

Efficient and Accurate Rotation of Finite Spherical Harmonics Expansions

C. Lessig, T. de Witt, E. Fiume

Department of Computer Science, University of Toronto, Toronto, ON, Canada

Abstract

Spherical harmonics are employed in a wide range of applications in computational science and physics, and many of them require the rotation of functions. We present an efficient and accurate algorithm for the rotation of finite spherical harmonics expansions. Exploiting the pointwise action of the rotation group on functions on the sphere, we obtain the spherical harmonics expansion of a rotated signal from function values at rotated sampling points. The number of sampling points and their location permits one to balance performance and accuracy, making our technique well-suited for a wide range of applications. Numerical experiments comparing different sampling schemes and various techniques from the literature are presented, making this the first thorough evaluation of spherical harmonics rotation algorithms.

Keywords: Spherical Harmonics, Reproducing kernel, Rotation

1. Introduction

Spherical harmonics are employed in a wide range of applications in computational science. In many areas, such as quantum chemistry, astronomy, geoscience, scattering theory, electromagnetics, and computer graphics, the rotation of the functions is also required: given the expansion

$$f(\omega) = \sum_{l=0}^L \sum_{m=-l}^l \alpha_{lm} y_{lm}(\omega),$$

Email address: lessig@dgp.toronto.edu, tyler@dgp.toronto.edu, elf@dgp.toronto.edu (C. Lessig, T. de Witt, E. Fiume)

of a function $f : \mathcal{S}^2 \rightarrow \mathbb{R}$ (or \mathbb{C}) in the Hilbert space $\mathcal{H}_{\leq L}(\mathcal{S}^2)$ spanned by all spherical harmonics y_{lm} up to band L , one seeks the spherical harmonics basis function coefficients $\bar{\alpha} = (\bar{\alpha}_{0,0}, \dots, \bar{\alpha}_{l,m}, \dots, \bar{\alpha}_{l,l})^T$ of the rotated function $\bar{f} = \mathbf{R}f$ when the rotation group $\text{SO}(3)$ acts pointwise on functions as $\bar{f}(\omega) = f(\mathbf{R}^{-1}\omega)$ for arbitrary $\mathbf{R} \in \text{SO}(3)$.

The rotated basis function coefficients $\bar{\alpha}_l$ in the l^{th} spherical harmonics band \mathcal{H}_l are a linear combination of the unrotated coefficients α_l in the band and can hence be obtained with a matrix \mathbf{R}_l as $\bar{\alpha}_l = \mathbf{R}_l \alpha_l$. