq \epsilon \max_{0 \leq j \leq N} H_{jj}$, then stop</p> <p>end</p>

FIGURE 3. Pseudocode for pivoted Cholesky algorithm, which costs $\mathcal{O}(K^2N)$ operations. This algorithm only requires the matrix H to be real, symmetric, and positive semidefinite, not necessarily Hankel.

We calculate this by first constructing $H_{:,p_r}$ and by applying each of the previous $r - 1$ Cholesky steps to $H_{:,p_r}$ (see also [19]). Figure 3 presents a summary of the algorithm. A simple operation count reveals that the algorithm costs $\mathcal{O}(K^2N)$ operations.

One may be legitimately concerned that the low rank approximants constructed by the pivoted Cholesky algorithm in Figure 3 are of poor quality, as they are not strictly the best low rank approximants. More precisely, suppose that H_K is the best rank K approximant of H in the matrix 2-norm computed via the SVD and $\tilde{H}_K = \sum_{r=1}^K a_r \underline{\ell}_r \underline{\ell}_r^T$ is constructed via the pivoted Cholesky algorithm. Loosely speaking, is it possible that $\|H - H_K\|_2$ is much smaller than $\|H - \tilde{H}_K\|_2$? When H has moderately decaying singular values, we believe not. One representative numerical experiment is shown in Figure 4 (left), where the low rank approximants constructed by the pivoted Cholesky algorithm are compared against those from the SVD for the Hankel part of the matrix M in (2) when $N = 1000$. All other investigations have revealed similar results [19, 42]. The best mathematical statement we know of is [19, Thm. 3.2], which shows that $\|H - \tilde{H}_K\|_2 \leq 4^K(N + 1)\|H - H_K\|_2$. This suggests that the pivoted Cholesky procedure could compute low rank approximants of H that are severely suboptimal. A more descriptive inequality between $\|H - \tilde{H}_K\|_2$ and $\|H - H_K\|_2$ that adequately describes the power of the pivoted Cholesky algorithm for constructing low rank approximants is a remaining mathematical challenge and may require a much deeper understanding on the numerical stability of Gaussian elimination (see [42, Chap. 4]). In principle, we could have presented a variant on the pivoted Cholesky algorithm to ensure that it is a so-called strong

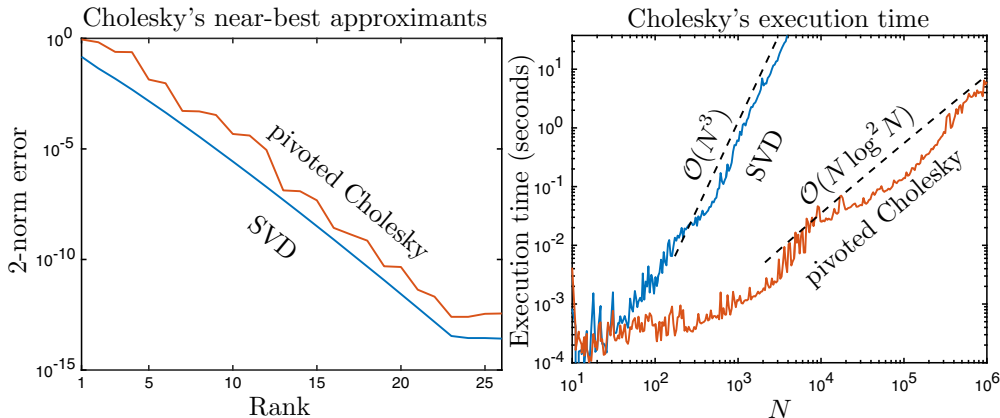


FIGURE 4. Left: The low rank approximants constructed by the pivoted Cholesky are often near-best. Here, the Hankel part of M in (2) is considered when $N = 1000$ and the 2-norm error between the SVD and pivoted Cholesky rank K approximants is compared. Right: The pivoted Cholesky algorithm is computationally more efficient than the SVD at constructing low rank approximants. Here, the execution times are compared for the Hankel part of M when $10 \leq N \leq 10^6$. At each N the algorithms construct a low rank approximant that is accurate to essentially double precision. Since $K = \mathcal{O}(\log N)$ for the Hankel part of M , the cost of the pivoted Cholesky algorithm is $\mathcal{O}(N \log^2 N)$ operations, while the cost of the SVD is $\mathcal{O}(N^3)$ operations.

rank-revealing algorithm [15]. However, this makes the pivoted Cholesky algorithm have a cost of $\mathcal{O}(N^2)$ operations and such a modification seems unnecessary in practice.

Our pivoted Cholesky algorithm costs $\mathcal{O}(K^2 N)$ operations. In Figure 4 (right) we compare the execution time for computing the SVD and the pivoted Cholesky algorithm on the Hankel part of the matrix M in (2) for obtaining an accuracy of essentially double precision. One can see that even when $N = 10^6$, the pivoted Cholesky algorithm can be employed to construct (4), whereas the SVD is limited to $N \leq 5000$. In some applications the size of the matrix M is fixed and one wishes to convert between two polynomial bases for many different polynomials of the same degree. In such a situation, the pivoted Cholesky algorithm can construct a low rank approximation for the Hankel part just once and it can be reused for each matrix-vector product.

3. THE HANKEL PART OF THE LEGENDRE-TO-CHEBYSHEV MATRIX

Let N be an integer. The $(N + 1) \times (N + 1)$ Legendre-to-Chebyshev matrix, denoted by M , in (2) can be written as a diagonally-scaled Toeplitz-dot-Hankel matrix. That is, $M = D_1(T \circ H)D_2$, where $D_1 = \frac{2}{\pi} \text{diag}(\frac{1}{2}, 1, \dots, 1)$, $D_2 = I_{N+1}$

the $(N + 1) \times (N + 1)$ identity matrix, and

$$(7) \quad \begin{aligned} T_{jk} &= \begin{cases} \Lambda\left(\frac{k-j}{2}\right), & 0 \leq j \leq k \leq N, k - j \text{ even,} \\ 0, & \text{otherwise,} \end{cases} \\ H_{jk} &= \Lambda\left(\frac{j+k}{2}\right), \quad 0 \leq j, k \leq N. \end{aligned}$$

Here $\Lambda(z) = \Gamma(z + 1/2)/\Gamma(z + 1)$, where $\Gamma(z)$ is the gamma function. In this section we show that H is: (1) real, symmetric, and positive semidefinite (see Section 3.1) and (2) H can be approximated, up to an accuracy of $0 < \epsilon < 1$, by a rank $K = \mathcal{O}(\log N \log(1/\epsilon))$ matrix (see Section 3.2).

3.1. The Hankel part of the Legendre-to-Chebyshev matrix is positive semidefinite. The Hankel matrix H in (7) is immediately seen to be real and symmetric. To show that it is positive semidefinite, we recall that the Hamburger moment problem states that a real Hankel matrix is positive semidefinite if and only if it is associated to a nonnegative Borel measure supported on the real line.

Lemma 1 (Hamburger moment problem). *A real $(N + 1) \times (N + 1)$ Hankel matrix, H , is positive semidefinite if and only if there exists a nonnegative Borel measure μ_H supported on the real line such that*

$$(8) \quad H_{jk} = \int_{-\infty}^{\infty} x^{j+k} d\mu_H(x), \quad 0 \leq j, k \leq N.$$

Proof. For a proof, see [34, Thm. 7.1]. □

We show that the Hankel matrix in (7) is positive semidefinite by expressing its entries in the form of (8).

Theorem 2. *The Hankel matrix, H , in (7) is positive semidefinite.*

Proof. It can be verified that

$$H_{jk} = \Lambda\left(\frac{j+k}{2}\right) = \frac{1}{\sqrt{\pi}} B\left(\frac{j+k+1}{2}, \frac{1}{2}\right) = \frac{1}{\sqrt{\pi}} \int_0^1 t^{\frac{j+k+1}{2}} (1-t)^{-\frac{1}{2}} dt, \quad 0 \leq j, k \leq N,$$

where $B(\cdot, \cdot)$ is the Beta function [31, (5.12.1)]. The last equality follows directly from the integral representation for the Beta function [31, (5.12.1)]. By substituting $t = x^2$, we arrive at the formula

$$H_{jk} = \frac{2}{\sqrt{\pi}} \int_0^1 x^{j+k} (1 - x^2)^{-\frac{1}{2}} dx, \quad 0 \leq j, k \leq N.$$

By setting $d\mu_H(x) = (2/\sqrt{\pi})\chi_{(0,1)}(1 - x^2)^{-1/2} dx$ in Lemma 1, where $\chi_{(0,1)}$ is the characteristic function for the interval $(0, 1)$, we conclude that H is a positive semidefinite matrix. □

Theorem 2 shows that H is positive semidefinite and therefore, the pivoted Cholesky algorithm described in Section 2.1 is applicable. We now show that H can be well-approximated by a rank K matrix where $K \ll N$.

3.2. The Hankel part of the Legendre-to-Chebyshev matrix is numerically of low rank. In Section 3.1 we showed that H in (7) is real, symmetric, and positive semidefinite. Such Hankel matrices are severely ill-conditioned, and Beckermann has proved the remarkably tight bound of $\kappa_2(H) \geq 0.0037 \times (33.97)^N / N$ [4], where $\kappa_2(H)$ denotes the condition number of H in the matrix 2-norm. This shows that H is in fact exponentially ill-conditioned, so one might expect that the singular values $\sigma_1(H), \dots, \sigma_{N+1}(H)$ of H decay rapidly to zero. Indeed, they do, as the following theorem shows.

Theorem 3. *Let N be an integer and H an $(N+1) \times (N+1)$ real positive definite Hankel matrix. Then, its singular values rapidly decay:*

$$\sigma_{j+2k}(H) \leq 16 \left[\exp \left(\frac{\pi^2}{4 \log(8 \lfloor (N+1)/2 \rfloor / \pi)} \right) \right]^{-2k+2} \sigma_j(H), \quad 1 \leq j+2k \leq N+1.$$

Moreover, for any $0 < \epsilon < 1$, there exists a rank K matrix \hat{H} such that $\|H - \hat{H}\|_2 \leq \epsilon \|H\|_2$, where

$$(9) \quad K = 2 \left\lceil \frac{2 \log(8 \lfloor (N+1)/2 \rfloor / \pi) \log(16/\epsilon)}{\pi^2} \right\rceil + 2 = \mathcal{O}(\log N \log(1/\epsilon)).$$

Proof. See [6, Sec. 5.2]. □

While the result in [6] is stated for real positive definite Hankel matrices, it immediately includes positive semidefinite Hankel matrices too. This is because a positive semidefinite Hankel matrix can be arbitrarily approximated by a positive definite Hankel matrix. This means that all the Hankel parts of the conversion matrices in this paper have a numerical rank of at most $\mathcal{O}(\log N \log(1/\epsilon))$, where the implicit constants are known from (9).

In practice, we use the pivoted Cholesky algorithm (see Section 2.1) to construct a low rank approximant for H in $\mathcal{O}(K^2 N)$ operations. Using the formula in (5) we can then calculate $\underline{c}^{cheb} = M \underline{c}^{leg}$ via a sum of $\mathcal{O}(K)$ diagonally-scaled Toeplitz matrix-vector products. Hence, we have described an $\mathcal{O}(N \log^2 N)$ algorithm for computing $\underline{c}^{cheb} = M \underline{c}^{leg}$ (see Figure 1).

3.3. Error analysis for the Legendre-to-Chebyshev matrix. We now investigate the error incurred in the matrix-vector product resulting from approximating a matrix of the form $D_1(T \circ H)D_2$ by $D_1(T \circ \hat{H})D_2$. Our analysis below shows that it is a valid approximation for the Legendre-to-Chebyshev matrix, but a more careful analysis is required for some of the other conversion matrices.

Theorem 4. *Let $0 < \epsilon < 1$ and set K be the integer in (9). Let \hat{H} be the rank K approximation to H so that $\|H - \hat{H}\|_2 \leq \epsilon \|H\|_2$. Then, for any $(N+1) \times 1$ vector we have*

$$\|D_1(T \circ H - T \circ \hat{H})D_2 \underline{v}\|_2 \leq \|D_1\|_{\max} \|D_2\|_{\max} \|T\|_{\max} 16\sqrt{2} \frac{\rho}{(\rho^2 - 1)^{1/2}} \epsilon \|H\|_2 \|v\|_2,$$

where $\|A\|_{\max}$ denotes the maximum absolute entry of A and ρ is given in (10).

Proof. We first derive a bound on $\|H - \hat{H}\|_F$, where $\|\cdot\|_F$ denotes the Frobenius matrix norm. From the definition of the Frobenius norm, we find that

$$\|H - \hat{H}\|_F^2 = \sum_{k=K+1}^N \sigma_k(H)^2 = \sigma_{K+1}(H)^2 \sum_{k=0}^{N-K-1} \frac{\sigma_{(K+1)+k}(H)^2}{\sigma_{K+1}(H)^2}.$$

To use Theorem 3, we note that $\sigma_{k+1}(H) \leq \sigma_k(H)$ so that

$$\|H - \hat{H}\|_F^2 \leq 2\sigma_{K+1}(H)^2 \sum_{k=0}^{\infty} \frac{\sigma_{(K+1)+2k}(H)^2}{\sigma_{K+1}(H)^2},$$

where $\sigma_j(H) = 0$ if $j > N$. From Theorem 3 we conclude that

(10)

$$\|H - \hat{H}\|_F^2 \leq 256 \sum_{k=0}^{\infty} \rho^{-2k+2} \sigma_{K+1}(H)^2, \quad \rho = \exp\left(\frac{\pi^2}{4 \log(8 \lfloor (N+1)/2 \rfloor / \pi)}\right).$$

Moreover, since $\sigma_{K+1}(H) \leq \epsilon \|H\|_2$, we obtain the bound

$$\|H - \hat{H}\|_F \leq 16\sqrt{2} \frac{\rho}{(\rho^2 - 1)^{1/2}} \epsilon \|H\|_2.$$

Since we have

$$\left\| D_1(T \circ H - T \circ \hat{H}) D_2 \underline{v} \right\|_2 \leq \|D_1\|_2 \|D_2\|_2 \|T \circ H - T \circ \hat{H}\|_2 \|\underline{v}\|_2,$$

the result follows by noting that $\|D_1\|_2 = \|D_1\|_{\max}$, $\|D_2\|_2 = \|D_2\|_{\max}$, and

$$\|T \circ H - T \circ \hat{H}\|_2 \leq \|T \circ H - T \circ \hat{H}\|_F \leq \|T\|_{\max} \|H - \hat{H}\|_F. \quad \square$$

For any $N \leq 10^8$, we find that $\rho/(\rho^2 - 1)^{1/2} \leq 2$. Moreover, in the special case of the Legendre-to-Chebyshev conversion matrix (see (7)), we have $\|D_1\|_{\max} = 1$, $\|D_2\|_{\max} = 1$, $\|T\|_{\max} = \sqrt{\pi}$ and $\|H\|_2 \leq 2\sqrt{N+1}$. Thus,

$$(11) \quad \left\| D_1(T \circ H - T \circ \hat{H}) D_2 \underline{v} \right\|_2 \leq 64\sqrt{2\pi(N+1)} \epsilon \|v\|_2.$$

In Figure 5 we observe an absolute error of $\mathcal{O}(\sqrt{N} \log N)$ in our transform when \underline{v} has entries drawn from independent identically distributed Gaussians. Therefore, we observe a better error than that predicted by (11) since $\|v\|_2$ is expected to be of $\mathcal{O}(\sqrt{N})$ in size. When the vector \underline{v} has decaying entries, as is common when the vector contains coefficients derived from expansion coefficients of a smooth function, we observe (see Figure 5) that the error in our Legendre-to-Chebyshev transform is essentially independent of N .

4. NUMERICAL RESULTS FOR LEGENDRE-TO-CHEBYSHEV CONVERSION

All the numerical results were performed on a 3.1 GHz Intel Core i7 Macbook Pro 2015 with MATLAB 2015b or Julia v0.4.5 [7]. In these numerical experiments we employ three different algorithms for computing the matrix-vector product $\underline{c}^{cheb} = M \underline{c}^{leg}$:

- **Direct:** The direct algorithm computes \underline{c}^{cheb} by first constructing the $(N+1) \times (N+1)$ matrix M one row at a time and then calculating the dot product with \underline{c}^{leg} . Therefore, the vector \underline{c}^{cheb} is computed entry-by-entry, costing a total of $\mathcal{O}(N^2)$ operations and requiring $\mathcal{O}(N)$ storage.

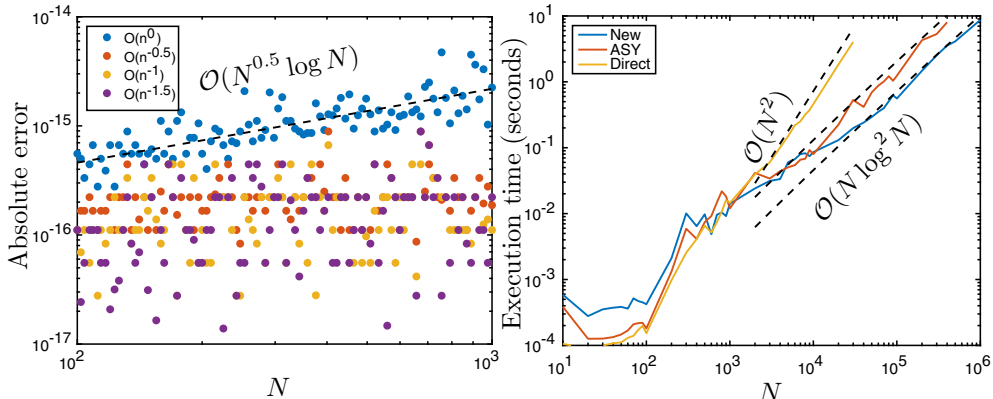


FIGURE 5. Left: Errors in computing $\underline{c}^{cheb} = M\underline{c}^{leg}$ with various decay rates in \underline{c}^{leg} . A vector is created using `randn(N)` in Julia and then the vector is scaled so the n th entry is $\mathcal{O}(n^0)$, $\mathcal{O}(n^{-0.5})$, $\mathcal{O}(n^{-1})$, $\mathcal{O}(n^{-1.5})$. The dashed line shows the observed error growth in the case when \underline{c}^{leg} does not decay. Right: Execution times between the direct (yellow), ASY (red), and the new algorithm described in this paper (blue). The new algorithm presented in this paper is about 2 or 3 times faster than ASY, while being more accurate and only requiring a handful of lines of code to implement.

- **ASY:** The algorithm that we call ASY here is described in [16]. It computes the matrix-vector product $\underline{c}^{cheb} = M\underline{c}^{leg}$ in $\mathcal{O}(N \log^2 N / \log \log N)$ operations by using a trigonometric asymptotic formula for Legendre polynomials. Before this paper, it was the algorithm employed in the `leg2cheb` command in Chebfun [11].
- **New:** The algorithm described in this paper. It is summarized in Figure 1, costing $\mathcal{O}(N \log^2 N)$ operations.

Other algorithms for computing the matrix-vector product $\underline{c}^{cheb} = M\underline{c}^{leg}$ in fewer than $\mathcal{O}(N^2)$ operations are given in the pioneering paper by Alpert and Rokhlin [1] as well as [24, 33].

As a first test we take arbitrarily distributed vectors \underline{c}^{leg} with various rates of decay and consider the accuracy of our algorithm described in this paper against an extended precision computation (performed using the `BigFloat` type in Julia).³ Figure 5 (left) shows the absolute maximum errors in the computed vectors \underline{c}^{cheb} for $10 \leq N \leq 10^4$. In [16, Fig. 5.1] analogous errors were calculated for the direct and ASY algorithms. In Table 1 we summarise the observed error growth in the absolute maximum for the three different algorithms. In many applications the Legendre expansion in (1) represents a polynomial interpolant of a smooth function. In this setting, if the function is Hölder continuous with parameter greater than 0, then we observe that our new algorithm has essentially no error growth with N .

³In particular, the vector corresponding to, say, $N = 100$ with $\mathcal{O}(n^{-1})$ decay can be reproduced exactly by the Julia code `srand(0); c = randn(101)./(1:101)`. The fixed random seed is used for the sake of reproducibility.

TABLE 1. Observed absolute maximum error growth in computing $\underline{c}^{cheb} = M\underline{c}^{leg}$ when \underline{c}^{leg} has different decay rates. The error growth of the ASY algorithm given in [16, Fig. 5.1] were computed using the implementation of `leg2cheb` found in version 5.3 of Chebfun [11]. These are worse than those presented here because the ASY algorithm now uses Reinsch’s modification of the Legendre recurrence as recommended in [39]. We observe that if the Legendre expansion in (1) comes from a polynomial interpolant of a Hölder continuous function with parameter > 0 , then our new algorithm has essentially no error growth.

	Decay of the vector \underline{c}^{leg}			
	$\mathcal{O}(1)$	$\mathcal{O}(n^{-0.5})$	$\mathcal{O}(n^{-1})$	$\mathcal{O}(n^{-1.5})$
Direct	$\mathcal{O}(N^{1.5}/\log N)$	$\mathcal{O}(N/\log N)$	$\mathcal{O}(N^{0.5}/\log N)$	$\mathcal{O}(\log N)$
ASY	$\mathcal{O}(N)$	$\mathcal{O}(N^{0.5})$	$\mathcal{O}(\log N)$	$\mathcal{O}(1)$
New	$\mathcal{O}(N^{0.5} \log N)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$

For a second test, in Figure 5 (right) we compare the execution times for the three algorithms. Despite the direct algorithm requiring $\mathcal{O}(N^2)$ operations, it is computationally more efficient when $N \leq 512$. The new algorithm presented here is 2 or 3 times faster than the ASY algorithm for large N , while being conceptually simpler and more accurate for nondecaying vectors \underline{c}^{leg} . Based on these numerical experiments, the `leg2cheb` command in Chebfun [11] and the `leg2cheb` in `Fast-Transforms.jl` [41] use the direct algorithm when $N \leq 512$ and the new algorithm otherwise.

5. OTHER POLYNOMIAL BASIS CONVERSIONS

So far the paper has focused on the task of converting Legendre coefficients for p_N in (1) to Chebyshev coefficients. In this section we consider other standard polynomial basis conversions, showing how our $\mathcal{O}(N \log^2 N)$ algorithm summarized in Figure 1 remains applicable.

5.1. Chebyshev-to-Legendre conversion. To compute the Legendre coefficients of a given polynomial p_N in fewer than $\mathcal{O}(N^2)$ operations, one can first compute the Chebyshev coefficients using the discrete cosine transform (DCT) of its values at Chebyshev points in $\mathcal{O}(N \log N)$ operations [14], then use a fast Chebyshev-to-Legendre conversion. Alternatively, a direct transform taking values of the polynomial in the complex plane to Legendre coefficients is given in [22] and a fast transform for converting values of the polynomial at Legendre points to Legendre coefficients is given in [18].

