An algorithm for the numerical evaluation of the associated Legendre functions that runs in time independent of degree and order

James Bremer[a]

[a]*Department of Mathematics, University of California, Davis*

Abstract

We describe a method for the numerical evaluation of normalized versions of the associated Legendre functions $P_\nu^{-\mu}$ and $Q_\nu^{-\mu}$ of degrees $0 \leq \nu \leq 1,000,000$ and orders $-\nu \leq \mu \leq \nu$ for arguments in the interval $(-1, 1)$. Our algorithm, which runs in time independent of ν and μ, is based on the fact that while the associated Legendre functions themselves are extremely expensive to represent via polynomial expansions, the logarithms of certain solutions of the differential equation defining them are not. We exploit this by numerically precomputing the logarithms of carefully chosen solutions of the associated Legendre differential equation and representing them via piecewise trivariate Chebyshev expansions. These precomputed expansions, which allow for the rapid evaluation of the associated Legendre functions over a large swath of parameter domain mentioned above, are supplemented with asymptotic and series expansions in order to cover it entirely. The results of numerical experiments demonstrating the efficacy of our approach are presented, and our code for evaluating the associated Legendre functions is publicly available.

Keywords: associated Legendre functions, fast algorithms, butterfly algorithms, special functions, nonoscillatory phase functions, asymptotic methods

1. Introduction

The associated Legendre functions arise in many contexts in applied mathematics and physics. They are perhaps most commonly encountered in connection with spherical harmonics, which are tensor products of associated Legendre functions and exponential functions. Among other things, the spherical harmonics are used to efficiently represent smooth functions given on the surface of the sphere, and in spectral methods for the solution of partial differential equations.

In this paper, we describe an algorithm for the numerical evaluation of versions of the associated Legendre functions of the first and seconds kinds. Following standard convention, we will denote by P_ν^μ the associated Legendre function of the first kind of degree ν and order μ, and by Q_ν^μ the associated Legendre function of the second kind of degree ν and order μ (see, for instance, Section 5.15 of [23] or Section 3.4 of [7] for definitions). Since the magnitudes of the associated Legendre functions are excessively large, even for parameters of relatively small magnitude, the normalized

Email address: bremer@math.ucdavis.edu (James Bremer)

associated Legendre functions defined via

$$\bar{P}_\nu^\mu(x) = \sqrt{\left(\nu + \frac{1}{2}\right) \frac{\Gamma(\nu - \mu + 1)}{\Gamma(\nu + \mu + 1)}}\; P_\nu^\mu(x) \tag{1}$$

and

$$\bar{Q}_\nu^\mu(x) = \frac{2}{\pi}\sqrt{\left(\nu + \frac{1}{2}\right) \frac{\Gamma(\nu - \mu + 1)}{\Gamma(\nu + \mu + 1)}}\; Q_\nu^\mu(x) \tag{2}$$

are often used in lieu of P_ν^μ and Q_ν^μ. Our interest in associated Legendre functions stems principally from their connection with spherical harmonics, and for many calculations involving the spherical harmonics it is more convenient to work with the functions $\tilde{P}_\nu^\mu$ and $\tilde{Q}_\nu^\mu$ defined by the formulas

$$\tilde{P}_\nu^\mu(t) = \bar{P}_\nu^\mu\left(\cos(t)\right)\sqrt{\sin(t)} \tag{3}$$

and

$$\tilde{Q}_\nu^\mu(t) = \bar{Q}_\nu^\mu\left(\cos(t)\right)\sqrt{\sin(t)} \tag{4}$$

(the variable t corresponds to one of the coordinates in the standard parameterization of the unit sphere S^2). Our algorithm, which runs in time independent of the degree ν and order μ, allows for the evaluation of $\tilde{P}_\nu^{-\mu}$ and $\tilde{Q}_\nu^{-\mu}$ when $0 \le \nu \le 1,000,000$, $0 \le \mu \le \nu$ and $0 < t \le \frac{\pi}{2}$ (in particular, both ν and μ can take on noninteger values). It is a consequence of standard connection formulas (such as those appearing in Section 3.4 of [7]) that this suffices for the evaluation of $\bar{P}_\nu^\mu(x)$ and $\bar{Q}_\nu^\mu(x)$ for any $0 \le \nu \le 1,000,000$, $-\nu \le \mu \le \nu$ and $-1 < x < 1$.

1.1. Overview of our algorithm

The functions $\tilde{P}_\nu^{-\mu}$ and $\tilde{Q}_\nu^{-\mu}$ satisfy the second order linear ordinary differential equation

$$y''(t) + \left(\lambda^2 - \eta^2 \csc^2(t)\right) y(t) = 0 \quad \text{for all} \;\; 0 < t < \frac{\pi}{2} \tag{5}$$

with $\lambda = \nu + \frac{1}{2}$ and $\eta^2 = \mu^2 - \frac{1}{4}$. By a slight abuse of terminology, we will refer to (5) as the associated Legendre differential equation. When $0 \le \mu \le \frac{1}{2}$, the coefficient of y in (5) is positive on the interval $\left(0, \frac{\pi}{2}\right)$, whereas when $\mu > \frac{1}{2}$ it is negative on the interval

$$\left(0, \arcsin\left(\frac{\eta}{\lambda}\right)\right) \tag{6}$$

and positive on

$$\left(\arcsin\left(\frac{\eta}{\lambda}\right), \frac{\pi}{2}\right). \tag{7}$$

It follows from these observations and well-known WKB estimates (see, for example, [9]) that when $\mu > \frac{1}{2}$, the solutions of (5) behave roughly like combinations of increasing or decreasing exponentials on (6) and are oscillatory on (7), whereas when $0 \le \mu \le \frac{1}{2}$, they are oscillatory on all of $\left(0, \frac{\pi}{2}\right)$. We will refer to the subset

$$\begin{aligned} \mathcal{O} = &\left\{(\nu, \mu, t) : \nu \ge 0,\; 0 \le \mu \le \min\left\{\nu, \frac{1}{2}\right\} \text{ and } 0 < t \le \frac{\pi}{2}\right\} \bigcup \\ &\left\{(\nu, \mu, t) : \nu > \frac{1}{2},\quad \frac{1}{2} < \mu \le \nu \;\text{ and }\; \arcsin\left(\frac{\eta}{\lambda}\right) \le t \le \frac{\pi}{2}\right\} \end{aligned} \tag{8}$$

of $\mathbb{R}^3$ as the oscillatory region, and to the subset

$$\mathcal{N} = \left\{ (\nu, \mu, t) : \nu \geq 0,\ \mu > \frac{1}{2} \ \text{ and } \ 0 < t < \arcsin\left(\frac{\eta}{\lambda}\right) \right\} \tag{9}$$

as the nonoscillatory region. When $\nu \gg \mu > \frac{1}{2}$, the solutions of (5) are highly oscillatory on (7), and when $\nu \geq \mu \gg \frac{1}{2}$, they behave roughly like combinations of rapidly decreasing and increasing exponentials on (6). Consequently, they cannot be effectively represented via polynomial expansions in the variables ν, μ and t on either of the sets $\mathcal{O}$ or $\mathcal{N}$, at least for large values of the parameters.

Nonetheless, the logarithms of certain solutions of (5) can be represented efficiently via polynomial expansions on the sets $\mathcal{N}$ and $\mathcal{O}$. This observation is related to the well-known fact that the associated Legendre differential equation admits a nonoscillatory phase function. Many special functions of interest posses this property as well, at least in an asymptotic sense [21, 6]. However, the sheer effectiveness with which nonoscillatory phase functions can represent solutions of the general equation

$$y''(t) + \lambda^2 q(t) y(t) = 0 \ \text{ for all } \ a < t < b \tag{10}$$

in which the coefficient q is smooth and positive appears to have been overlooked. Indeed, under mild conditions on q, it is shown in [5] that there exist a positive real number σ, a nonoscillatory function α and a basis of solutions $\{u, v\}$ of (10) such that

$$u(t) = \frac{\cos\left(\alpha(t)\right)}{\sqrt{|\alpha'(t)|}} + O\left(\exp(-\sigma\lambda)\right) \tag{11}$$

and

$$v(t) = \frac{\sin\left(\alpha(t)\right)}{\sqrt{|\alpha'(t)|}} + O\left(\exp(-\sigma\lambda)\right). \tag{12}$$

The constant σ is a measure of the extent to which q oscillates, with larger values of σ corresponding to greater smoothness on the part of q. The function α is nonoscillatory in the sense that it can be represented using various series expansions using a number of terms in which is independent of λ. That is, $O(\exp(-\sigma\lambda))$ accuracy is obtained using an $O(1)$-term expansion. The results of [5] are akin to standard results on WKB approximation in that they apply to the more general case in which q varies with the parameter λ assuming only that q satisfies certain innocuous hypotheses independent of λ. An effective numerical algorithm for the computation of nonoscillatory phase functions for fairly general second order differential equations is described in [3], although we will not need it here. We will instead use specialized formulas which apply only in the case of associated Legendre functions. However, the existence of the algorithm of [3] and results of [5] mean that the approach of this paper can be applied to a large class of special functions satisfying second order differential equations.

The algorithm of this paper operates by numerically calculating the logarithms of certain solutions of the associated Legendre differential equation. We represent them via trivariate Chebyshev expansions — or rather, truncated versions of these expansions which we call "compressed" trivariate Chebyshev expansions — the coefficients of which are stored in a table on the disk. This table is computed only once and is loaded into memory and used to evaluate the associated Legendre functions rapidly. The table used in the experiments described in this paper is approximately 138 MB in size. We supplement these precomputed expansions with series and asymptotic expansions in order to cover the entire parameter domain mentioned above. In addition to the values of the

functions $\tilde{P}_\nu^{-\mu}$ and $\tilde{Q}_\nu^{-\mu}$, our algorithm also produces the values of a nonoscillatory phase function for (5) and its derivative when (ν, μ, t) is in the oscillatory region $\mathcal{O}$ and the values of the logarithms of $\tilde{P}_\nu^{-\mu}$ and $\tilde{Q}_\nu^{-\mu}$ when (ν, μ, t) is in the nonoscillatory region $\mathcal{N}$. The phase function is useful for, among other things, calculating the roots of the associated Legendre functions and applying special function transforms involving the associated Legendre functions. Calculating the values of the logarithms of (3) and (4) obviates many problems which arise from numerical overflow and underflow.

1.2. Prior work

There are extensive literatures both on the representation of associated Legendre functions via asymptotic expansions and on their numerical evaluation via the three-term recurrence relations they satisfy. Nonetheless, both such approaches to the evaluation of associated Legendre functions suffer from significant drawbacks which warrant the development of new approaches to the problem.

The most obvious disadvantage of the three-term recurrence relations is that a sequence of associated Legendre functions of various orders and/or degrees must be evaluated in order to obtain the value of the function of a specified degree and order. In [26], it was observed that butterfly algorithms, which are schemes for the rapid application of certain classes of matrices (e.g., those satisfying the complementary low-rank property), can be used for the rapid application of spherical harmonic transforms. The algorithm of [26] uses the three-term recurrence relations to evaluate the associated Legendre functions. However, the classical butterfly algorithm requires that every element of a matrix being applied be evaluated, and the computations of [26] are arranged so that the use of the three-term recurrence relations does not result in any inefficiency. In cases in which an $\mathcal{O}(1)$ procedure is available for evaluating the entries of the matrix, more recent versions of the butterfly algorithm, such as [18] and [17], can greatly accelerate computations through interpolation. The algorithm of this paper is intended, among other things, to be combined with the schemes of [18] and [17] in order to rapidly apply spherical harmonic transforms of large orders. The availability of an $\mathcal{O}(1)$ procedure is also essential for parallelizing the workloads of butterfly algorithms, which is necessary in order to apply spherical harmonic transforms of large orders. There is considerable interest in doing so; to give one example, transforms of order larger than 10^5 are often used in calculations simulating cosmic microwave background [8, 24].

There are further disadvantages of using the three-term recurrence relations to evaluate the associated Legendre functions numerically. Although [22] describes a recurrence-based method for evaluating associated Legendre functions of noninteger degrees and integer orders, the author is unaware of any similar approach for the evaluation of associated Legendre functions of noninteger degree and noninteger orders and most schemes based on the recurrence relations are restricted to integer orders and degrees. More seriously, the recurrence relations are numerically unstable, with the consequence that most numerically viable implementations of them use some form of extended precision arithmetic (examples include [11] and [22]). Usually, such algorithms proceeds by representing real numbers in the form $x \cdot R^y$ with R a chosen radix and x and y double precision numbers. We note, though, that the algorithm of [26] avoids the use of extended precision arithmetic by only evaluating the associated Legendre function at the nodes of certain Gauss-Jacobi quadrature rules at which the recurrence relations appear to be stable.

It is often suggested that existing asymptotic expansions suffice for the numerical evaluation of the associated Legendre functions. While many highly effective approximations are available, it

appears to be quite difficult to produce a numerical algorithm which is efficient and accurate over the entire range of the variables ν, μ and t considered here. The trigonometric expansions

$$\tilde{P}_{\nu}^{-\mu}(t) = \sqrt{\frac{2\left(\nu+\frac{1}{2}\right)\Gamma(\nu+\mu+1)\Gamma(\nu-\mu+1)}{\pi\sin(t)}}\frac{1}{\Gamma\left(\nu+\frac{3}{2}\right)} \times \sum_{k=0}^{\infty}\left((-1)^k\frac{\left(\mu+\frac{1}{2}\right)_k\left(-\mu+\frac{1}{2}\right)_k}{\left(2\sin(t)\right)^k\Gamma(k+1)\left(\nu+\frac{3}{2}\right)_k}\sin\left(\left(\nu+k+\frac{1}{2}\right)t+\frac{\pi}{2}\left(k-\mu\right)+\frac{\pi}{4}\right)\right) \tag{13}$$

and

$$\tilde{Q}_{\nu}^{-\mu}(t) = \sqrt{\frac{2\left(\nu+\frac{1}{2}\right)\Gamma(\nu+\mu+1)\Gamma(\nu-\mu+1)}{\pi\sin(t)}}\frac{1}{\Gamma\left(\nu+\frac{3}{2}\right)} \times \sum_{k=0}^{\infty}\left((-1)^k\frac{\left(\mu+\frac{1}{2}\right)_k\left(-\mu+\frac{1}{2}\right)_k}{\left(2\sin(t)\right)^k\Gamma(k+1)\left(\nu+\frac{3}{2}\right)_k}\cos\left(\left(\nu+k+\frac{1}{2}\right)t+\frac{\pi}{2}\left(k-\mu\right)+\frac{\pi}{4}\right)\right), \tag{14}$$

which can be found in Section 3.5 of [7], illustrate some of difficulties that arise. In (13) and (14), $(x)_k$ is the Pochhammer symbol defined via

$$(x)_k = \frac{\Gamma(x+k)}{\Gamma(x)}. \tag{15}$$

These expansions converge when $\frac{\pi}{6} < t < \frac{5\pi}{6}$, andthey can be used as asymptotic approximations when t is outside of that interval. However, they require on the order of $|\mu|$ terms in order to achieve a fixed accuracy. Even more seriously, when $|\mu|$ is not small relative to ν, the terms in (13) and (14) are of large magnitude and alternate signs, with the consequence that the numerical evaluation of (13) and (14) generally results in catastrophic cancellation errors.