The inverse of the Legendre-to-Chebyshev matrix M^{-1} , denoted by L in [1], converts Chebyshev coefficients to Legendre coefficients, i.e., $\underline{c}^{leg} = L\underline{c}^{cheb}$ in (1).

Explicit formulas for the entries of $L = M^{-1}$ are given as follows [1, (2.19)]:

$$(12) \quad L_{jk} = \begin{cases} 1, & j = k = 0, \\ \frac{\sqrt{\pi}}{2\Lambda(j)}, & 0 < j = k \leq N, \\ -k(j + \frac{1}{2}) \left(\frac{\Lambda(\frac{k-j-2}{2})}{k-j} \right) \left(\frac{\Lambda(\frac{j+k-1}{2})}{j+k+1} \right), & 0 \leq j < k \leq N, \quad k - j \text{ even,} \\ 0, & \text{otherwise,} \end{cases}$$

where $\Lambda(z) = \Gamma(z + 1/2)/\Gamma(z + 1)$ and $\Gamma(z)$ is the gamma function. The fact that L can be written as $L = D_1(T \circ H)D_2$ is almost immediate from (12). In particular, we have $D_1 = \text{diag}(\frac{1}{2}, \frac{3}{2}, \frac{5}{2}, \dots, \frac{2N+1}{2})$, $D_2 = -\frac{1}{4}\text{diag}(\frac{4}{\sqrt{\pi}}, 1, 2, 3, \dots, N)$, and

$$T_{jk} = \begin{cases} \Gamma\left(\frac{k-j-1}{2}\right) / \Gamma\left(\frac{k-j+2}{2}\right), & 0 \leq j \leq k \leq N, \quad k - j \text{ even,} \\ 0, & \text{otherwise,} \end{cases}$$

$$H_{jk} = \begin{cases} \Gamma\left(\frac{j+k}{2}\right) / \Gamma\left(\frac{j+k+3}{2}\right), & 0 \leq j, k \leq N, \quad j + k > 0, \\ 1, & j = k = 0. \end{cases}$$

Unfortunately, the matrix H is not positive semidefinite. This turns out not to matter, because the submatrix $\tilde{H} = (H_{jk})_{1 \leq j, k \leq N}$ is positive semidefinite by the following identity (see Lemma 1),

$$H_{jk} = \frac{4}{\sqrt{\pi}} \int_0^1 x^{j+k-1} (1-x^2)^{1/2} dx, \quad 1 \leq j, k \leq N,$$

Moreover, Theorem 3 shows that \tilde{H} can be approximated, up to an error of $0 < \epsilon < 1$, by a rank $\mathcal{O}(\log N \log(1/\epsilon))$ matrix. Hence, when computing $\underline{c}^{leg} = L\underline{c}^{cheb}$, we compute the first entry of \underline{c}^{leg} directly, and use the algorithm described in Figure 1 on $(L_{jk})_{1 \leq j, k \leq N}$ to compute the remaining entries. The resulting algorithm is implemented in the `cheb2leg` commands in Chebfun [11] and FastTransforms.jl [41].

In Figure 6 we repeat the same experiments as for the Legendre-to-Chebyshev conversion in Section 4. In Figure 6 (left) we compute the maximum error of the resulting vector \underline{c}^{leg} for different decay rates in \underline{c}^{cheb} . Due to the $\mathcal{O}(N)$ growth in the entries of D_1 , we find that the conversion requires a decay faster than $\mathcal{O}(n^{-1})$ in \underline{c}^{cheb} to have essentially no error growth. This holds when p_N in (1) is a Chebyshev interpolant of a Hölder continuous function with Hölder parameter $> 1/2$. The observed error growth is less than that observed for the Chebyshev-to-Legendre algorithm in [16, Fig. 5.2].

In Figure 6 (right) we show the execution times of the three algorithms: (1) Direct, an algorithm that costs $\mathcal{O}(N^2)$ operations and requires $\mathcal{O}(N)$ memory based on generating the whole matrix L one row at a time, (2) ASY, an $\mathcal{O}(N \log^2 N / \log \log N)$ complexity algorithm described in [16, Sec. 4], and (3) New, the algorithm described in this paper (see Figure 1). Our new algorithm is faster than direct when $N \geq 1,000$ and is about 2 or 3 times faster than the ASY algorithm for large N .

5.2. Ultraspherical-to-ultraspherical conversion. The ultraspherical polynomial of degree k and parameter $\lambda > 0$ is denoted by $C_k^{(\lambda)}$ [31, Tab. 18.3.1]. If

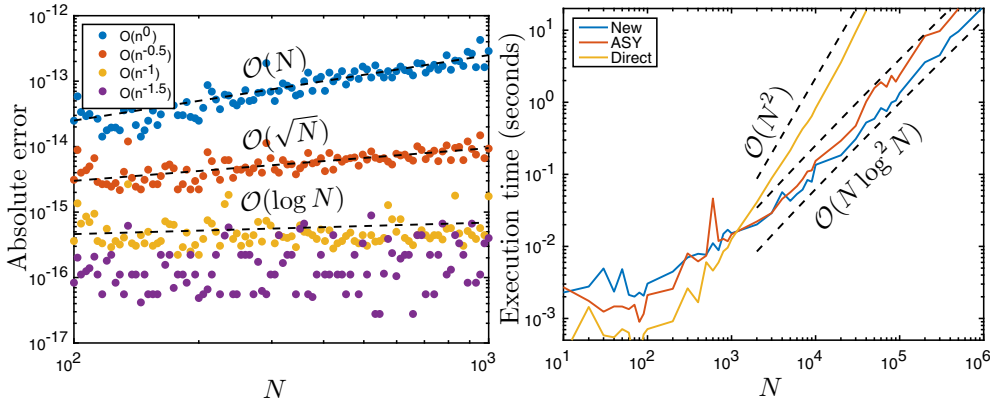


FIGURE 6. Left: Errors in computing $\underline{c}^{leg} = L\underline{c}^{cheb}$ with various decay rates in \underline{c}^{cheb} . Arbitrary vectors are generated in precisely the same manner as in Section 4. The dashed lines show the observed error growth for the various decay rates. Right: Execution times between the direct (yellow), ASY (red), and the new algorithm described in this paper (blue).

$\lambda > 0$, then $\{C_0^{(\lambda)}, C_1^{(\lambda)}, \dots\}$ is a family of orthogonal polynomials that are orthogonal with respect to the weight function $(1 - x^2)^{\lambda-1/2}$ on $[-1, 1]$. Ultraspherical polynomials can be seen as a generalization of Legendre polynomials since $P_k(x) = C_k^{(1/2)}(x)$ [31, (18.7.9)]. For papers on computing and converting ultraspherical expansions see, for example, [8, 23].

Let $\lambda_1 > 0$ and $\lambda_2 > 0$. The degree N polynomial p_N in (1) can be expanded in the two ultraspherical polynomial bases associated to λ_1 and λ_2 , i.e.,

$$p_N(x) = \sum_{k=0}^N c_k^{(\lambda_1)} C_k^{(\lambda_1)}(x) = \sum_{k=0}^N c_k^{(\lambda_2)} C_k^{(\lambda_2)}(x).$$

There is an upper-triangular conversion matrix, A , such that $\underline{c}^{(\lambda_2)} = A\underline{c}^{(\lambda_1)}$. We desire a fast algorithm for computing the matrix-vector product $\underline{c}^{(\lambda_2)} = A\underline{c}^{(\lambda_1)}$. There are several cases to consider: (1) $\lambda_1 - \lambda_2$ is an integer, (2) $|\lambda_1 - \lambda_2| < 1$, and (3) $|\lambda_1 - \lambda_2| > 1$, but the difference is not an integer. Despite the same formulas for the entries of the conversion matrix in all cases, we must consider each case separately because the Hankel part of the conversion matrix is only positive semidefinite when $|\lambda_1 - \lambda_2| < 1$.

Case 1: $\lambda_1 - \lambda_2$ is an integer. If $\lambda_1 - \lambda_2$ is an integer, then an $\mathcal{O}(N)$ matrix-vector product is immediate from the recurrence relation [31, (18.9.7)]. For example, if $\lambda_2 - \lambda_1$ is a positive integer, then we can factor A into the product of sparse matrices as follows:

(13)

$$A = P_N S_{\lambda_2-1} \cdots S_{\lambda_1+1} S_{\lambda_1} P_N^T, \quad S_\lambda = \begin{pmatrix} 1 & & & & \\ & \frac{\lambda}{\lambda+1} & & & \\ & & -\frac{\lambda}{\lambda+2} & & \\ & & & -\frac{\lambda}{\lambda+3} & \\ & & & & -\frac{\lambda}{\lambda+4} & \\ & & & & & \ddots & \\ & & & & & & \ddots & \ddots \end{pmatrix}.$$

where $P_N = \begin{bmatrix} I_{N+1} & O \end{bmatrix}$ is the $(N+1) \times \infty$ projection matrix. The matrix-vector product $\underline{c}^{(\lambda_2)} = A \underline{c}^{(\lambda_1)}$ can be computed in $\mathcal{O}(N)$ operations by applying each truncated sparse factor in turn. Since S_λ is banded and upper-triangular for all $\lambda > 0$, the matrix-vector product $\underline{c}^{(\lambda_2)} = A \underline{c}^{(\lambda_1)}$ can also be computed in $\mathcal{O}(N)$ operations when $\lambda_2 < \lambda_1$ by using backward substitution. The factorization in (13) is one key decomposition for the ultraspherical spectral method [32].

Case 2: $|\lambda_1 - \lambda_2| < 1$. Now assume that $|\lambda_1 - \lambda_2| < 1$, then the conversion matrix, A , in $\underline{c}^{(\lambda_2)} = A \underline{c}^{(\lambda_1)}$ has the following explicit formula [23, (3.6)]:

$$(14) \quad A_{jk} = \begin{cases} \frac{\Gamma(\lambda_2)\Gamma(j+\lambda_2)}{\Gamma(\lambda_1)\Gamma(\lambda_1-\lambda_2)} \left(\frac{\Gamma(\frac{k-j}{2}+\lambda_1-\lambda_2)}{\Gamma(\frac{k-j}{2}+1)} \right) \left(\frac{\Gamma(\frac{k+j}{2}+\lambda_1)}{\Gamma(\frac{k+j}{2}+\lambda_2+1)} \right), & 0 \leq j \leq k, \ k-j \text{ even,} \\ 0, & \text{otherwise.} \end{cases}$$

The formula in (14) reveals that the matrix A can be written as a diagonally-scaled Toeplitz-dot-Hankel matrix. More precisely, let $D_1 = \text{diag}(\lambda_2, \lambda_2 + 1, \dots, \lambda_2 + N)$, let $D_2 = I_{N+1}$ be the $(N+1) \times (N+1)$ identity matrix, and

$$T_{jk} = \begin{cases} \frac{\Gamma(\lambda_1-\lambda_2)\Gamma(\frac{k-j}{2}+\lambda_1-\lambda_2)}{\Gamma(\frac{k-j}{2}+1)}, & 0 \leq j \leq k \leq N, \ k-j \text{ even,} \\ 0, & \text{otherwise,} \end{cases}$$

$$H_{jk} = \frac{\Gamma(\lambda_2)\Gamma(\frac{k+j}{2}+\lambda_1)}{\Gamma(\lambda_1)\Gamma(\frac{k+j}{2}+\lambda_2+1)}, \quad 0 \leq j, k \leq N.$$

Since the entries of the Hankel part can be expressed as

$$H_{jk} = \frac{2\Gamma(\lambda_2)}{\Gamma(\lambda_1)\Gamma(\lambda_2 - \lambda_1 + 1)} \int_0^1 x^{j+k} x^{2(\lambda_1-1/2)} (1-x^2)^{\lambda_2-\lambda_1} dx, \quad 0 \leq j, k \leq N,$$

we know from Lemma 1 that H is real, symmetric, and positive semidefinite (note that $\lambda_2 - \lambda_1 > -1$ and $2(\lambda_1 - 1/2) > -1$ so that the measure μ_H is locally finite). From Theorem 3 we know that H can be approximated, up to an error of ϵ , by a rank $K = \mathcal{O}(\log N \log(1/\epsilon))$ matrix. Therefore, the algorithm that is summarized in Figure 1 is applicable in this case and can be used to compute $\underline{c}^{(\lambda_2)} = A \underline{c}^{(\lambda_1)}$ in $\mathcal{O}(N \log^2 N)$ operations.

Case 3: $|\lambda_1 - \lambda_2| > 1$, but the difference is not an integer. If $|\lambda_1 - \lambda_2| > 1$, then we reduce the quantity $|\lambda_1 - \lambda_2|$ by converting to either increase or reduce λ_1 by one (see Case 1). This is repeated if necessary until $|\lambda_1 - \lambda_2| < 1$ and the criterion for Case 2 is satisfied.

Conversions such as ultraspherical-to-Chebyshev and Chebyshev-to-ultraspherical are associated to upper-triangular matrices that can also be written as diagonally-scaled Toeplitz-dot-Hankel matrices. Fast $\mathcal{O}(N \log^2 N)$ algorithms based on Figure 1 for these conversions are also possible using the formulas in [23, Lem. 3.1].

5.3. Jacobi-to-Jacobi conversion. The Jacobi polynomial of degree k and parameter (α, β) is denoted by $P_k^{(\alpha, \beta)}$, where $\alpha, \beta > -1$ [31, Tab. 18.3.1]. The family of orthogonal polynomials $\{P_0^{(\alpha, \beta)}, P_1^{(\alpha, \beta)}, \dots\}$ is orthogonal with respect to the weight function $(1-x)^\alpha(1+x)^\beta$ on $[-1, 1]$. Here, $P_k^{(\alpha, \beta)}$ have the simplest and

standard normalization of $P_k^{(\alpha,\beta)}(1) = \binom{\alpha+k}{k}$ [31, (18.6.1)]. Mathematically, alternative normalizations only introduces multiplication with two additional diagonal matrices; however, numerically one should be concerned with overflow and underflow issues. The Jacobi polynomials can be seen as a generalization of Chebyshev, Legendre, and ultraspherical polynomials; see [31, Sec. 18.7].

For $\alpha, \beta, \gamma, \delta > -1$, the degree N polynomial $p_N(x)$ in (1) can be expanded in Jacobi bases as follows:

$$(15) \quad p_N(x) = \sum_{k=0}^N c_k^{(\alpha,\beta)} P_k^{(\alpha,\beta)}(x) = \sum_{k=0}^N c_k^{(\gamma,\delta)} P_k^{(\gamma,\delta)}(x),$$

where there is an upper-triangular matrix, A , such that $\underline{c}^{(\gamma,\delta)} = A\underline{c}^{(\alpha,\beta)}$. By the reflection formula $P_k^{(\alpha,\beta)}(-x) = (-1)^k P_k^{(\beta,\alpha)}(x)$ [31, Tab. 18.6.1], it is sufficient to assume that $\beta = \delta$ in (15).⁴ A different fast algorithm using off-diagonal low rank structure of the conversion matrix, which has a fast quasilinear complexity online cost and an $\mathcal{O}(N^2)$ precomputation is given in [38]. Another fast algorithm for computing Jacobi expansions coefficients of analytic functions is described in [44].