Liouville-Green methods can be used to obtain asymptotic expansions of the associated Legendre functions which are uniform in the argument t and apply in the event that $0 \le \mu \le \nu$. In [2], the new dependent variable ζ defined via the implicit relations

$$\int_{\gamma^2}^{\zeta}\frac{\sqrt{\xi^2-\gamma^2}}{2\xi}\,d\xi = -\int_{\sqrt{1-\gamma^2}}^{x}\frac{\sqrt{1-\gamma^2-\tau^2}}{1-\tau^2}\,d\tau \quad \text{if} \quad x < \sqrt{1-\gamma^2} \tag{16}$$

and

$$\int_{\gamma^2}^{\zeta}\frac{\sqrt{\xi^2-\gamma^2}}{2\xi}\,d\xi = -\int_{\sqrt{1-\gamma^2}}^{x}\frac{\sqrt{\tau^2-1+\gamma^2}}{1-\tau^2}\,d\tau \quad \text{if} \quad x > \sqrt{1-\gamma^2}, \tag{17}$$

where $\lambda = \nu + \frac{1}{2}$ and $\gamma = \frac{\mu}{\lambda}$, is introduced to obtain the uniform asymptotic expansions

$$\bar{P}_{\nu}^{-\mu}(x) \sim \left(\frac{\zeta-\gamma^2}{1-\gamma^2-x^2}\right)^{\frac{1}{4}}\left(\sqrt{\lambda}J_{\mu}\left(\lambda\sqrt{\zeta}\right)\sum_{k=0}^{\infty}\frac{F_k(\zeta)}{\lambda^{2s}} + \left(\frac{\zeta}{\lambda}\right)^{\frac{1}{2}}J'_{\mu}\left(\lambda\sqrt{\zeta}\right)\sum_{k=0}^{\infty}\frac{\tilde{F}_k(\zeta)}{\lambda^{2s}}\right) \tag{18}$$

as $\nu \to \infty$ and

$$\bar{Q}_{\nu}^{-\mu}(x) \sim -\left(\frac{\zeta-\gamma^2}{1-\gamma^2-x^2}\right)^{\frac{1}{4}}\left(\sqrt{\lambda}Y_{\mu}\left(\lambda\sqrt{\zeta}\right)\sum_{k=0}^{\infty}\frac{F_k(\zeta)}{\lambda^{2s}} + \left(\frac{\zeta}{\lambda}\right)^{\frac{1}{2}}Y'_{\mu}\left(\lambda\sqrt{\zeta}\right)\sum_{k=0}^{\infty}\frac{\tilde{F}_k(\zeta)}{\lambda^{2s}}\right) \tag{19}$$

as $\nu \to \infty$. The coefficients are given by $F_0(\zeta) = 1$, as well as by the formulas

$$\tilde{F}_k(\zeta) = \left|\zeta - \alpha^2\right|^{-\frac{1}{2}} \int_{\alpha^2}^{\zeta} \left|\xi - \alpha^2\right|^{-\frac{1}{2}} \left(\xi F_k''(\xi) + F_k'(\xi) - \psi(\xi) F_k(\xi)\right)\ d\xi \tag{20}$$

and

$$F_{k+1}(\zeta) = -\zeta \tilde{F}_k'(\zeta) + \int_{\alpha^2}^{\zeta} \psi(\xi) \tilde{F}_k'(\xi)\ d\xi + \zeta \tilde{F}_k'(\alpha^2), \tag{21}$$

where

$$\psi(\zeta) = \frac{1}{16} \left(\zeta - \alpha^2\right)^{-2} \left(\zeta + 4\alpha^2 + \frac{1 - x^2}{\zeta} \left(\frac{\zeta - \alpha^2}{1 - \alpha^2 - x^2} \right)^3 \left(\left(1 - 4\alpha^2\right) x^2 + \left(\alpha^4 - 1\right) \right) \right). \tag{22}$$

Note that the variable x in (22) depends on ζ through either (16) or (17); if this is neglected and x is treated as a constant in (20) and (21), then these integrals diverge. While (18) and (19) are powerful expansions in the sense that they converge rapidly for a large range of parameters and arguments, it is not clear how to accurately and efficiently compute the variable ζ defined via the relations (16) and (17) given x. Nor is it obvious how to evaluate the coefficients in these expansions. Only the first few are known analytically, and the numerical calculation of the remaining coefficients is complicated by the delicate cancellations of singularities which occur in these formulas. Alternate Liouville-Green expansions for (3) and (4) are given in Chapter 12 of [23]. However, the coefficients in these expansions appear to be no easier to compute than those in (18) and (19), and they have the unfortunate property that are only applicable when μ is small relative to ν.

Asymptotic expansions are a promising approach to the numerical evaluation of associated Legendre functions, but there are substantial difficulties which must be addressed in order to construct a viable numerical method using them.

1.3. Outline of this paper

The remainder of this paper is structured as follows. In Section 2, we review certain mathematical facts and numerical procedures which are used in the rest of this article. In Section 3, we describe a numerical method for the solution of the differential equation (5) which runs in time independent of the parameters ν and μ. The construction of a precomputed table of expansions of the associated Legendre functions which makes use of the algorithm of Section 3 is discussed in Section 4. In Section 5, we detail our algorithm for the numerical calculation of the associated Legendre functions. Section 6 describes the results of numerical experiments carried out to verify the efficacy of our algorithm. We close with a few brief comments in Section 7.

2. Mathematical and numerical preliminaries

2.1. The condition number of the evaluation of a function

The condition number of the evaluation of a differentiable function $f : \mathbb{R} \to \mathbb{R}$ at the point x is commonly defined to be

$$\kappa_f(x) = \left| \frac{x f'(x)}{f(x)} \right| \tag{23}$$

(see, for instance, Section 1.6 of [14]). This quantity measures the ratio of the magnitude of the relative change in $f(x)$ induced by a small change in the argument x to the magnitude of the

relative change in x in the sense that

$$\left|\frac{f(x+\delta)-f(x)}{f(x)}\right| \approx \kappa_f(x) \left|\frac{\delta}{x}\right| \tag{24}$$

for small δ. Since almost all quantities which arise in the course of numerical calculations are subject to perturbations with relative magnitudes on the order of machine epsilon, we consider

$$\kappa_f(x)\epsilon_0, \tag{25}$$

where ϵ_0 denotes machine epsilon, to be a rough estimate of the relative accuracy one should expect when evaluating $f(x)$ numerically (in fact, it tends to be a slightly pessimistic estimate). In the rest of this paper, we take ϵ_0 to be

$$\epsilon_0 = 2^{-52} \approx 2.22044604925031 \times 10^{-16}. \tag{26}$$

It is immediately clear from (23) that when $f'(x_0)x_0 \neq 0$ and $f(x_0) = 0$, $\kappa_f(x)$ diverges to ∞ as $x \to x_0$. One consequence of this is that there is often a significant loss of relative accuracy when a function is evaluated near one of its roots. For the most part, we avoid this issue by representing the solutions of the associated Legendre differential equation via functions which are bounded away from 0.

2.2. Trivariate Chebyshev expansions

For each nonnegative integer n, the Chebyshev polynomial of degree n is defined for $-1 < x < 1$ via the formula

$$T_n(x) = \cos\left(n \arccos(x)\right). \tag{27}$$

The trivariate Chebyshev series of a continuous function $f : [-1,1]^3 \to \mathbb{R}$ is

$$\sum_{i=0}^{\infty}{}' \sum_{j=0}^{\infty}{}' \sum_{k=0}^{\infty}{}' a_{i,j,k} T_i(x) T_j(y) T_k(z), \tag{28}$$

where the coefficients are defined via the formula

$$a_{i,j,k} = \frac{8}{\pi^3} \int_{-1}^{1}\int_{-1}^{1}\int_{-1}^{1} f(x,y,z) T_i(x) T_j(y) T_k(z) \frac{dx}{\sqrt{1-x^2}} \frac{dy}{\sqrt{1-y^2}} \frac{dz}{\sqrt{1-z^2}} \tag{29}$$

and the dashes next to the summation symbols indicate that the first term in each sum is halved. The well-known relationship between Chebyshev and Fourier series (see, for instance, [20]), together with the results of [10] on the pointwise almost everywhere convergence of multiple Fourier series immediately imply that

$$\lim_{N\to\infty} \sum_{i=0}^{N}{}' \sum_{j=0}^{N}{}' \sum_{k=0}^{N}{}' a_{i,j,k} T_i(x) T_j(y) T_k(z) = f(x,y,z) \tag{30}$$

for almost all $(x,y,z) \in [-1,1]^3$. As in the case of univariate Chebyshev series, under mild smoothness conditions on f, the convergence of (28) is uniform. See, for instance, Theorem 5.9 in [20].

If $f(x,y,z)$ is analytic on the set

$$\left\{(x,y,z) \in \mathbb{C}^3 : \left|x + \sqrt{x^2-1}\right| < r_1, \quad \left|y + \sqrt{y^2-1}\right| < r_2, \quad \left|z + \sqrt{z^2-1}\right| < r_3\right\}, \tag{31}$$

where $r_1, r_2, r_3 > 1$, then $|a_{i,j,k}| = O\left(r_1^{-i} r_2^{-j} r_3^{-k}\right)$ (this is according to Theorem 11 in Chapter V of [1]), with the consequence that the limit in (30) converges rapidly to f when f is analytic in a

large neighborhood containing $[-1,1]^3$.

For each nonnegative integer n, we refer to the collection of points

$$\rho_{j,n} = -\cos\left(\frac{\pi j}{n}\right), \quad j = 0, 1, \ldots, n, \tag{32}$$

as the $(n+1)$-point Chebyshev grid on the interval $[-1,1]$, and we call individual elements of this set Chebyshev nodes or points. One discrete version of the well-known orthogonality relation

$$\int_{-1}^{1} \frac{T_i(x)T_j(x)}{\sqrt{1-x^2}}\, dx = \begin{cases} 0 & \text{if } \ i \neq j \\ \frac{\pi}{2} & \text{if } \ i = j > 0 \\ \pi & \text{if } \ i = j = 0. \end{cases} \tag{33}$$

is

$$\sum_{l\,=0}^{n}{}'' T_i(\rho_{l,n})T_j(\rho_{l,n}) = \begin{cases} 0 & \text{if } \ 0 \leq i, j \leq n \ \text{ and } i \neq j \\ \frac{n}{2} & \text{if } \ 0 < i = j < n \\ n & \text{if } \ i = j = 0 \ \text{ or } \ i = j = n. \end{cases} \tag{34}$$

Here, the double dash next to the summation sign indicates that the first and last term in the series are halved. Formula (34) can be found in a slightly different form in Chapter 4 of [20].

It follows easily from (33) and (34) that any trivariate polynomial f of degree less than or equal to n can be represented in the form

$$f(x,y,z) = \sum_{i=0}^{n}{}'' \sum_{j=0}^{n}{}'' \sum_{k=0}^{n}{}'' b_{i,j,k}T_i(x)T_j(y)T_k(z), \tag{35}$$

where

$$b_{i,j,k} = \frac{8}{n^3} \sum_{r=0}^{n}{}'' \sum_{s=0}^{n}{}'' \sum_{t=0}^{n}{}'' T_i\left(\rho_{r,n}\right) T_j\left(\rho_{s,n}\right) T_k\left(\rho_{t,n}\right) f\left(\rho_{r,n}, \rho_{s,n}, \rho_{t,n}\right). \tag{36}$$

If f is not a polynomial or is a polynomial of degree greater than n, then the representation (35) is no longer exact. However, in this event, there is a well-known relationship between the coefficients defined via (36) and those given by (28). In particular,

$$b_{i,j,k} = a_{ijk} + \sum_{l_1=1}^{\infty}\sum_{l_2=1}^{\infty}\sum_{l_3=1}^{\infty} \left(a_{i+2l_1n,j+2l_2n,k+2l_3n} + a_{-i+2l_1n,-j+2l_2n,-k+2l_3n}\right) \tag{37}$$

for all $0 \leq i, j, k \leq n$ (the one-dimensional version of this result can be found, for instance, in [20]). Using (37) it is easy to show that there exists a constant C such that

$$\begin{aligned} &\sup_{x\,\in[-1,1]} \left| f(x,y,z) - \sum_{i=0}^{n}{}'' \sum_{j=0}^{n}{}'' \sum_{k=0}^{n}{}'' b_{i,j,k}T_i(x)T_j(y)T_k(z) \right| \\ &\leq C \sup_{x\in[-1,1]} \left| f(x,y,z) - \sum_{i=0}^{n}{}'' \sum_{j=0}^{n}{}'' \sum_{k=0}^{n}{}'' a_{i,j,k}T_i(x)T_j(y)T_k(z) \right|. \end{aligned} \tag{38}$$

It follows, of course, that the sum (36) converges rapidly to f when f is analytic in a large neighbor-

hood of $[-1,1]^3$. By a slight abuse of terminology, we will refer to (35) as the n^{th} order Chebyshev expansion for the function f.

2.3. Compressed trivariate Chebyshev expansions

It often happens that many of the coefficients in the trivariate Chebyshev expansion (35) of a function $f : [-1,1]^3 \to \mathbb{R}$ are of negligible magnitude. In order to reduce the cost of storing such expansions as well as the cost of evaluating them, we use the following construction to reduce the number of coefficients which need to be considered.

Suppose that $\epsilon > 0$, and that

$$\sum_{i=0}^{n}{}'' \sum_{j=0}^{n}{}'' \sum_{k=0}^{n}{}'' b_{i,j,k} T_i(x) T_j(y) T_k(z) \tag{39}$$

is the n^{th} order Chebyshev expansion for $f : [-1,1]^3 \to \mathbb{R}$. We let M denote the least nonnegative integer which is less than or equal to n and such that

$$|b_{i,j,k}| < \epsilon \text{ for all } i = M+1, \ldots, n, \quad j = 0, \ldots, n \text{ and } k = 0, \ldots, n, \tag{40}$$

assuming such an integer exists. If not, then we take $M = n$. For each $i = 0, \ldots, M$, we let m_i be the least nonnegative integer less than or equal to n such that

$$|b_{i,j,k}| < \epsilon \text{ for all } j = m_i + 1, \ldots, n, \quad \text{and } k = 0, \ldots, n \tag{41}$$

if such an integer exists, and we let $m_i = n$ otherwise. Finally, for each pair (i,j) such that $0 \le i \le M$ and $0 \le j \le m_i$, we let $n_{i,j}$ be the least nonnegative integer such that

$$|b_{i,j,k}| < \epsilon \text{ for all } k = 0, \ldots, n_{i,j}. \tag{42}$$

We refer to the series

$$\sum_{i=0}^{M} \sum_{j=0}^{m_i} \sum_{k=0}^{n_{i,j}} \widetilde{b}_{i,j,k} T_i(x) T_j(y) T_k(z), \tag{43}$$

where $\widetilde{b}_{i,j,k}$ is defined via

$$\begin{aligned} \widetilde{b}_{i,j,k} = b_{i,j,k} &\left(1 - \frac{1}{2}\delta(i-n)\right)\left(1 - \frac{1}{2}\delta(i-0)\right)\left(1 - \frac{1}{2}\delta(j-0)\right)\left(1 - \frac{1}{2}\delta(j-n)\right) \\ &\left(1 - \frac{1}{2}\delta(k-0)\right)\left(1 - \frac{1}{2}\delta(k-n)\right), \end{aligned} \tag{44}$$

as the ϵ-compressed n^{th} order Chebyshev expansion of f. Here, $\delta(x)$ denotes the function which is equal to 1 when $x = 0$ and is 0 for all other values of x. The factors in (44) involing δ are necessary because the first and last terms in (39) are scaled by $\frac{1}{2}$.

Obviously, the results discussed in this and the preceding section can be modified in a straightforward fashion so as to apply to functions given on an arbitrary rectangular prism $[a,b] \times [c,d] \times [e,f]$.

2.4. Series expansions of the associated Legendre functions and connection formulas

When $\nu \geq 0$, $-\nu \leq \mu \leq \nu$ and $-1 \leq x < 1$, the associated Legendre function of the first kind of degree ν and order $-\mu$ is given by

$$P_\nu^{-\mu}(x) = \left(\frac{1-x}{1+x}\right)^{\mu/2} \sum_{n=0}^{\infty} (-1)^n \frac{\Gamma(\nu+n+1)}{\Gamma(\nu-n+1)} \frac{\left(\frac{1}{2}-\frac{x}{2}\right)^n}{\Gamma(n+1)\Gamma(n+\mu+1)}. \tag{45}$$

Here, we have adopted the convention that

$$\frac{1}{\Gamma(k)} = 0 \tag{46}$$

whenever k is an integer which is less than or equal to 0. The trigonometric form

$$P_\nu^{-\mu}(\cos(t)) = \left(\tan\left(\frac{t}{2}\right)\right)^{\mu} \sum_{n=0}^{\infty} (-1)^n \frac{\Gamma(\nu+n+1)}{\Gamma(\nu-n+1)} \frac{\left(\sin\left(\frac{t}{2}\right)\right)^{2n}}{\Gamma(n+1)\Gamma(n+\mu+1)} \tag{47}$$

of (45) is obtained by letting $x = \cos(t)$ and making use of elementary identities. When the parameters ν and μ are of small magnitude, the coefficients in (47) decay rapidly as n increases, with the consequence that only a small number of terms of (47) are required to accurately evaluate $P_\nu^{-\mu}$. Likewise, even when the parameters are of large magnitude the coefficients in this expansion decay rapidly with n if t is sufficiently small, so that (47) is efficient in this regime as well. For extremely large values of ν, we found numerical roundoff error to be a problem in the evaluation of (47). For this reason, we only use this series expansion in the event that ν is less than $10,000$.

One futher potential difficulty with the use of (47) as a numerical tool is that underflow can occur when the parameters are large and t is small. To obviate such problems, we use a truncation of the formula

$$\begin{aligned} \log\left(\tilde{P}_\nu^{-\mu}(t)\right) = \log\left(\frac{\Gamma(\nu-\mu+1)}{\Gamma(\nu+\mu+1)}\right) + \log\left(\nu+\frac{1}{2}\right) + \frac{1}{2}\log\left(\sin(t)\right) + \mu\log\left(\tan\left(\frac{t}{2}\right)\right) \\ -\log\left(\frac{1}{\Gamma\left(\mu+1\right)}\right) + \log\left(\sum_{n=0}^{\infty} (-1)^n \frac{\Gamma(\nu+n+1)}{\Gamma(\nu-n+1)} \frac{\Gamma(\mu+1)}{\Gamma(\mu+n+1)} \frac{\left(\sin\left(\frac{t}{2}\right)\right)^{2n}}{\Gamma(n+1)}\right), \end{aligned} \tag{48}$$

which is easily obtained from (3) and (47), to evaluate the logarithm of the associated Legendre function of the first kind in this regime. We note that $\tilde{P}_\nu^{-\mu}(t)$ is necessarily positive when t is sufficiently small, so that this logarithm is sensible. When μ is equal to a negative integer, say $-m$, the first m terms of the sum in (47) are 0 and we use a version of (48) which is modified accordingly.

Remark 1. *The naive evaluation of the first term in (48) can lead to numerical cancellation when ν is large and μ is small relative to ν. In this event, we use the first sixteen terms of the asymptotic approximation*

$$\begin{aligned} \log\left(\frac{\Gamma(x-y)}{\Gamma(x+y)}\right) \sim -2y\log\left(x\right) + \log\left(1 - \frac{y}{x} - \frac{y\left(2y^2-3y+1\right)}{6x^2} + \frac{y^2\left(2y^2-3y+1\right)}{6x^3}\right. \\ \left. + \frac{y\left(20y^5-96y^4+155y^3-90y^2+5y+6\right)}{360x^4} + \cdots\right) \end{aligned}$$

in order to evaluate it.

For $\nu \geq 0$, $-1 \leq x < 1$ and $-\nu \leq \mu \leq \nu$ not an integer, the associated Legendre function of the second kind of degree ν and order $-\mu$ is given by

$$Q_\nu^{-\mu}(x) = \frac{\pi}{2}\left(\frac{\Gamma(\nu-\mu+1)}{\Gamma(\nu+\mu+1)}\sec(\mu\pi)P_\nu^{\mu}(x) - \cot(\mu\pi)P_\nu^{-\mu}(x)\right). \tag{49}$$

The normalized versions of the associated Legendre functions satisfy the somewhat simpler relation

$$\tilde{Q}_\nu^{-\mu}(t) = \sec(\mu\pi)\tilde{P}_\nu^{\mu}(t) - \cot(\mu\pi)\tilde{P}_\nu^{-\mu}(t), \tag{50}$$

which is an immediate consequence of (49), (3) and (4). Similarly to the case of (47), the use of (50) can lead to numerical overflow when t is small. Accordingly, we generally compute the logarithm of $\tilde{Q}_\nu^{-\mu}$ via the less delicate formula

$$\begin{aligned}\log\left(\tilde{Q}_\nu^{-\mu}(t)\right) = \log\left(\tilde{P}_\nu^{-\mu}(t)\right) + \log\Big(\sec(\mu\pi) - \cot(\mu\pi)\,\mathrm{sign}\left(\tilde{P}_\nu^{\mu}(t)\right) \\ \exp\left(\log\left(\left|\tilde{P}_\nu^{\mu}(t)\right|\right) - \log\left(\tilde{P}_\nu^{-\mu}(t)\right)\right)\Big)\end{aligned} \tag{51}$$

in this regime. We note that for sufficiently small t, the function $\tilde{Q}_\nu^{-\mu}(t)$ is positive.

When μ is an integer, (50) and (51) lose their meanings. Various series expansions for $Q_\nu^{-m}(x)$ with m a positive integer can be obtained (see, for instance, Section 3.6 of [7]), but they are somewhat cumbersome and do not address a second problem with the use of (50) as a numerical method for the evaluation of $Q_\nu^{-\mu}$. Namely, that when μ is close to, but does not coincide with, an integer, the evaluation (50) results in severe loss of precision due to numerical cancellation. However, since $\tilde{Q}_\nu^{-\mu}$ is an analytic function of the parameter μ, it can be efficiently interpolated in the μ variable. For instance, when μ is close to, or coincides with, an integer m, the value of $\tilde{Q}_\nu^{-\mu}(t)$ can be calculated by first evaluating

$$\tilde{Q}_\nu^{\xi_1}(t), \ldots, \tilde{Q}_\nu^{\xi_{2n+1}}(t) \tag{52}$$

with $\xi_1, \ldots, \xi_{2n}$ the nodes of the $(2n)$-point Chebyshev grid on the interval $[m-\epsilon], [m+\epsilon]$, and then using Chebyshev interpolation to calculate $\tilde{Q}_\nu^{-\mu}(t)$. Here, ϵ is an appropriate chosen positive real number and n is a positive integer. An even number of nodes is chosen in order to ensure that none coincide with the integer m. In the code used in this paper, we apply this procedure when μ is within a distance of 0.001 of an integer, and we take $n = 6$ and $\epsilon = 0.1$. Of course, the same approach can be used to evaluate $\log(\tilde{Q}_\nu^{-\mu}(t))$.

Indeed, many other connection formulas for the associated Legendre functions can be handled in a similar fashion, such as the identity

$$\tilde{Q}_\nu^{-\mu}(t) = \sec(\mu\pi)\tilde{Q}_\nu^{\mu}(t) + \tan(\mu\pi)\tilde{P}_\nu^{-\mu}(t), \tag{53}$$

which follows easily from a formula found in in Section 3.4 of [7]. On the other hand, the connection formulas

$$\tilde{P}_\nu^{-\mu}(\pi - t) = \cos(\pi(\nu-\mu))\tilde{P}_\nu^{-\mu}(t) - \sin(\pi(\nu-\mu))\tilde{Q}_\nu^{-\mu}(t) \tag{54}$$

and

$$\tilde{Q}_\nu^{-\mu}(\pi - t) = -\cos(\pi(\nu-\mu))\tilde{Q}_\nu^{-\mu}(t) - \sin(\pi(\nu-\mu))\tilde{P}_\nu^{-\mu}(t), \tag{55}$$

which also appear (in a slightly different form) in Section 3.4 of [7], are immune from such problems.

2.5. Macdonald's asymptotic expansions

In [19], an asymptotic formula for $P_\nu^{-\mu}(\cos(t))$ which is accurate when ν is large, $0 \leq \mu \leq \nu$ and t is small is derived by replacing the ratio of Gamma functions

$$\frac{\Gamma(\nu+n+1)}{\Gamma(\nu-n+1)} \tag{56}$$

appearing in (47) with a finite truncation of the series expansion

$$\frac{\Gamma(\nu+n+1)}{\Gamma(\nu-n+1)} = \lambda^{2n} - G_1\lambda^{2n-2} + G_2\lambda^{2n-4} - G_3\lambda^{2n-6} + \cdots, \tag{57}$$

where $\lambda = \nu + \frac{1}{2}$. The coefficients in this expansion can be easily derived starting from the formula

$$\frac{\Gamma\left(x+n+\frac{1}{2}\right)}{\Gamma\left(x-n+\frac{1}{2}\right)} = \frac{\left(4x^2-1^2\right)\left(4x^2-3^2\right)\cdots\left(4x^2-(2n-1)^2\right)}{4^n}, \tag{58}$$

which can be found (for instance) in Section 8.339 of [12]. The first few are

$$G_1 = \frac{1}{3}\frac{\Gamma(n+1)}{\Gamma(n-2)} + \frac{\Gamma(n+1)}{\Gamma(n-1)} + \frac{1}{4}\frac{\Gamma(n+1)}{\Gamma(n)},$$

$$G_2 = \frac{1}{18}\frac{\Gamma(n+1)}{\Gamma(n-5)} + \frac{11}{15}\frac{\Gamma(n+1)}{\Gamma(n-4)} + \frac{31}{12}\frac{\Gamma(n+1)}{\Gamma(n-3)} + \frac{29}{12}\frac{\Gamma(n+1)}{\Gamma(n-2)} + \frac{9}{32}\frac{\Gamma(n+1)}{\Gamma(n-1)},$$

and

$$\begin{aligned} G_3 = {} & \frac{\Gamma(n+1)}{162\Gamma(n-8)} + \frac{17\Gamma(n+1)}{90\Gamma(n-7)} + \frac{4943\Gamma(n+1)}{2520\Gamma(n-6)} + \frac{1513\Gamma(n+1)}{180\Gamma(n-5)} \\ & + \frac{1381\Gamma(n+1)}{96\Gamma(n-4)} + \frac{751\Gamma(n+1)}{96\Gamma(n-3)} + \frac{75\Gamma(n+1)}{128\Gamma(n-2)}. \end{aligned}$$

The first three terms in the asymptotic expansion of $P_\nu^{-\mu}$ obtained in this fashion are

$$P_\nu^{-\mu}(\cos(t)) \approx \left(\lambda\cos\left(\frac{t}{2}\right)\right)^{-\mu} \cdot \left(J_\mu(\eta) + \sin^2\left(\frac{t}{2}\right)H_1 + \sin^4\left(\frac{t}{2}\right)H_2 + \sin^6\left(\frac{t}{2}\right)H_3\right), \tag{59}$$

where $\eta = (2\nu+1)\sin(\frac{t}{2})$, J_γ denotes the Bessel function of the first kind of order γ,

$$H_1 = \frac{\eta}{6}J_{\mu+3}(\eta) - J_{\mu+2}(\eta) + \frac{1}{2\eta}J_{\mu+1}(\eta),$$

$$H_2 = \frac{\eta^2}{72}J_{\mu+6}(\eta) - \frac{11\eta}{30}J_{\mu+5}(\eta) + \frac{31}{12}J_{\mu+4}(\eta) - \frac{29}{6\eta}J_{\mu+3}(\eta) + \frac{9}{8\eta^2}J_{\mu+2}(\eta),$$

and

$$\begin{aligned} H_3 = {} & \frac{75}{16\eta^3}J_{\mu+2}(\eta) - \frac{751}{24\eta^2}J_{\mu+3}(\eta) + \frac{1381}{48\eta}J_{\mu+4}(\eta) - \frac{1513}{180}J_{\mu+5}(\eta) \\ & + \frac{4943\eta}{5040}J_{\mu+6}(\eta) - \frac{17\eta^2}{360}J_{\mu+7}(\eta) + \frac{\eta^3}{1296}J_{\mu+8}(\eta). \end{aligned}$$