As in Section 5.2 there are several cases to consider when computing $\underline{c}^{(\gamma,\delta)} = A\underline{c}^{(\alpha,\beta)}$, where $\beta = \delta$, in fewer than $\mathcal{O}(N^2)$ operations: (1) $\alpha - \gamma$ is an integer, (2) $|\alpha - \gamma| < 1$ and $\alpha + \beta > -1$, (3) $|\alpha - \gamma| < 1$ and $\alpha + \beta \leq -1$, and (4) $|\alpha - \gamma| > 1$, but the difference is a not an integer.

Case 1: $\alpha - \gamma$ is an integer. First, suppose that $\alpha - \gamma$ is an integer. A fast matrix-vector product for $\underline{c}^{(\gamma,\beta)} = A\underline{c}^{(\alpha,\beta)}$ is almost immediate via the recurrence relation [31, 18.9.5] and [31, Tab. 18.6.1]. That is, assuming that $\gamma > \alpha$ we can factor A as follows:

$$A = P_N S_{(\gamma-1,\beta)} \cdots S_{(\alpha,\beta)} P_N^T, \quad S_{(\alpha,\beta)} = \begin{pmatrix} 1 & -\frac{\alpha+1}{\alpha+\beta+3} & & & & \\ & \frac{\alpha+\beta+2}{\alpha+\beta+3} & -\frac{\alpha+2}{\alpha+\beta+5} & & & \\ & & \frac{\alpha+\beta+3}{\alpha+\beta+5} & -\frac{\alpha+3}{\alpha+\beta+7} & & \\ & & & \ddots & \ddots & \\ & & & & \ddots & \ddots \end{pmatrix},$$

where $P_N = \begin{bmatrix} I_{N+1} & 0 \end{bmatrix}$ is the $(N+1) \times \infty$ projection matrix. The matrix-vector product $\underline{c}^{(\gamma,\beta)} = A\underline{c}^{(\alpha,\beta)}$ can be computed in $\mathcal{O}(N)$ operations by applying each truncated sparse factor in turn. If $\gamma < \alpha$, then since each $S_{(\alpha,\beta)}$ is bidiagonal and upper-triangular $\underline{c}^{(\gamma,\beta)} = A\underline{c}^{(\alpha,\beta)}$ can still be computed in $\mathcal{O}(N)$ operations by using backward substitution.

Case 2: $|\alpha - \gamma| < 1$ and $\alpha + \beta > -1$. When $|\alpha - \gamma| < 1$ there is no known *sparse* factorization for the conversion matrix as in Case 1. However, the following explicit formula for its entries is known [24, 25]:

$$A_{jk} = \frac{(2j + \gamma + \beta + 1)}{\Gamma(\alpha - \gamma)} \frac{\Gamma(k + \beta + 1)}{\Gamma(k + \alpha + \beta + 1)} \frac{\Gamma(j + \gamma + \beta + 1)}{\Gamma(j + \beta + 1)} \\ \times \frac{\Gamma(k - j + \alpha - \gamma)}{\Gamma(k - j + 1)} \frac{\Gamma(k + j + \alpha + \beta + 1)}{\Gamma(k + j + \gamma + \beta + 2)},$$

⁴If $\beta \neq \delta$, then a Jacobi (α, β) expansion can first be converted to a Jacobi (α, δ) expansion and then a Jacobi (γ, δ) expansion.

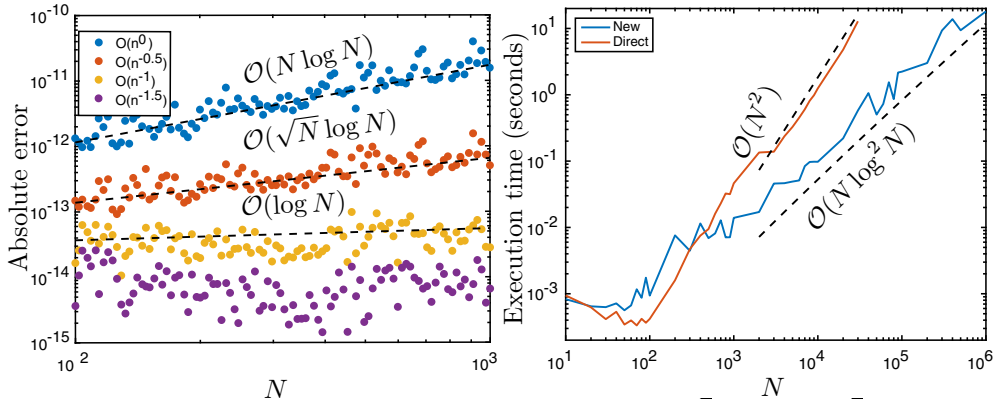


FIGURE 7. Left: Errors in computing $\underline{c}^{(-1/4, \sqrt{2}/2)} = A\underline{c}^{(0, \sqrt{2}/2)}$ with various decay rates in $\underline{c}^{(0, \sqrt{2}/2)}$. Arbitrary vectors are generated in precisely the same manner as in Section 4 and Section 5.1. The dashed lines show the observed error growth for the various decay rates. Right: Execution times between the direct (yellow), ASY (red), and the new algorithm described in this paper (blue).

where $0 \leq j \leq k \leq N$. The entries of A are zero otherwise. A careful inspection of this formula reveals that it can also be expressed as a diagonally-scaled Toeplitz-dot-Hankel matrix.

Moreover, the entries of the Hankel matrix can be expressed as follows:

$$(16) \quad H_{jk} = \frac{1}{\Gamma(\gamma - \alpha + 1)} \int_0^1 x^{j+k} x^{\alpha+\beta} (1-x)^{\gamma-\alpha} dx, \quad 0 \leq j, k \leq N,$$

proving that H is real, symmetric, and positive semidefinite (see Lemma 1) since $\alpha + \beta > -1$ and $\gamma - \alpha > -1$. Theorem 3 shows that H can be approximated, up to an error of ϵ , by a rank $\mathcal{O}(\log N \log(1/\epsilon))$ matrix. Therefore, the $\mathcal{O}(N \log^2 N)$ complexity algorithm summarized in Figure 1 can be employed for Jacobi conversion in this case.

The `jac2jac` algorithm in `Chebfun` and `FastTransforms.jl` implements this and the other three cases. In Figure 7 (left) we test the accuracy of our algorithm by using it to compute $\underline{c}^{(-1/4, \sqrt{2}/2)} = A\underline{c}^{(0, \sqrt{2}/2)}$ for various decay rates in the vector $\underline{c}^{(0, \sqrt{2}/2)}$. In Figure 7 (right) we compare the execution time of this algorithm and a direct approach. We observe that our algorithm is faster in this case when $N \geq 512$.

Case 3: $|\alpha - \gamma| < 1$ and $\alpha + \beta \leq -1$. This is a situation where the Jacobi conversion matrix in $\underline{c}^{(\gamma, \beta)} = A\underline{c}^{(\alpha, \beta)}$ can be written as a diagonally-scaled Toeplitz-dot-Hankel matrix; however, the Hankel part is not positive semidefinite (see (16)). This is similar to what happens in Section 5.1. Indeed, the submatrix $\tilde{H} = (H_{jk})_{1 \leq j, k \leq N}$ is in fact positive semidefinite because $\alpha + \beta + 1 > -1$. Hence, we can do the same trick when we compute $\underline{c}^{(\gamma, \beta)} = A\underline{c}^{(\alpha, \beta)}$: apply the first row directly, and use the algorithm described in Figure 1 on $(A_{jk})_{1 \leq j, k \leq N}$ for the remaining entries.