The first few terms of the analogous expansion of the associated Legendre function of the second kind, which is applicable when $0 \leq \mu \leq \nu$, are

$$Q_\nu^\mu(\cos(t)) \approx -\frac{\pi}{2}\left(\lambda\cos\left(\frac{t}{2}\right)\right)^\mu \cdot \left(Y_{-\mu}(\eta) + \sin^2\left(\frac{t}{2}\right)N_1 + \sin^4\left(\frac{t}{2}\right)N_2 + \sin^6\left(\frac{t}{2}\right)N_3\right), \quad (60)$$

where Y_γ denotes the Bessel function of the second kind of order γ,

$$N_1 = \frac{\eta}{6}Y_{-\mu+3}(\eta) - Y_{-\mu+2}(\eta) + \frac{1}{2\eta}Y_{-\mu+1}(\eta),$$

$$N_2 = \frac{\eta^2}{72}Y_{-\mu+6}(\eta) - \frac{11\eta}{30}Y_{-\mu+5}(\eta) + \frac{31}{12}Y_{-\mu+4}(\eta) - \frac{29}{6\eta}Y_{-\mu+3}(\eta) + \frac{9}{8\eta^2}Y_{-\mu+2}(\eta),$$

and

$$\begin{aligned} N_3 = {} & \frac{75}{16\eta^3}Y_{-\mu+2}(\eta) - \frac{751}{24\eta^2}Y_{-\mu+3}(\eta) + \frac{1381}{48\eta}Y_{-\mu+4}(\eta) - \frac{1513}{180}Y_{-\mu+5}(\eta) \\ & + \frac{4943\eta}{5040}Y_{-\mu+6}(\eta) - \frac{17\eta^2}{360}Y_{-\mu+7}(\eta) + \frac{\eta^3}{1296}Y_{-\mu+8}(\eta). \end{aligned}$$

To reduce the potential for numerical underflow in the evaluation of (59), we evaluate $\log(\tilde{P}_\nu^{-\mu}(t))$ using the following formula instead:

$$\begin{aligned} \log\left(\tilde{P}_\nu^{-\mu}(t)\right) \approx {} & \log\left(\frac{\Gamma(\nu-\mu+1)}{\Gamma(\nu+\mu+1)}\right) + \log\left(\nu+\frac{1}{2}\right) + \frac{1}{2}\log\left(\sin(t)\right) - \mu\log\left(\lambda\cos\left(\frac{t}{2}\right)\right) \\ & + \log\left(J_\mu(\eta)\right) + \log\left(1 + \sin^2\left(\frac{t}{2}\right)\tilde{H}_1 + \sin^4\left(\frac{t}{2}\right)\tilde{H}_2 + \sin^6\left(\frac{t}{2}\right)\tilde{H}_3\right), \end{aligned} \quad (61)$$

where

$$\begin{aligned} \tilde{H}_1 = {} & \frac{\eta}{6}\exp\left(\log\left(J_{\mu+3}(\eta)\right) - \log\left(J_\mu(\eta)\right)\right) - \exp\left(\log\left(J_{\mu+2}(\eta)\right) - \log\left(J_\mu(\eta)\right)\right) \\ & + \frac{1}{2\eta}\exp\left(\log\left(J_{\mu+1}(\eta)\right) - \log\left(J_\mu(\eta)\right)\right), \end{aligned}$$

$$\begin{aligned} \tilde{H}_2 = {} & \frac{\eta^2}{72}\exp\left(\log\left(J_{\mu+6}(\eta)\right) - \log\left(J_\mu(\eta)\right)\right) - \frac{11\eta}{30}\exp\left(\log\left(J_{\mu+5}(\eta)\right) - \log\left(J_\mu(\eta)\right)\right) \\ & + \frac{31}{12}\exp\left(\log\left(J_{\mu+4}(\eta)\right) - \log\left(J_\mu(\eta)\right)\right) - \frac{29}{6\eta}\exp\left(\log\left(J_{\mu+3}(\eta)\right) - \log\left(J_\mu(\eta)\right)\right) \\ & + \frac{9}{8\eta^2}\exp\left(\log\left(J_{\mu+2}(\eta)\right) - \log\left(J_\mu(\eta)\right)\right), \end{aligned}$$

and

$$\begin{aligned}\tilde{H}_3 = &\frac{75}{16\eta^3}\exp\left(\log\left(J_{\mu+2}(\eta)\right)-\log\left(J_\mu(\eta)\right)\right)-\frac{751}{24\eta^2}\exp\left(\log\left(J_{\mu+3}(\eta)\right)-\log\left(J_\mu(\eta)\right)\right)\\ &+\frac{1381}{48\eta}\exp\left(\log\left(J_{\mu+4}(\eta)\right)-\log\left(J_\mu(\eta)\right)\right)-\frac{1513}{180}\exp\left(\log\left(J_{\mu+5}(\eta)\right)-\log\left(J_\mu(\eta)\right)\right)\\ &+\frac{4943\eta}{5040}\exp\left(\log\left(J_{\mu+6}(\eta)\right)-\log\left(J_\mu(\eta)\right)\right)\\ &-\frac{17\eta^2}{360}\exp\left(\log\left(J_{\mu+7}(\eta)\right)-\log\left(J_\mu(\eta)\right)\right)+\frac{\eta^3}{1296}\exp\left(\log\left(J_{\mu+8}(\eta)\right)-\log\left(J_\mu(\eta)\right)\right).\end{aligned}$$

As discussed in Remark 1, some care must be taken in evaluating the first term in (61). We use an analogous form of (60) in order to evaluate $\log(\tilde{Q}_\nu^\mu(t))$. The logarithms of the Bessel functions appearing in these formulas are calculated via the algorithm of [4].

We are unaware of a simple analog of (60) for the function $\tilde{Q}_\nu^{-\mu}$. When we say that we evaluate $\log(\tilde{Q}_\nu^{-\mu}(t))$ via Macdonald's expansions, we mean that we combine the above formulas for $\log(\tilde{P}_\nu^{-\mu}(t))$ and $\log(\tilde{Q}_\nu^{\mu}(t))$ with the connection formula (53) in order to calculate $\log(\tilde{Q}_\nu^{-\mu}(t))$.

2.6. Riccati's equation, Kummer's equation and phase functions

In this section, we suppose that q is a smooth, real-valued function defined on an open interval $I \subset \mathbb{R}$. In the event that q is strictly negative on I, two linearly independent solutions of the second order differential equation

$$y''(t) + q(t)y(t) = 0 \quad \text{for all} \quad t \in I \tag{62}$$

both of which are positive on I can be found. This follows easily from standard proofs of Picard's theorem on the existence and uniqueness of solutions of ordinary differential equations (see, for instance, Section 2.3 of [15]). Any positive solution y of (62) can be represented in the form $y = \exp(r(t))$ with r real-valued, and a straightforward computation shows that r must satisfy

$$r''(t) + (r'(t))^2 + q(t) = 0 \quad \text{for all} \quad t \in I. \tag{63}$$

Equation (63) is known as Riccati's equation; a detailed discussion of it can be found in [15], among many other sources.

When q is positive on I, the solutions of (62) oscillate and their logarithms are complex-valued. In this case it is convenient to represent them via a phase function. We say that a smooth function α defined on I is a phase function for the second order differential equation (62) provided α' does not vanish on I and the pair

$$u(t) = \frac{\cos\left(\alpha(t)\right)}{\sqrt{|\alpha'(t)|}} \tag{64}$$

and

$$v(t) = \frac{\sin\left(\alpha(t)\right)}{\sqrt{|\alpha'(t)|}} \tag{65}$$

form a basis in the space of solutions of (62). Of course, from (64) and (65), it is immediate that α is simply the imaginary part of the logarithm of the solution $u(t) + iv(t)$. We do not use a phase function to represent solutions of (62) in intervals on which q is negative since in that event the phase function cannot be efficiently represented via polynomial expansions (see Figure 1).

Proofs of the following elementary results regarding phase functions can be found in [13] and [4].

Theorem 1. *Suppose that I is an open interval in $\mathbb{R}$, and that q is a smooth, real-valued function defined on I. Suppose also that α is a smooth, real-valued function defined on I whose first derivative does not vanish there. Then α is a phase function for the second order differential equation (62) if and only if its derivative α' satisfies the second order nonlinear differential equation*

$$q(t) - (\alpha'(t))^2 - \frac{1}{2}\left(\frac{\alpha'''(t)}{\alpha'(t)}\right) + \frac{3}{4}\left(\frac{\alpha''(t)}{\alpha'(t)}\right)^2 = 0 \quad \textit{for all} \quad t \in I. \tag{66}$$

Theorem 2. *Suppose that u, v is a pair of smooth, real-valued solutions of (62) whose (necessarily constant) Wronskian W is nonzero. Then there is a phase function α for (62) such that*

$$u(t) = \sqrt{W}\frac{\cos(\alpha(t))}{\sqrt{|\alpha'(t)|}} \tag{67}$$

and

$$v(t) = \sqrt{W}\frac{\sin(\alpha(t))}{\sqrt{|\alpha'(t)|}}. \tag{68}$$

Moreover, the derivative of α is given by

$$\alpha'(t) = \frac{W}{(u(t))^2 + (v(t))^2} \quad \textit{for all} \quad t \in I, \tag{69}$$

and α is unique up to addition by an integer multiple of 2π. That is, $\tilde{\alpha}$ is a phase function for (62) such that (67) and (68) hold if and only there exists an integer L such that

$$\tilde{\alpha}(t) = \alpha(t) + 2\pi L \quad \textit{for all} \quad t \in I. \tag{70}$$

We will refer to (66) as Kummer's equation, after E. E. Kummer who studied it in [16].

2.7. A nonoscillatory phase function for the associated Legendre differential equation

From Theorem 2, we see that there is a phase function $\alpha_{\nu,\mu}$ for (5) such that

$$\tilde{P}_\nu^{-\mu}(t) = \sqrt{\frac{2\left(\nu + \frac{1}{2}\right)}{\pi}}\frac{\cos(\alpha_{\nu,\mu}(t))}{\sqrt{\alpha'_{\nu,\mu}(t)}} \tag{71}$$

and

$$\tilde{Q}_\nu^{-\mu}(t) = \sqrt{\frac{2\left(\nu + \frac{1}{2}\right)}{\pi}}\frac{\sin(\alpha_{\nu,\mu}(t))}{\sqrt{\alpha'_{\nu,\mu}(t)}}, \tag{72}$$

and whose derivative is given by

$$\alpha'_{\nu,\mu}(t) = \frac{2}{\pi}\left(\nu + \frac{1}{2}\right)\frac{1}{\left(\tilde{P}_\nu^{-\mu}(t)\right)^2 + \left(\tilde{Q}_\nu^{-\mu}(t)\right)^2}. \tag{73}$$

We have made use of the fact (which can be found in a slightly different form in Section 3.4 of [7]) that the Wronskian of the pair $\tilde{P}_\nu^{-\mu}, \tilde{Q}_\nu^{-\mu}$ is $\frac{2}{\pi}\left(\nu + \frac{1}{2}\right)$. It has long been known that the function

(73) is nonoscillatory. Indeed, it is immediate from (18) and (19) that

$$\alpha'_{\nu,\mu}(\arccos(x)) \sim \frac{2}{\pi}\sqrt{1-x^2}\left(\frac{\zeta-\gamma^2}{1-\gamma^2-x^2}\right)^{-\frac{1}{2}} \frac{1}{\left(J_\mu\left(\lambda\sqrt{\zeta}\right)\right)^2+\left(Y_\mu\left(\lambda\sqrt{\zeta}\right)\right)^2} \quad \text{as} \quad \nu\to\infty, \tag{74}$$

where ζ is the variable defined implicitly by (16) and (17), $\lambda = \nu + \frac{1}{2}$ and $\gamma = \frac{\mu}{\lambda}$. A cursory inspection of Nicholson's integral formula

$$J_\mu^2(z) + Y_\mu^2(z) = \frac{8}{\pi^2}\int_0^\infty K_0(2z\sinh(t))\cosh(2\mu t)\ dt, \tag{75}$$

a derivation of which can be found in Section 13.73 of [27], reveals that the function $(J_\mu(x))^2 + (Y_\mu(x))^2$ is nonoscillatory. We note that the nonoscillatory phase function for the associated Legendre differential equation is unique up to a constant. That is, all phase functions for the associated Legendre differential equation which are are not equal to $\alpha_{\nu,\mu} + C$ for some constant C oscillate. Figure 1 compares $\alpha'_{\nu,\mu}$ with the derivatives of oscillatory phase functions for (5).

It follows from (73) that there exists a constant C such that

$$\alpha_{\nu,\mu}(t) = C + \int_{\frac{\pi}{2}}^t \alpha'_{\nu,\mu}(s)\ ds. \tag{76}$$

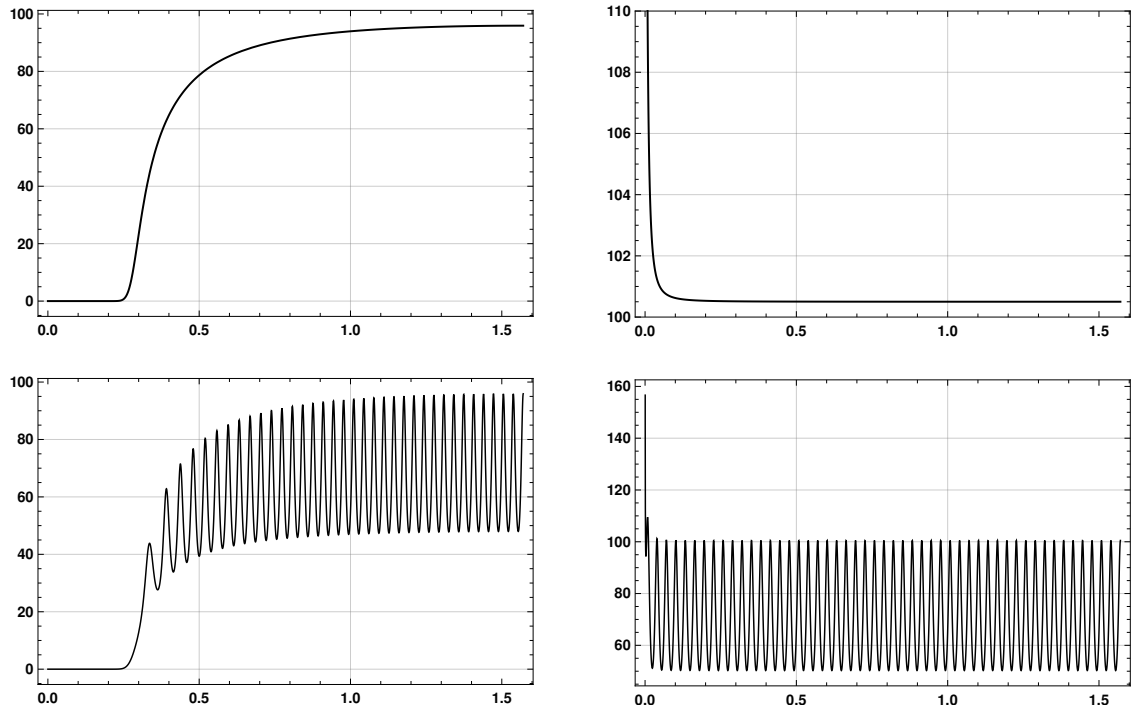

Figure 1: Top left: a plot of the function $\alpha'_{\nu,\mu}$ defined via (73) when $\nu = 100$ and $\mu = 30$. Top right: a plot of $\alpha'_{\nu,\mu}$ when $\nu = 100$ and $\mu = 0$. Bottom left: a plot of the derivative of a typical oscillatory phase function for (5) when $\nu = 100$ and $\mu = 30$. Bottom right: a plot of the derivative of a typical oscillatory phase function for (5) when $\nu = 100$ and $\mu = 0$.

In fact, using the formulas

$$\tilde{P}_\nu^{-\mu}\left(\frac{\pi}{2}\right) = \frac{2^{-\mu}}{\sqrt{\pi}}\sqrt{\left(\nu+\frac{1}{2}\right)\frac{\Gamma(\nu+\mu+1)}{\Gamma(\nu-\mu+1)}}\frac{\cos\left(\frac{1}{2}\pi(\nu-\mu)\right)\Gamma\left(\frac{1}{2}(-\mu+\nu+1)\right)}{\Gamma\left(\frac{1}{2}(\mu+\nu+2)\right)} \tag{77}$$

and

$$\tilde{Q}_\nu^{-\mu}\left(\frac{\pi}{2}\right) = -\frac{2^{-\mu}}{\sqrt{\pi}}\sqrt{\left(\nu+\frac{1}{2}\right)\frac{\Gamma(\nu+\mu+1)}{\Gamma(\nu-\mu+1)}}\frac{\sin\left(\frac{1}{2}\pi(\nu-\mu)\right)\Gamma\left(\frac{1}{2}(-\mu+\nu+1)\right)}{\Gamma\left(\frac{1}{2}(\mu+\nu+2)\right)}, \tag{78}$$

which appear in a slightly different form in Section 3.4 of [7], we see that (71) and (72) hold so long as the constant C in (76) is taken to be

$$C = \frac{\pi}{2}(\nu-\mu) + 2\pi L \tag{79}$$

with L an integer. In the remainder of this paper, we let $\alpha_{\nu,\mu}$ denote the phase function defined via the formula

$$\alpha_{\nu,\mu}(t) = 2\pi + \frac{\pi}{2}(\nu-\mu) + \int_{\frac{\pi}{2}}^{t} \alpha'_{\nu,\mu}(s)\, ds. \tag{80}$$

We set $L = 1$ in order to ensure that $\alpha_{\nu,\mu}$ is bounded away from 0 on the interval $\left(0, \frac{\pi}{2}\right)$. In this way, we avoid certain difficulties which arise because the condition number of evaluation of a function is generally infinite near one of its roots (as per the discussion in Section 2.1). By inserting (77) and (78) into (73), we see that

$$\alpha'_{\nu,\mu}\left(\frac{\pi}{2}\right) = \frac{2\Gamma\left(\frac{1}{2}(\nu-\mu+2)\right)\Gamma\left(\frac{1}{2}(\nu+\mu+2)\right)}{\Gamma\left(\frac{1}{2}(\nu-\mu+1)\right)\Gamma\left(\frac{1}{2}(\nu+\mu+1)\right)}. \tag{81}$$

Expressions for the values of the derivatives of the functions $\tilde{P}_\nu^{-\mu}$ and $\tilde{Q}_\nu^{-\mu}$ at the point $\frac{\pi}{2}$ can be easily derived from formulas appearing in Section 3.4 of [7]. A tedious computation which makes use of them in addition to (77) and (78) shows that

$$\alpha''_{\nu,\mu}\left(\frac{\pi}{2}\right) = 0. \tag{82}$$

2.8. An adaptive discretization procedure

We now briefly describe a fairly standard procedure for adaptively discretizing a smooth function $f : [a,b] \to \mathbb{R}$. It takes as input a desired precision $\epsilon > 0$, a positive integer n and a subroutine for evaluating f. The goal of this procedure is to construct a partition

$$a = \gamma_0 < \gamma_1 < \cdots < \gamma_m = b \tag{83}$$

of $[a,b]$ such that the n^{th} order Chebyshev expansion of f on each of the subintervals $[\gamma_j, \gamma_{j+1}]$ of $[a,b]$ approximates f with accuracy ϵ. That is, for each $j = 0, \ldots, m-1$ we aim to achieve

$$\sup_{x\in[\gamma_j,\gamma_{j+1}]}\left|f(x) - {\sum_{i=0}^{n}}'' b_{i,j} T_i\left(\frac{2}{\gamma_{j+1}-\gamma_j}x + \frac{\gamma_{j+1}+\gamma_j}{\gamma_j-\gamma_{j+1}}\right)\right| < \epsilon, \tag{84}$$

where $b_{0,j}, b_{1,j} \ldots, b_{n,j}$ are the coefficients in the n^{th} order Chebyshev expansion of f on the interval $[\gamma_j, \gamma_{j+1}]$. These coefficients are defined by the formula

$$b_{i,j} = \frac{2}{n}{\sum_{l=0}^{n}}'' T_i(\rho_{l,n})\, f\left(\frac{\gamma_j-\gamma_{j+1}}{2}\cos\left(\frac{\pi l}{n}\right) + \frac{\gamma_{j+1}+\gamma_j}{2}\right). \tag{85}$$

During the procedure, two lists of subintervals are maintained: a list of subintervals which are to be processed and a list of output subintervals. Initially, the list of subintervals to be processed consists of $[a, b]$ and the list of output subintervals is empty. The procedure terminates when the list of subintervals to be processed is empty or when the number of subintervals in this list exceeds a present limit (we usually take this limit to be 300). In the latter case, the procedure is deemed to have failed. As long as the list of subintervals to process is nonempty and its length does not exceed the preset maximum, the algorithm proceeds by removing a subinterval $[\eta_1, \eta_2]$ from that list and performing the following operations:

1. Compute the coefficients $b_0, \ldots, b_n$ in the n^{th} order Chebyshev expansion of the restriction of f to the interval $[\eta_1, \eta_2]$.

2. Compute the quantity
$$\Delta = \frac{\max\left\{\left|b_{\frac{n}{2}+1}\right|, \left|b_{\frac{n}{2}+2}\right|, \ldots, |b_n|\right\}}{\max\left\{|b_0|, |b_1|, \ldots, |b_n|\right\}}. \tag{86}$$

3. If $\Delta < \epsilon$ then the subinterval $[\eta_1, \eta_2]$ is added to the list of output subintervals.

4. If $\Delta \geq \epsilon$, then the subintervals
$$\left[\eta_1, \frac{\eta_1 + \eta_2}{2}\right] \quad \text{and} \quad \left[\frac{\eta_1 + \eta_2}{2}, \eta_2\right] \tag{87}$$
are added to the list of subintervals to be processed.

This algorithm is heuristic in the sense that there is no guarantee that (84) will be achieved, but similar adaptive discretization procedures are widely used with great success.

There is one common circumstance which leads to the failure of this procedure. The quantity Δ is an attempt to estimate the relative accuracy with which the Chebyshev expansion of f on the interval $[\eta_1, \eta_2]$ approximates f. In cases in which the condition number of the evaluation of f is larger than ϵ on some part of $[a, b]$, the procedure will generally fail or an excessive number of subintervals will be generated. Particular care needs to be taken when f has a zero in $[a, b]$. In most cases, for x near a zero of f, the condition number of evaluation of $f(x)$ (as defined in Section 2.1) is large. In this article, we avoid such difficulties by only applying this procedure to functions which are bounded away from 0.

3. A method for the rapid numerical solution of the associated Legendre differential equation

In this section, we describe an algorithm for the numerical solution of the associated Legendre differential equation which runs in time independent of ν and μ. It is a crucial component of the scheme of the following section for the construction of a table which allows for the rapid numerical evaluation of the associated Legendre functions.

The algorithm makes use of a solver for nonlinear second order ordinary differential equations of the form
$$y''(t) = f(t, y(t), y'(t)) \quad \text{for all} \quad a < t < b \tag{88}$$

which is described in detail in Section 4 of [4]. That solver is designed to be extremely robust, but not necessarily highly efficient. It takes as input a subroutine for evaluating the function f and its derivatives with respect to t, y and y', a positive integer k, a precision $\epsilon > 0$ for the calculations, and either initial or terminal conditions for the desired solution y. It returns a collection of subintervals

$$[\gamma_1, \gamma_2], \ldots, [\gamma_{m-1}, \gamma_m] \tag{89}$$

and the values of the functions y, y' and y'' at the $(k+1)$-point Chebyshev grid on each of the subintervals (89). In particular, the functions y, y' and y'' are represented via piecewise kth order Chebyshev expansions. Given this data, the value of any one of these functions at any point on the interval (a, b) can be computed using Chebyshev interpolation (see, for instance, [25] for a thorough discussion of such techniques). The collection of subintervals is determined adaptively in the course of solving (88) using an approach which attempts to achieve relative accuracy in the expansions of y, y' and y'' on the order of the specified precision ϵ. The algorithm is heuristic and offers no accuracy guarantees, but similar approaches are commonly used with great success. There is one situation in which this solver is likely to fail. When the condition number of evaluation of the solution of (89) is large, it is not possible to represent it with high relative accuracy using Chebyshev expansions. In this event, the solver tends to produce an excessive number of subintervals or fail altogether. Since the condition number of evaluation of a function is generally large near one of its roots, we only apply this solver in cases in which the solution is bounded away from 0.

Our algorithm for the numerical solution of (5) takes as input real numbers ν and μ such that $0 \leq \mu \leq \nu$, a desired precision $\epsilon > 0$, and a positive integer k specifying the order of the Chebyshev expansions to use. It proceeds in three stages.

Stage one: computation of the nonoscillatory phase function $\alpha_{\nu,\mu}$

In this stage, we calculate the values of the nonoscillatory phase function (80) on the interval

$$\left(t^*_{\nu,\mu}, \frac{\pi}{2}\right), \tag{90}$$

where

$$t^*_{\nu,\mu} = \arcsin\left(\frac{\sqrt{\mu^2 - \frac{1}{4}}}{\nu + \frac{1}{2}}\right) \tag{91}$$

is the turning point of (5) if $\mu \geq 1$ and

$$t^*_{\nu,\mu} = \frac{1}{\nu^{\frac{3}{2}}} \tag{92}$$

if $0 \leq \mu < 1$. The rationale for using (92) as the left endpoint for the interval on which the phase function is calculated when $0 < \mu < 1$ is to avoid a discontinuity in $t_{\nu,\mu}$ when μ crosses the threshold $\mu = \frac{1}{2}$.

We first construct $\alpha'_{\nu,\mu}$ by solving a terminal value problem for Kummer's equation (66) using the solver of Section 4 of [4]. The values of $\alpha'_{\nu,\mu}$ and $\alpha''_{\nu,\mu}$ at the point $\frac{\pi}{2}$ are obtained using (81) and (82). The required precision for these computations is taken to be ϵ. Next, $\alpha_{\nu,\mu}$ is constructed through Formula (80). Since $\alpha_{\nu,\mu}$ is represented via its values at the $(k+1)$-point Chebyshev nodes on a collection of intervals, it is easy to evaluate the required integral via spectral integration.

Upon the completion of this stage, the values of $\alpha_{\nu,\mu}$ and its first two derivatives are known at the

nodes of the $(k+1)$-point Chebyshev points on each interval in a collection of subintervals of (90). Using standard Chebyshev interpolation methods, the values of these functions can be calculated in a stable fashion anywhere on the interval (90).

Stage two: computation of $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$

In the event that $\mu \geq 1$, we calculate the function $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$ on the interval

$$\left(\frac{t^*_{\nu,\mu}}{100}, t^*_{\nu,\mu}\right) \tag{93}$$

by solving a terminal boundary value problem for Riccati's equation (63) using the solver described in Section 4 of [4]. In fact, we solve the terminal boundary value problem on the slightly larger interval

$$\left(\frac{t^*_{\nu,\mu}}{100}, t_0\right), \tag{94}$$

where t_0 is the solution of the nonlinear equation

$$\alpha_{\nu,\mu}(t_0) = \frac{\pi}{4} + 2\pi. \tag{95}$$

The functions $\alpha_{\nu,\mu}$ and its derivative having been calculated in the preceding stage, there is no difficulty in using Newton's method to solve (95). From (72) and (95), we see that

$$\tilde{Q}_\nu^{-\mu}(t_0) = \sqrt{\frac{\left(\nu + \frac{1}{2}\right)}{\pi \alpha'_{\nu,\mu}(t)}} \tag{96}$$

and

$$\frac{d\tilde{Q}_\nu^{-\mu}}{dt}(t_0) = \sqrt{\frac{\nu + \frac{1}{2}}{\pi}} \left(\sqrt{\alpha'(t_0)} - \frac{\alpha''(t_0)}{2\left(\alpha'(t_0)\right)^{\frac{3}{2}}}\right). \tag{97}$$

The rationale for introducing t_0 is to ensure that the terminal value of $\tilde{Q}_\nu^{-\mu}$ and its derivative used in the solution of Riccati's equation are computed accurately. The condition number of evaluation of the function $\tilde{Q}_\nu^{-\mu}$ is large when the parameters ν and μ are of large magnitude, with the consequence that its numerical evaluation will generally result in a loss of precision in this event. In the case of (72), the evaluation of a trigonometric functions at a large argument is the specific mechanism by which this loss of precision takes place. By evaluating $\tilde{Q}_\nu^{-\mu}$ at a point t_0 at which the value of the phase function is known, however, we avoid this loss of precision entirely. This can be seen from (96) and (97). They involve only the evaluation of $\alpha'_{\nu,\mu}$ and $\alpha''_{\nu,\mu}$, the condition number of evaluation of which is small independent of ν and μ.