5.3.1. *Case 4:* $|\alpha - \gamma| > 1$, but the difference is not an integer. If $|\alpha - \gamma| > 1$, then either $\alpha < \gamma - 1$ or $\alpha > \gamma + 1$. If $\alpha < \gamma - 1$, then we convert the Jacobi (α, β) expansion to $(\alpha + 1, \beta)$ using Case 1, repeating if necessary until $|\alpha - \gamma| < 1$.

Similarly, if $\alpha > \gamma + 1$, then we convert the Jacobi (α, β) expansion to $(\alpha - 1, \beta)$ using Case 1. Again, repeating until $|\alpha - \gamma| < 1$. Thus, this case reduces the difference between α and γ until the criterion for either Case 2 or Case 3 is applicable.

The four cases above are implemented in the `jac2jac` commands in Chebfun [11] and FastTransforms.jl [41] with the syntax `jac2jac(v, a, b, g, d)`. Based on the particular values of \mathbf{a} , \mathbf{b} , \mathbf{g} , and \mathbf{d} various cases above are exercised. For all parameter ranges the cost of the conversion is at most $\mathcal{O}(N \log^2 N)$ operations. This algorithm is also employed for the commands `jac2cheb` and `cheb2jac` by exploiting the fact that $T_k(x) = P_k^{(-1/2, -1/2)}(x)/P_k^{(-1/2, -1/2)}(1)$ [31, (18.7.3)]. One can also compute the Jacobi-to-Chebyshev and Chebyshev-to-Jacobi conversions in $\mathcal{O}(N \log^2 N / \log \log N)$ operations using asymptotic expansions of Jacobi polynomials [39].

5.4. Laguerre-to-Laguerre conversion. We are not aware of major applications for Laguerre-to-Laguerre conversions, though related conversions are discussed in [9]. Due to the simplicity of the conversion, we include it in this section.

For $\alpha > -1$ the generalized Laguerre polynomial of degree k is given by $L_k^{(\alpha)}(x) = \frac{\Gamma(n+\alpha+1)}{\Gamma(\alpha+1)\Gamma(n+1)} {}_1F_1(-n, \alpha+1; x)$ [31, Tab. 18.5.12]. The sequence $L_0^{(\alpha)}, L_1^{(\alpha)}, \dots$ forms a family of polynomials that are orthogonal with respect to the weight function $x^\alpha e^{-x}$ on $[0, \infty)$.

Suppose that $\alpha_1 > -1$ and $\alpha_2 > -1$. Then, there is an upper-triangular matrix, A , that converts expansion coefficients in the $L^{(\alpha_1)}$ basis to coefficients in the $L^{(\alpha_2)}$ basis. If $\alpha_1 - \alpha_2$ is an integer, then an $\mathcal{O}(N)$ complexity algorithm for computing the matrix-vector product is almost immediate thanks to the recurrence relation given in [31, (18.9.13)]. If $\alpha_1 - \alpha_2$ is not an integer, then there is an explicit formula for the entries of A given by [31, (18.18.18)]

$$A_{jk} = \begin{cases} \frac{1}{\Gamma(\alpha_1 - \alpha_2)} \frac{\Gamma(k-j+\alpha_1-\alpha_2)}{\Gamma(k-j+1)}, & 0 \leq j \leq k \leq N, \\ 0, & \text{otherwise.} \end{cases}$$

One observes that this conversion matrix is a diagonally-scaled Toeplitz matrix, which is also a diagonally-scaled Toeplitz-dot-Hankel matrix by taking the Hankel part as the matrix of all ones. A fast $\mathcal{O}(N \log N)$ algorithm follows by a fast Toeplitz matrix-vector product based on the FFT [29].

5.5. Conversion between other polynomial bases. In related work it was found that the conversion operator between the Chebyshev polynomials of the second kind $U_k(x)$, and normalized orthogonal polynomials whose Jacobi operator is a trace-class perturbation of the Jacobi operator for $U_k(x)$ (of which all Jacobi polynomials are examples), is Toeplitz-plus-compact in an appropriate norm [45]. For this broader class of orthogonal polynomials, whether after dividing out the Toeplitz part of the conversion matrix we are left with a matrix of numerically low rank (such as a diagonally-scaled positive semidefinite Hankel matrix) is an open question.

One may also consider conversion between two polynomial bases where the supports of the two measures of orthogonality are not equal, e.g., Hermite-to-Chebyshev conversion. Such conversion matrices can be exponentially ill-conditioned [5, p. 93] and the conversion may be ill-advised for computational

purposes. That being said, there may be structure, such as Toeplitz-dot-Hankel or otherwise, in these conversion matrices. At present we have not explored this.

Expressions for connection coefficients are also known between many other families of orthogonal polynomials including the Wilson and Racah polynomials [36]. We hope that the associated conversion matrices are also structured.

CONCLUSION

Many of the standard conversion matrices for converting between expansions coefficients in orthogonal polynomial bases can be written as a diagonally-scaled Hadamard product between a Toeplitz and Hankel matrix. This leads to an $\mathcal{O}(N \log^2 N)$ complexity for basis conversion for a polynomial of degree N . The resulting algorithm is conceptually simple, FFT-based, and requires no precomputation, while being competitive in terms of computational time with existing fast algorithms. We hope this paper motivates fast matrix-vector products involving matrices that has the decomposition $A \circ B$, where A has a known fast matrix-vector product and B can be approximated by a low rank matrix.

ACKNOWLEDGMENTS

We thank the School of Mathematics and Statistics at The University of Sydney for awarding Alex Townsend a travel grant that allowed him to visit the school in February 2016. We thank the Cecil King Foundation and the London Mathematical Society for awarding Marcus Webb the Cecil King Travel Scholarship to visit The University of Sydney from January to April 2016. We thank Laurent Demanet and Haihao Lu for serendipitously discussing Toeplitz-dot-Hankel matrices with the first author a year earlier. We are grateful to Bernhard Beckermann for a discussion on the singular values of real, symmetric, and positive semidefinite Hankel matrices. We also thank Nick Hale and Mikael Slevinsky for discussions on related topics over many years and providing excellent comments that improved this paper and our implementations of the algorithms. We thank the referees and editor for their time and consideration.

REFERENCES

- [1] B. K. Alpert and V. Rokhlin, *A fast algorithm for the evaluation of Legendre expansions*, SIAM J. Sci. Statist. Comput. **12** (1991), no. 1, 158–179. MR1078802
- [2] C. Anderson and M. D. Dahleh, *Rapid computation of the discrete Fourier transform*, SIAM J. Sci. Comput. **17** (1996), no. 4, 913–919. MR1395355
- [3] R. Askey, *Orthogonal polynomials and special functions*, Society for Industrial and Applied Mathematics, Philadelphia, Pa., 1975. MR0481145
- [4] B. Beckermann, *The condition number of real Vandermonde, Krylov and positive definite Hankel matrices*, Numer. Math. **85** (2000), no. 4, 553–577. MR1771780
- [5] B. Beckermann and E. Bourreau, *How to choose modified moments?*, J. Comput. Appl. Math. **98** (1998), no. 1, 81–98. MR1656990
- [6] B. Beckermann and A. Townsend, *On the Singular Values of Matrices with Displacement Structure*, SIAM J. Matrix Anal. Appl. **38** (2017), no. 4, 1227–1248. MR3717820
- [7] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, *Julia: a fresh approach to numerical computing*, SIAM Rev. **59** (2017), no. 1, 65–98. MR3605826
- [8] M. J. Cantero and A. Iserles, *On rapid computation of expansions in ultraspherical polynomials*, SIAM J. Numer. Anal. **50** (2012), no. 1, 307–327. MR2888315
- [9] G. S. CHIRIKJIAN AND A. B. KYATKIN, *Harmonic Analysis for Engineers and Applied Scientists*, CRC Press, Second edition, 2000