We construct $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$ in lieu of $\log(\tilde{Q}_\nu^{-\mu}(t))$ because the former is bounded away from 0 on the interval (93) while the latter is not. We note that the addition of this constant has no impact on the accuracy with which values of $Q_\nu^{-\mu}$ are computed. It is strictly a mechanism for avoiding numerical difficulties in our solver for ordinary differential equations and in the construction of our precomputed expansions. Upon the completion of this stage, the values of $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$ at the nodes of the $(k+1)$-point Chebyshev points on each of a collection of subintervals which cover (90) are known. Using standard Chebyshev interpolation methods, the values of this function can be calculated in a stable fashion anywhere on the interval (93).

Stage three: computation of $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$

Assuming that $\mu \geq 1$, we now compute $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$ on the interval (93). Proceeding here as we did in the calculation of $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$ would be problematic. Unlike $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$, which increases rapidly as t goes to 0 from the right, $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$ decreases rapidly as t goes to 0 from the right. Consequently, it is recessive when solving Riccati's equation in the backward direction and attempts to approximate it numerically by solving a terminal value problem for Riccati's equation lead to excessively large errors.

Instead, we solve an initial value problem for Riccati's equation on the interval (93) in order to calculate $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$. This is numerically viable since it is a dominant solution of Riccati's equation when solving in the forward direction. When $\nu \leq 10,000$, we use a truncation of the series expansion (48) in order to generate the necessary initial values. For $\nu > 10,000$, we calculate initial values via (61) instead since (48) can lead to numerical roundoff errors when ν is large.

As in the case of the function of the second kind, the reason for computing $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$ in lieu of $\log(\tilde{P}_\nu^{-\mu}(t))$ is that the former is bounded away from 0 on the interval (93) while the latter is not. Upon the completion of this stage, the values of $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$ at the nodes of the $(k+1)$-point Chebyshev points on each interval in a collection of subintervals of (90) are known. Using standard Chebyshev interpolation methods, the values of this function can be calculated in a stable fashion anywhere on the interval (93).

4. The numerical construction of the precomputed table

In this section, we describe the procedure used to construct the precomputed table which allows for the rapid numerical evaluation of the associated Legendre functions $\tilde{P}_\nu^{-\mu}$ and $\tilde{Q}_\nu^{-\mu}$ for a large range of ν, μ and t. This table stores the coefficients in the compressed piecewise trivariate Chebyshev expansions of eights pair of functions.

A first pair of functions A_1, B_1 allows for the evaluation of the phase function $\alpha_{\nu,\mu}$ and its derivative on the subset

$$\mathcal{O}_1 = \left\{(\nu, \mu, t) : 10 \leq \nu \leq 1,000,000, \quad 1 \leq \mu \leq \nu \text{ and } \quad t^*_{\nu,\mu} \leq t \leq \frac{\pi}{2}\right\} \tag{98}$$

of the oscillatory region $\mathcal{O}$. Here, $t^*_{\nu,\mu}$ is as in (91) and (92). A second pair A_2, B_2 allows for the evaluation of the phase function $\alpha_{\nu,\mu}$ and its derivative on the subset

$$\mathcal{O}_2 = \left\{(\nu, \mu, t) : 10 \leq \nu \leq 1,000,000, \quad 0 \leq \mu < 1 \text{ and } \quad t^*_{\nu,\mu} \leq t \leq \frac{\pi}{2}\right\} \tag{99}$$

of the oscillatory region $\mathcal{O}$. The functions A_3 and B_3 allow for the evaluation of $\alpha_{\nu,\mu}$ and $\alpha'_{\nu,\mu}$ on

$$\mathcal{O}_3 = \left\{(\nu, \mu, t) : 2 \leq \nu \leq 10, \quad 1 \leq \mu \leq \nu \text{ and } \quad t^*_{\nu,\mu} \leq t \leq \frac{\pi}{2}\right\}, \tag{100}$$

and a fourth pair A_4, B_4 allows for the evaluation of the phase function $\alpha_{\nu,\mu}$ and its derivative on

$$\mathcal{O}_4 = \left\{(\nu, \mu, t) : 2 \leq \nu \leq 10, \ 0 \leq \mu < 1 \text{ and } \quad t^*_{\nu,\mu} \leq t \leq \frac{\pi}{2}\right\}. \tag{101}$$

We divided the range of the parameter ν because it is more efficient to represent $\alpha_{\nu,\mu}$ via polynomial expansions in $\frac{1}{\nu}$ when ν is large, and via expansions in ν when ν is small.

A fifth set of functions C_1 and D_1 allows for the evaluation of the functions

$$\log\left(\tilde{P}_\nu^{-\mu}(t)\right) - \nu \quad \text{and} \quad \log\left(\tilde{Q}_\nu^{-\mu}(t)\right) + \nu \tag{102}$$

on the subset

$$\mathcal{N}_1 = \left\{(\nu, \mu, t) : 10 \leq \nu \leq 1,000,000, \quad 1 < \mu \leq \nu \text{ and } \quad 0 \leq t < t^*_{\nu,\mu}\right\} \tag{103}$$

of the nonoscillatory region $\mathcal{N}$. A sixth pair of functions C_2 and D_2 allows for the evaluation of the functions (102) on

$$\mathcal{N}_2 = \left\{(\nu, \mu, t) : 10 \leq \nu \leq 1,000,000, \quad 0 \leq \mu < 1 \text{ and } \quad 0 \leq t < t^*_{\nu,\mu}\right\}. \tag{104}$$

The seventh pair of functions C_3, D_3 allows for the evaluation of (102) on

$$\mathcal{N}_3 = \left\{(\nu, \mu, t) : 2 \leq \nu \leq 10, \quad 1 \leq \mu \leq \nu \text{ and } \quad 0 \leq t < t^*_{\nu,\mu}\right\}. \tag{105}$$

The eighth and final pair of functions C_4, D_4 allows for the evaluation of the functions (102) on

$$\mathcal{N}_4 = \left\{(\nu, \mu, t) : 2 \leq \nu \leq 10, \quad 0 \leq \mu < 1 \text{ and } \quad 0 \leq t < t^*_{\nu,\mu}\right\}. \tag{106}$$

We construct expansions of the functions (102) rather than expansions of $\log(\tilde{P}_\nu^{-\mu}(t))$ and $\log(\tilde{Q}_\nu^{-\mu}(t))$ because the former are bounded away from 0 on the sets in which we consider them while the latter are not. This ensures that their condition number of evaluation is not large because of the presence of roots.

These computations were conducted in IEEE quadruple precision arithmetic in order to ensure high accuracy. The resulting table, which consists of the coefficients in the expansions of the functions $A_1, \ldots, A_4$, $B_1, \ldots, B_4$, $C_1, \ldots, C_4$, $D_1, \ldots, D_4$ is roughly 138 MB in size. The precomputed table allows for the evaluation of $\alpha_{\nu,\mu}$, $\alpha'_{\nu,\mu}$, $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$ and $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$ with roughly double precision accuracy (see the experiments of Section 6). The code was written in Fortran with OpenMP extensions and compiled with version 4.8.4 of the GNU Fortran compiler. It was executed on a computer equipped with 28 Intel Xeon E5-2697 processor cores running at 2.6 GHz. The construction of the table took approximately 24 hours on this machine.

Here, we describe only the construction of the functions A_1, B_1, C_1 and D_1. The construction of the others is extremely similar. The procedure proceeded in four stages as follows:

Stage one: construction of the phase functions and logarithms

We began this stage of the procedure by constructing a partition

$$\xi_1 < \xi_2 < \ldots < \xi_{11} \tag{107}$$

which divides the interval

$$\left[\frac{1}{1,000,000}, \frac{1}{10}\right] \tag{108}$$

over which $\frac{1}{\nu}$ is allowed to vary into 10 subintervals. The precise locations of the nodes ξ_j are not critically important; reasonable choices were arrived at quickly through trial and error. Next, we constructed a partition

$$0 = \tau_1 < \tau_2 < \ldots < \tau_{18} = 1 \tag{109}$$

which divides the interval $[0, 1]$ into 17 subintervals. Again, the precise distribution of the nodes τ_j is not critical and reasonable choices were arrived at quickly through trial and error.

For each $i = 1, \ldots, 10$ and $j = 1, \ldots, 17$, we processed the tensor product of intervals $[\xi_i, \xi_{i+1}] \times [\tau_j, \tau_{j+1}]$ as follows. We let $x_1, \ldots, x_{31}$ be the nodes of the 31-point Chebyshev grid on the interval $[\xi_i, \xi_{i+1}]$ and $y_1, \ldots, y_{31}$ the nodes of the 31-point Chebyshev grid on the interval $[\tau_j, \tau_{j+1}]$. For each pair x_k, y_l, the algorithm of Section 3 was used to calculate $\alpha_{\nu,\mu}$, $\alpha'_{\nu,\mu}$, $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$, and $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$ with ν and μ taken to be

$$\nu = \frac{1}{x_k} \quad \text{and} \quad \mu = 1 + (\nu - 1)y_l. \tag{110}$$

The precision for the computations was $\epsilon = 10^{-17}$. The functions $\alpha_{\nu,\mu}$ and $\alpha'_{\nu,\mu}$ were represented as 30th order piecewise Chebyshev expansions on some adaptively determined collection of subintervals of $\left[t^*_{\nu,\mu}, \frac{\pi}{2}\right]$, while the functions $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$ and $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$ were represented as 30th order piecewise Chebyshev expansions on some adaptively determined collection of subintervals of $\left[0, t^*_{\nu,\mu}\right]$. Using this data, the nonoscillatory phase function and its derivative can be evaluated for any triple (ν, μ, t) in the region $\mathcal{O}_1$ via Chebyshev interpolation. Likewise, the logarithms of the associated Legendre functions can be evaluated at any point in $\mathcal{N}_1$.

In this stage of the procedure, the differential equation (5) was solved via the algorithm of Section 3 for 163,370 different pairs of the parameters (ν, μ), many of which were large in magnitude. Obviously, this was only possible because our solver runs in time independent of ν and μ.

Stage two: formation of unified discretizations

For each pair of points ξ and τ such that ξ is one of the Chebyshev nodes in one of the subintervals defined by the partition (107) and τ is one of the Chebyshev nodes in one of the subintervals defined by the partition (109), we used the procedure of Section 2.8 to adaptively form discretization of the functions

$$f_{\xi,\tau}(u) = \alpha'_{\nu,\mu}(t(u)) \tag{111}$$

and

$$g_{\xi,\tau}(u) = \alpha_{\nu,\mu}(t(u)), \tag{112}$$

where

$$\nu = \frac{1}{\xi}, \quad \mu = 1 + (\nu - 1)\tau \quad \text{and} \quad t(u) = t^*_{\nu,\mu} + \left(\frac{\pi}{2} - t^*_{\nu,\mu}\right) u. \tag{113}$$

The functions $\alpha_{\nu,\mu}$ and $\alpha'_{\nu,\mu}$ are evaluated via Chebyshev interpolation using the data constructed in the first stage of these calculations. We requested $\epsilon = 10^{-17}$ accuracy and took the parameter n to be 30. For each ξ and τ considered, this results in a collection of subintervals of $[0, 1]$ on which $f_{\xi,\tau}$ is represented with relative accuracy roughly ϵ via a 30th order Chebyshev expansion and another collection of subintervals of $[0, 1]$ on which $g_{\xi,\tau}$ is represented with relative accuracy roughly ϵ via a 30th order Chebyshev expansion. We then formed a unified discretization

$$[a_0, a_1], \ [a_1, a_2], \ [a_2, a_3], \ \ldots, \ [a_{20}, a_{21}] \tag{114}$$

of $[0, 1]$ by merging these discretizations; that is, by ensuring that each subinterval in the discretization of one of the functions $f_{\xi,\tau}$ or $g_{\xi,\tau}$ is the union of some set of subintervals of (114).

A unified discretization

$$[b_0, b_1], \ [b_1, b_2], \ [b_2, b_3], \ \ldots, \ [b_{17}, b_{18}] \tag{115}$$

for the functions

$$\tilde{f}_{\xi,\tau}(u) = \log\left(\tilde{P}_{\nu}^{-\mu}(t)\right) - \nu \tag{116}$$

and

$$\tilde{g}_{\xi,\tau}(u) = \log\left(\tilde{Q}_{\nu}^{-\mu}(t)\right) + \nu \tag{117}$$

with ν, μ and t related to ξ, τ and u via (113) was formed in the same fashion.

Stage three: Construction of the functions A_1 and B_1

The function A_1 is defined via the formula

$$A_1(\xi,\tau,u) = \frac{1}{\nu}\alpha_{\nu,\mu}(t), \tag{118}$$

where

$$\nu = \frac{1}{\xi}, \quad \mu = 1 + (\nu-1)\tau \quad \text{and} \quad t = t^*_{\nu,\mu} + \left(\frac{\pi}{2} - t^*_{\nu,\mu}\right)u. \tag{119}$$

Likewise, B_1 is defined via

$$B_1(\xi,\tau,u) = \frac{1}{\nu}\alpha'_{\nu,\mu}(t), \tag{120}$$

with ν, μ and t given by (119). In this way, we ensure that A_1 and B_1 are defined on the rectangular prism

$$\left[\frac{1}{1,000,000}, \frac{1}{10}\right] \times [0,1] \times [0,1], \tag{121}$$

and hence suitable for representation via a collection of piecewise trivariate Chebyshev expansions.