- [10] W. S. Don and D. Gottlieb, *The Chebyshev-Legendre method: implementing Legendre methods on Chebyshev points*, SIAM J. Numer. Anal. **31** (1994), no. 6, 1519–1534. MR1302673
- [11] T. A. DRISCOLL, N. HALE, AND L. N. TREFETHEN, *Chebfun Guide*, Pafnuty Publications, Oxford, 2014.
- [12] M. FRIGO AND S. G. JOHNSON, *The design and implementation of FFTW3*, Proc. IEEE, **93** (2005), pp. 216–231.
- [13] G. H. Golub and C. F. Van Loan, *Matrix computations*, 4th ed., Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 2013. MR3024913
- [14] W. M. Gentleman, *Implementing Clenshaw-Curtis quadrature. II. Computing the cosine transformation*, Comm. ACM **15** (1972), 343–346. MR0327002
- [15] M. Gu and L. Miranian, *Strong rank revealing Cholesky factorization*, Electron. Trans. Numer. Anal. **17** (2004), 76–92. MR2040798
- [16] N. Hale and A. Townsend, *A fast, simple, and stable Chebyshev-Legendre transform using an asymptotic formula*, SIAM J. Sci. Comput. **36** (2014), no. 1, A148–A167. MR3163249
- [17] N. Hale and A. Townsend, *An algorithm for the convolution of Legendre series*, SIAM J. Sci. Comput. **36** (2014), no. 3, A1207–A1220. MR3217220
- [18] N. Hale and A. Townsend, *A fast FFT-based discrete Legendre transform*, IMA J. Numer. Anal. **36** (2016), no. 4, 1670–1684. MR3556400
- [19] H. Harbrecht, M. Peters, and R. Schneider, *On the low-rank approximation by the pivoted Cholesky decomposition*, Appl. Numer. Math. **62** (2012), no. 4, 428–440. MR2899254
- [20] N. J. Higham, *Analysis of the Cholesky decomposition of a semi-definite matrix*, Reliable numerical computation, Oxford Sci. Publ., Oxford Univ. Press, New York, 1990, pp. 161–185. MR1098323
- [21] S. Kunis and I. Melzer, *Fast evaluation of real and complex exponential sums*, Electron. Trans. Numer. Anal. **46** (2017), 23–35. MR3614469
- [22] A. Iserles, *A fast and simple algorithm for the computation of Legendre coefficients*, Numer. Math. **117** (2011), no. 3, 529–553. MR2772418
- [23] J. Keiner, *Computing with expansions in Gegenbauer polynomials*, SIAM J. Sci. Comput. **31** (2009), no. 3, 2151–2171. MR2516147
- [24] J. KEINER, *Fast Polynomial Transforms*, Logos Verlag Berlin GmbH, 2011.
- [25] P. Maroni and Z. da Rocha, *Connection coefficients between orthogonal polynomials and the canonical sequence: an approach based on symbolic computation*, Numer. Algorithms **47** (2008), no. 3, 291–314. MR2385739
- [26] J. C. Mason and D. C. Handscomb, *Chebyshev polynomials*, Chapman & Hall/CRC, Boca Raton, FL, 2003. MR1937591
- [27] G. Marsaglia and G. P. H. Styan, *Equalities and inequalities for ranks of matrices*, Linear and Multilinear Algebra **2** (1974/75), 269–292. MR0384840
- [28] A. Mori, R. Suda, and M. Sugihara, *An improvement on Orszag’s fast algorithm for Legendre polynomial transform* (Japanese, with English and Japanese summaries), Trans. Inform. Process. Soc. Japan **40** (1999), no. 9, 3612–3615. MR1724286
- [29] M. K. Ng, *Iterative methods for Toeplitz systems*, Numerical Mathematics and Scientific Computation, Oxford University Press, New York, 2004. MR2108963
- [30] S. OLVER, R. M. SLEVINSKY, et al., <https://github.com/ApproxFun/ApproxFun.jl>, v0.1.0, 2016.
- [31] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark (eds.), *NIST handbook of mathematical functions*, U.S. Department of Commerce, National Institute of Standards and Technology, Washington, DC; Cambridge University Press, Cambridge, 2010. With 1 CD-ROM (Windows, Macintosh and UNIX). MR2723248
- [32] S. Olver and A. Townsend, *A fast and well-conditioned spectral method*, SIAM Rev. **55** (2013), no. 3, 462–489. MR3089410
- [33] S. A. ORSZAG, *Fast eigenfunction transforms*, Science and Computers, Academic Press, New York, (1986), pp. 23–30.
- [34] V. V. Peller, *Hankel operators and their applications*, Springer Monographs in Mathematics, Springer-Verlag, New York, 2003. MR1949210
- [35] D. Potts, G. Steidl, and M. Tasche, *Fast algorithms for discrete polynomial transforms*, Math. Comp. **67** (1998), no. 224, 1577–1590. MR1474655
- [36] J. Sánchez-Ruiz and J. S. Dehesa, *Some connection and linearization problems for polynomials in and beyond the Askey scheme*, Proceedings of the Fifth International Symposium on

- Orthogonal Polynomials, Special Functions and their Applications (Patras, 1999), J. Comput. Appl. Math. **133** (2001), no. 1-2, 579–591. MR1858314
- [37] J. Shen, *Efficient spectral-Galerkin method. I. Direct solvers of second- and fourth-order equations using Legendre polynomials*, SIAM J. Sci. Comput. **15** (1994), no. 6, 1489–1505. MR1298626
- [38] J. SHEN, Y. WANG, AND J. XIA, *Fast structured Jacobi-Jacobi transforms*, preprint, 2016.
- [39] R. M. SLEVINSKY, *On the use of Hahn's asymptotic formula and stabilized recurrence for a fast, simple, and stable Chebyshev–Jacobi transform*, to appear in IMA Numer. Anal., 2017.
- [40] G. Szegő, *Orthogonal polynomials*, American Mathematical Society Colloquium Publications, Vol. 23. Revised ed, American Mathematical Society, Providence, R.I., 1959. MR0106295
- [41] R. M. SLEVINSKY, S. OLVER, et al. <https://github.com/MikaelSlevinsky/FastTransforms.jl>, v0.0.6, 2016.
- [42] A. Townsend and L. N. Trefethen, *Continuous analogues of matrix factorizations*, Proc. A. **471** (2015), no. 2173, 20140585, 21. MR3285359
- [43] L. N. Trefethen, *Approximation theory and approximation practice*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2013. MR3012510
- [44] H. WANG AND D. HUYBRECHS, *Fast and highly accurate computation of Chebyshev expansion coefficients of analytic functions*, arXiv preprint. arXiv: 1404.2463 (2014).
- [45] M. WEBB AND S. OLVER, *Spectra of Jacobi operators via connection coefficient matrices*, arXiv preprint. arXiv: 1702.03095 (2017).
- [46] W. Xu and S. Qiao, *A fast symmetric SVD algorithm for square Hankel matrices*, Linear Algebra Appl. **428** (2008), no. 2-3, 550–563. MR2374566

DEPARTMENT OF MATHEMATICS, CORNELL UNIVERSITY, ITHACA, NEW YORK 14853
E-mail address: townsend@cornell.edu

DEPARTMENT OF COMPUTER SCIENCE, KU LEUVEN, 3001 LEUVEN, BELGIUM
E-mail address: marcus.webb@cs.kuleuven.be

DEPARTMENT OF MATHEMATICS, IMPERIAL COLLEGE, LONDON SW7 2AZ, UNITED KINGDOM
E-mail address: s.olver@imperial.ac.uk