For each $i = 1,\ldots,10$, $j = 1,\ldots,17$ and $k = 1,\ldots,21$, we formed the 30th order compressed trivariate Chebyshev expansions (as defined in Section 2.3) for the functions A_1 and B_1 on the rectangular prism

$$[\xi_i, \xi_{i+1}] \times [\tau_j, \tau_{j+1}] \times [a_k, a_{k+1}]. \tag{122}$$

There are $3,570$ such rectangular prisms. Since an uncompressed 30th order trivariate Chebyshev expansion has $29,791$ coefficients, a total of $212,707,740$ coefficients would be required to store the uncompressed Chebyshev expansions of the functions A_1 and B_1. If each coefficient were stored as an IEEE double precision number, roughly 1.5 GB of memory would be required to store these expansions. Fortunately, the compressed 30th order Chebyshev expansions were far more efficient. The compressed Chebyshev expansions for A_1 and B_1 had only $7,839,620$ coefficients.

Stage four: construction of the functions C_1 and D_1

The function C_1 is defined via

$$C_1(\xi,\tau,u) = \log\left(\tilde{P}_{\nu}^{-\mu}(t)\right) - \nu, \tag{123}$$

where

$$\nu = \frac{1}{\xi}, \quad \mu = 1 + (\nu-1)\tau \quad \text{and} \quad t = \frac{t^*_{\nu,\mu}}{100} + \left(t^*_{\nu,\mu} - \frac{t^*_{\nu,\mu}}{100}\right)u. \tag{124}$$

Finally, D_1 is defined via

$$D_1(\xi, \tau, u) = \log\left(\tilde{Q}_\nu^{-\mu}(t)\right) + \nu, \tag{125}$$

with ν, μ and t as in (124). Obviously, C_1 and D_1 are also given on the rectangular prism (121).

For each $i = 1, \ldots, 10$, $j = 1, \ldots, 17$ and $k = 1, \ldots, 17$, we formed the 30th order compressed trivariate Chebyshev expansions (as defined in Section 2.3) for the functions C_1 and D_1 on the rectangular prism

$$[\xi_i, \xi_{i+1}] \times [\tau_j, \tau_{j+1}] \times [b_k, b_{k+1}]\,. \tag{126}$$

There are $2,890$ such rectangular prisms and a total of $172,191,980$ coefficients would be required to store the uncompressed Chebyshev expansions of the functions C_1 and D_1. If each coefficient were stored as an IEEE double precision number, roughly 1.3 GB of memory would be required to store these expansions. The compressed 30th order Chebyshev expansions were far more efficient. They required only $4,441,063$ coefficients to store C_1 and D_1.

5. An algorithm for the rapid numerical evaluation of the associated Legendre functions

In this section, we describe the operation of our code for evaluating the associated Legendre functions $\tilde{P}_\nu^{-\mu}(t)$ and $\tilde{Q}_\nu^{-\mu}(t)$ when

$$0 \le \nu \le 1,000,000,\ 0 \le \mu \le \nu \ \text{ and } \ 0 < t \le \frac{\pi}{2}. \tag{127}$$

The code was written in Fortran and its interface to the user consists of two subroutines, one called `alegendre_eval_init` and the other `alegendre_eval`. The `alegendre_eval_init` routine reads the precomputed table constructed via the procedure of Section 4 from the disk into memory. The precomputed table used in the experiments described in this paper is approximately 138 megabytes in size. Once the precomputed table has been loaded, the `alegendre_eval` can be called. It takes as input a triple (ν, μ, t) satisfying the conditions (127). When (ν, μ, t) is in the oscillatory region $\mathcal{O}$, it returns the values of $\alpha_{\nu,\mu}(t)$ and $\alpha'_{\nu,\mu}(t)$ as well as those of $\tilde{P}_\nu^{-\mu}(t)$ and $\tilde{Q}_\nu^{-\mu}(t)$. When (ν, μ, t) is in the nonoscillatory region $\mathcal{N}$, it returns the values of $\log(\tilde{P}_\nu^{-\mu}(t))$ and $\log(\tilde{Q}_\nu^{-\mu}(t))$ as well as those of $\tilde{P}_\nu^{-\mu}(t)$ and $\tilde{Q}_\nu^{-\mu}(t)$. Of course, when t is close to 0, the latter values might not be representable via the IEEE double format arithmetic. In this event, 0 is returned for $\tilde{P}_\nu^{-\mu}(t)$ and ∞ for $\tilde{Q}_\nu^{-\mu}(t)$.

The `alegendre_eval` code is available from the GitHub repository at address

`http://github.com/JamesCBremerJr/ALegendreEval`.

It uses several different methods to evaluate the associated Legendre functions and the associated auxiliary functions, depending on the values of ν, μ and t. The following description of the operation of the `alegendre_eval` code is organized by listing each such method.

Method one: series expansions for $\tilde{P}_\nu^{-\mu}(t)$ and $\tilde{Q}_\nu^{-\mu}(t)$

This method is used when $\nu < 2$ and (ν, μ, t) is in the oscillatory region $\mathcal{O}$. It consists of evaluating $\tilde{P}_\nu^{-\mu}(t)$ via a truncation of the series expansion (47) and evaluating $\tilde{Q}_\nu^{-\mu}(t)$ via formula (50). As discussed in Section 2.4, when μ is close to or coincides with an integer, Chebyshev interpolation

in the variable μ is used to avoid roundoff error in the evaluation of (50). The value of $\alpha'_{\nu,\mu}(t)$ is calculated via (73) and $\alpha_{\nu,\mu}$ is computed using the formula

$$\alpha_{\nu,\mu} = \mathrm{Arg}\left(\tilde{P}_\nu^{-\mu}(t) + i\tilde{Q}_\nu^{-\mu}(t) \right) + 2\pi, \tag{128}$$

where $\mathrm{Arg}(z)$ denotes the principal value of the argument of the complex number z. The limitation on the range of parameters for which this method is used ensures that the principal branch of the argument function is the correct one.

Method two: series expansions for $\log(\tilde{P}_\nu^{-\mu}(t))$ *and* $\log(\tilde{Q}_\nu^{-\mu}(t))$ This method is used when $\nu < 10$ and t is in the nonoscillatory regime, and when $10 \leq \nu < 10,000$ and

$$0 < t < \frac{t^*_{\nu,\mu}}{100}. \tag{129}$$

It consists of evaluating $\log(\tilde{P}_\nu^{-\mu}(t))$ via a truncation of (48) and $\log(\tilde{Q}_\nu^{-\mu}(t))$ via (51). When μ coincides with or is close to an integer, Chebyshev interpolation in the parameter μ is used in the evaluation of (51). The values of $\tilde{P}_\nu^{-\mu}(t)$ and $\tilde{Q}_\nu^{-\mu}(t)$ are computed from their logarithms in the obvious fashion.

Method three: Macdonald's asymptotic expansions for $\log(\tilde{P}_\nu^{-\mu}(t))$ *and* $\log(\tilde{Q}_\nu^{-\mu}(t))$

This method is used when $\nu \geq 10,000$ and

$$0 < t < \frac{t^*_{\nu,\mu}}{100}. \tag{130}$$

It consists of evaluating $\log(\tilde{P}_\nu^{-\mu}(t))$ and $\log(\tilde{Q}_\nu^{-\mu})$ via Macdonald's asymptotic expansions (see Section 2.5). The values of $\tilde{P}_\nu^{-\mu}(t)$ and $\tilde{Q}_\nu^{-\mu}(t)$ are computed from their logarithms in the obvious fashion.

Method four: precomputed expansions

In all other cases, the precomputed expansions of the functions $A_1, \ldots, A_4$, $B_1, \ldots, B_4$, $C_1, \ldots, C_4$, whose construction is described in Section 4, are used to evaluate $\tilde{P}_\nu^{-\mu}(t)$, $\tilde{Q}_\nu^{-\mu}(t)$ and the appropriate auxiliary functions. Here, we describe the use of the functions A_1 and B_1 to evaluate $\alpha_{\nu,\mu}$ and $\alpha'_{\nu,\mu}$ in the event that (ν, μ, t) is in the set $\mathcal{O}_1$. The other cases are extremely similar.

First, we let $\xi = \frac{1}{\nu}$, $\tau = \frac{\mu-1}{\nu-1}$ and

$$u = \frac{t - t^*_{\nu,\mu}}{\frac{\pi}{2} - t^*_{\nu,\mu}}. \tag{131}$$

That is, we compute the values of ξ, τ and u defined by the mapping (119) given ν, μ and t. Next, we find the smallest positive integer i such that $\xi_i \leq \xi \leq \xi_{i+1}$, where $\xi_1, \ldots, \xi_{11}$ are the nodes of the partition (107), the smallest positive integer j such that $\tau_j \leq \tau \leq \tau_{i+1}$, where $\tau_1, \ldots, \tau_{17}$ are the nodes of the partition (109), and the smallest positive integer k such that $a_k \leq u \leq a_{k+1}$, where $a_1, a_2, \ldots, a_{21}$ are the nodes of the partition (114).

Having discovered that (ξ, τ, u) is in the set $[\xi_i, \xi_{i+1}] \times [\tau_j, \xi_{j+1}] \times [a_k, a_{k+1}]$, we evaluate the compressed trivariate Chebyshev expansions representing A_1 and C_1 on this rectangular prism. We

scale the results by ν to obtain the values of $\alpha_{\nu,\mu}$ and $\alpha'_{\nu,\mu}$. The values of $\tilde{P}_\nu^{-\mu}$ and $\tilde{Q}_\nu^{-\mu}$ are then calculated via (71) and (72).

6. Numerical experiments

In this section, we present the results of numerical experiments which were conducted to assess the performance of the `alegendre_eval` routine. The task of constructing reference values with which to compare our results was quite challenging. All existing packages of which the author is aware were prohibitively slow when evaluating associated Legendre functions with large noninteger parameters, and existing asymptotic expansions are either not viable (e.g., the Liouville-Green expansions (18) and (19) whose coefficients cannot be readily computed) or only applicable in the case of an extremely limited range of parameters (e.g., the trigonometric expansions (13) and (14) which are catastrophically unstable even for relatively small values of μ). As a result, we were quite limited in the extent to which we could verify our approach in the case of large noninteger parameters.

In the case of integer values of the parameters, the well-known three-term recurrence relations can be used to evaluate the associated Legendre function accurately, provided extended precision arithmetic is used to perform the computations. In our reference calculations, we represented real numbers in the form $x \exp(y)$ with x and y quadruple precision (Fortran REAL*16) numbers. This enabled us to test our code quite thoroughly in the case of integer parameters. We note that the time required to evaluate the associated Legendre functions using the recurrence relations grows with the magnitudes of the parameters, making such an approach uncompetitive with the algorithm of this paper in many cases.

These experiments were carried out on a laptop computer equipped with an Intel Core i7-5600U processor running at 2.6 GHz and 16 GB of memory. Our code was compiled with the GNU Fortran compiler version 5.2.1 using the "-Ofast" compiler optimization flag.

6.1. The accuracy with which $\alpha'_{\nu,\mu}$ is evaluated for small noninteger values of ν

In these experiments, we measured the accuracy with which `alegendre_eval` calculates $\alpha'_{\nu,\mu}$ in the oscillatory region. Reference values were calculated using version 11 of Wolfram's Mathematica package. The cost of the reference calculations was prohibitively expensive for large ν, with the consequence that we only considered values of ν between 0 and $1,000$.

In each experiment, we choose 10 pairs (ν, μ) by first picking a random value of ν in a given range, and then choosing a random value of μ in the interval $(0, \nu)$. For each pair chosen in this fashion, we evaluated $\alpha'_{\nu,\mu}$ at 100 equispaced points either in interval

$$\left(t^*_{\nu,\mu}, \frac{\pi}{2}\right) \tag{132}$$

or in the interval

$$\left(\frac{1}{1000}, \frac{\pi}{2}\right), \tag{133}$$

depending on whether $\mu > \frac{1}{2}$ or not. Table 1 reports the results. There, each row corresponds to one experiment and gives the largest relative error observed in $\alpha'_{\nu,\mu}$ as well as the average time taken by the `alegendre_eval` routine.

Range of ν	Maximum relative error in $\alpha'_{\nu,\mu}$	Average evaluation time (in seconds)
0 - 1	2.26×10^{-14}	2.88×10^{-06}
1 - 5	2.62×10^{-15}	1.90×10^{-06}
5 - 10	2.38×10^{-15}	1.62×10^{-06}
10 - 50	4.15×10^{-15}	3.21×10^{-06}
50 - 100	8.53×10^{-15}	2.29×10^{-06}
100 - 500	1.88×10^{-14}	2.35×10^{-06}
500 - 1,000	3.49×10^{-14}	1.44×10^{-06}

Table 1: The results of the experiments of Section 6.1 in which the accuracy with which `alegendre_eval` calculates $\alpha'_{\nu,\mu}$ for small values of ν is tested via comparison with Wolfram's Mathematica package.

6.2. The accuracy with which the logarithms are evaluated in the case of small noninteger values of ν

In these experiments, we measured the accuracy with which `alegendre_eval` calculates the functions

$$\log\left(\tilde{P}_\nu^{-\mu}(t)\right) - \nu \quad \text{and} \quad \log\left(\tilde{Q}_\nu^{-\mu}(t)\right) + \nu \tag{134}$$

in the nonoscillatory region. High accuracy reference values for these experiments were calculated using version 11 of Wolfram's Mathematica package. Again the high cost of the reference calculations led us to only consider values of ν between 0 and $1,000$.

In each experiment, we choose 10 pairs (ν, μ) by first picking a random value of ν in a given range, and then choosing a random value of μ in the interval $\left(\frac{1}{2}, \nu\right)$. For each pair chosen in this fashion, we evaluated the functions (134) at 100 equispaced points in the interval

$$\left(0, t^*_{\nu,\mu}\right). \tag{135}$$

Table 2 reports the results. There, each row corresponds to one experiment and gives the largest relative error observed in each of the functions (134), as well as the average time taken by the `alegendre_eval` routine.

Range of ν	Maximum relative error in $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$	Maximum relative error in $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$	Average evaluation time (in seconds)
0.5 - 1	3.36×10^{-16}	2.58×10^{-15}	1.32×10^{-06}
1 - 5	3.21×10^{-16}	9.28×10^{-16}	1.37×10^{-06}
5 - 10	8.85×10^{-16}	9.14×10^{-15}	1.69×10^{-06}
10 - 50	4.39×10^{-15}	4.43×10^{-15}	3.36×10^{-06}
50 - 100	2.58×10^{-15}	3.49×10^{-15}	1.61×10^{-06}
100 - 500	4.21×10^{-15}	4.47×10^{-15}	2.89×10^{-06}
500 - 1,000	2.54×10^{-15}	3.24×10^{-15}	1.59×10^{-06}

Table 2: The results of the experiments of Section 6.2 in which the accuracy with which `alegendre_eval` calculates $\alpha'_{\nu,\mu}$ for small values of ν is tested via comparison with Wolfram's Mathematica package.

6.3. The accuracy with which $\alpha'_{\nu,\mu}$ is evaluated in the case of large ν and small μ

In these experiments, we measured the accuracy with which `alegendre_eval` calculates $\alpha'_{\nu,\mu}$ in the oscillatory region by comparison with values obtained using the trigonometric expansions (13) and (14). Since these expansions are numerically unstable, and catastrophically so when μ is large, we considered only pairs of the parameters (ν, μ) with μ small in magnitude. Even so, 1,000 digit arithmetic was required in order to obtain accurate reference values for these experiments.

In each experiment, we choose 10 pairs (ν, μ) by first picking a random value of ν in a given range, and then choosing a random value of μ in the interval $\left(\frac{1}{2}, \frac{\nu}{100}\right)$. For each pair, the function $\alpha'_{\nu,\mu}$ was evaluated at 100 equispaced points in the interval

$$\left(\max\left(t^*_{\nu,\mu}\frac{\pi}{6}\right), \frac{\pi}{2}\right). \tag{136}$$

We note that the trigonometric expansions used here only converge in the interval $\left(\frac{\pi}{6}, \frac{5\pi}{6}\right)$, hence the choice of the interval (136).

Range of ν	Maximum relative error in $\alpha'_{\nu,\mu}$	Average evaluation time (in seconds)
1,000 - 5,000	1.95×10^{-15}	1.41×10^{-06}
5,000 - 10,000	1.45×10^{-15}	1.01×10^{-06}
10,000 - 50,000	1.06×10^{-15}	1.06×10^{-06}
50,000 - 100,000	9.70×10^{-16}	7.69×10^{-07}
100,000 - 500,000	8.66×10^{-16}	7.71×10^{-07}
500,000 - 1,000,000	9.73×10^{-16}	5.63×10^{-07}
1,000,000 - 5,000,000	7.30×10^{-16}	5.33×10^{-07}
5,000,000 - 10,000,000	8.05×10^{-16}	4.86×10^{-07}

Table 3: The results of the experiments of Section 6.3 in which the accuracy with which `alegendre_eval` calculates $\alpha'_{\nu,\mu}$ for large ν and small μ is tested via comparison with the trigonometric expansions (13) and (14).

6.4. The accuracy with which $\alpha'_{\nu,\mu}$ is evaluated in the case of integer parameters

In these experiments, the accuracy with which $\alpha'_{\nu,\mu}$ is evaluated in the oscillatory regime was measured by comparison with reference values calculated using the well-known three-term recurrence relations satisfied by the associated Legendre functions. The reference calculations were conducted in extended precision arithmetic in order to ensure accuracy.

The experiments of this section proceeded just as those described in Section 6.1, except only integer values of the parameters were considered. Table 4 displays the results.

6.5. The accuracy with which the logarithms are evaluated in the case of integer parameters

In these experiments, we measured the accuracy with which `alegendre_eval` calculates the functions (134) in the nonoscillatory regime. Reference values were calculated using the well-known three-term recurrence relations satisfied by the associated Legendre functions. The reference calculations were conducted in extended precision arithmetic in order to ensure accuracy.

These experiments proceeded just as those described in Section 6.2, except only integer values of the parameters were considered. Table 5 displays the results.

Range of ν	Maximum relative error in $\alpha'_{\nu,\mu}$	Average evaluation time (in seconds)
10 - 50	2.35×10^{-14}	2.90×10^{-06}
50 - 100	4.71×10^{-15}	1.84×10^{-06}
100 - 500	4.96×10^{-15}	3.17×10^{-06}
500 - 1,000	2.86×10^{-14}	2.13×10^{-06}
1,000 - 5,000	8.62×10^{-15}	1.36×10^{-06}
5,000 - 10,000	5.94×10^{-15}	1.00×10^{-06}
10,000 - 50,000	2.74×10^{-14}	1.22×10^{-06}
50,000 - 100,000	7.36×10^{-14}	9.47×10^{-07}
100,000 - 500,000	1.86×10^{-14}	1.01×10^{-06}
500,000 - 1,000,000	3.09×10^{-14}	8.16×10^{-07}

Table 4: The results of the experiments of Section 6.4 in which the accuracy with which `alegendre_eval` calculates $\alpha'_{\nu,\mu}$ for integers values of the parameters is tested.

Range of ν	Maximum relative error in $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$	Maximum relative error in $\log(\tilde{Q}_\nu^{-\mu}(t)) + \nu$	Average evaluation time (in seconds)
10 - 50	4.21×10^{-15}	4.65×10^{-15}	3.62×10^{-06}
50 - 100	3.42×10^{-15}	3.32×10^{-15}	2.96×10^{-06}
100 - 500	3.07×10^{-15}	4.07×10^{-15}	2.79×10^{-06}
500 - 1,000	2.95×10^{-15}	3.01×10^{-15}	1.71×10^{-06}
1,000 - 5,000	2.63×10^{-15}	4.14×10^{-15}	1.51×10^{-06}
5,000 - 10,000	1.98×10^{-15}	1.83×10^{-15}	1.03×10^{-06}
10,000 - 50,000	1.98×10^{-15}	2.68×10^{-15}	1.68×10^{-06}
50,000 - 100,000	1.63×10^{-15}	2.07×10^{-15}	1.26×10^{-06}
100,000 - 500,000	1.73×10^{-15}	1.63×10^{-15}	1.51×10^{-06}
500,000 - 1,000,000	1.67×10^{-15}	2.23×10^{-15}	1.13×10^{-06}

Table 5: The results of the experiments of Section 6.5 in which the accuracy with which `alegendre_eval` calculates the functions $\log(\tilde{P}_\nu^{-\mu}(t)) - \nu$ and $\log(\tilde{Q}_\nu^{-\mu}(t)) - \nu$ for integers values of the parameters is tested.

6.6. The accuracy with which the associated Legendre functions are evaluated in the case of integer parameters

In these experiments, we measured the accuracy with which `alegendre_eval` calculates the functions $\tilde{P}_\nu^{-\mu}$ and $\tilde{Q}_\nu^{-\mu}$ in the case of integer values of the parameters. Reference values were calculated using the three-term recurrence relations. As usual, extended precision arithmetic was used during the reference calculations in order to ensure their accuracy.

In each experiment, 10 pairs of the parameters (ν, μ) were constructed by first choosing an integer value of ν in a given range at random and then choosing an integer value of μ in the range $(0, \nu)$ at random. For each such pair, we evaluated the function $\tilde{P}_\nu^{-\mu}(t) + i\tilde{Q}_\nu^{-\mu}(t)$ at 100 equispaced points either in interval

$$\left(t^*_{\nu,\mu}, \frac{\pi}{2}\right) \tag{137}$$

or in the interval

$$\left(\frac{1}{1000}, \frac{\pi}{2}\right), \tag{138}$$

depending on whether $\mu > \frac{1}{2}$ or not. Table 6 reports the results. Each row corresponds to one experiment and reports the largest relative error which was observed as well as the average evaluation time. We note that we considered the function $\tilde{P}_\nu^{-\mu}(t) + i\tilde{Q}_\nu^{-\mu}(t)$ because, unlike $\tilde{P}_\nu^{-\mu}(t)$ and $\tilde{Q}_\nu^{-\mu}(t)$, its absolute value is nonoscillatory and does not have roots on the interval $(0, \pi)$.

Range of ν	Maximum relative error in $\tilde{P}_\nu^{-\mu}(t) + i\tilde{Q}_\nu^{-\mu}(t)$	Average evaluation time (in seconds)
10 - 50	2.62×10^{-13}	3.75×10^{-06}
50 - 100	4.20×10^{-13}	2.05×10^{-06}
100 - 500	1.20×10^{-12}	2.29×10^{-06}
500 - 1,000	1.72×10^{-12}	1.60×10^{-06}
1,000 - 5,000	8.57×10^{-12}	2.53×10^{-06}
5,000 - 10,000	1.38×10^{-11}	1.42×10^{-06}
10,000 - 50,000	8.51×10^{-11}	1.21×10^{-06}
50,000 - 100,000	9.07×10^{-11}	9.11×10^{-07}
100,000 - 500,000	9.83×10^{-10}	1.06×10^{-06}
500,000 - 1,000,000	8.25×10^{-10}	8.20×10^{-07}

Table 6: The results of the experiments of Section 6.6 in which the accuracy with which `alegendre_eval` evaluates the function $\tilde{P}_\nu^{-\mu}(t) + i\tilde{Q}_\nu^{-\mu}(t)$ for integers values of the parameters is tested.

From Table 6, we see that the relative errors in the calculated values of the associated Legendre functions increase as a function of the parameter ν. This is expected, and consistent with the condition number of the evaluation of the function $\tilde{P}_\nu^{-\mu}(t) + i\tilde{Q}_\nu^{-\mu}(t)$.

7. Conclusions

Nonoscillatory phase functions provide powerful theoretical tools for analyzing the solutions of second order differential equations as well as a framework for the design of simple and efficient numerical algorithms. Here, we have designed a scheme for the numerical evaluation of the associated Legendre functions on the cut using this framework. Our approach is simple-minded and highly effective. Moreover, by making use of the algorithms of [3] and [5], it can be applied in the case of many other special functions satisfying second order differential equations, such as the prolate spheroidal wave functions and the generalized Laguerre functions. The author will report on the use of the techniques of this paper to evaluate other special functions of interest at a later date.

In the nonoscillatory region, our algorithm calculates the logarithms of the associated Legendre functions as well as their values. This is useful in cases in which the magnitudes of those functions are too large or too small to be encoded using the IEEE double precision format. In the oscillatory region, in addition to the values of the associated Legendre functions, our algorithm also returns the values of a nonoscillatory phase function for the associated Legendre differential equation and its derivative. This is extremely helpful when computing the zeros of special functions, and when

applying special function transforms via the butterfly algorithm. The author will report on the use nonoscillatory phase functions to rapidly compute the roots of the associated Legendre functions and to rapidly apply the spherical harmonic transform at a later date.

8. Acknowledgments

The author is grateful to Zydrunas Gimbutas of NIST Boulder for providing his arbitrary precision arithmetic codes for evaluating the associated Legendre functions of large integer degrees and orders via the three-term recurrence relations they satisfy. This work was supported in part by a UC Davis Chancellor's Fellowship.

9. References

References

[1] Bochner, S., and Martin, W. *Several Complex Variables.* Princeton University Press, 1948.

[2] Boyd, W. G. C., and Dunster, T. M. Uniform asymptotic solutions of a class of second-order linear differential equations having a turning point and a regular singularity, with an application to Legendre functions. *SIAM Journal on Mathematical Analysis 17* (1986), 422–450.

[3] Bremer, J. On the numerical solution of second order differential equations in the high-frequency regime. *Applied and Computational Harmonic Analysis*, to appear.

[4] Bremer, J. An algorithm for the numerical evaluation of Bessel functions of real orders and arguments. *arxiv.org:1705.07820* (2017).

[5] Bremer, J., and Rokhlin, V. Improved estimates for nonoscillatory phase functions. *Discrete and Continuous Dynamical Systems, Series A 36* (2016), 4101–4131.

[6] *NIST Digital Library of Mathematical Functions.* http://dlmf.nist.gov/, Release 1.0.13 of 2016-09-16. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller and B. V. Saunders, eds.

[7] Erdélyi, A., et al. *Higher Transcendental Functions*, vol. I. McGraw-Hill, 1953.

[8] Fabbian, G., and Stompor, R. High precision simulations of weak lensing effect on cosmic microwave background polarization. *Astronomy and Astrophysics 556* (2013).

[9] Fedoryuk, M. V. *Asymptotic Analysis.* Springer-Verlag, 1993.

[10] Fefferman, C. On the convergence of multiple Fourier series. *Bulletin of the American Mathematical Society 77* (1971), 744–745.

[11] Fukushima, T. Numerical computation of spherical harmonics of arbitrary degree and order by extending exponent of floating point numbers. *Journal of Geodesy 86* (2012), 271–285.

[12] Gradshteyn, I. S., and Ryzhik, I. M. *Table of Integrals, Series and Products*, Eigth ed. Elsevier, 2015.

[13] Heitman, Z., Bremer, J., Rokhlin, V., and Vioreanu, B. On the asymptotics of Bessel functions in the Fresnel regime. *Applied and Computational Harmonic Analysis 39* (2015), 347–355.

[14] Higham, N. J. *Accuracy and Stability of Numerical Algorithms*, second ed. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2002.

[15] Hille, E. *Ordinary differential equations in the complex domain.* Wiley, New York, 1976.

[16] Kummer, E. De generali quadam aequatione differentiali tertti ordinis. *Progr. Evang. Köngil. Stadtgymnasium Liegnitz* (1834).

[17] Li, Y., and Yang, H. Interpolative butterfly factorization. *SIAM Journal on Scientific Computing 39* (2017), A503–A531.

[18] Li, Y., Yang, H., Martin, E., Ho, K. L., and Ying, L. Butterfly factorization. *SIAM Journal on Multiscale Modeling and Simulation 13* (2015), 714–732.

[19] Macdonald, H. Formulae for the spherical harmonic $P_n^{-m}(\mu)$, when $1-\mu$ is a small quantity. *Proceedings of the Royal Society of London* (1914), 220–221.

[20] Mason, J., and Handscomb, D. *Chebyshev Polynomials.* Chapman and Hall, 2003.

[21] Miller, J. On the choice of standard solutions for a homogeneous linear differential equation of the second order. *Quarterly Journal of Mechanics and Applied Mathematics 3* (1950), 225–235.

[22] Olver, F., and Smith, J. Associated legendre functions on the cut. *Journal of Computational Physics 51* (1983), 502–518.

[23] Olver, F. W. *Asymptotics and Special Functions.* A.K. Peters, Natick, MA, 1997.

[24] Reinecke, M., and Seljebotn, D. Libsharp - spherical harmonic transforms revisited. *Astronomy and Astrophysics 554* (2013).

[25] Trefethen, N. *Approximation Theory and Approximation Practice.* Society for Industrial and Applied Mathematics, 2013.

[26] Tygert, M. Fast algorithms for spherical harmonic expansions, III. *Journal of Computational Physics 229* (2010), 6181–6192.

[27] Watson, G. N. *A Treatise on the Theory of Bessel Functions*, second ed. Cambridge University Press, New York, 1995.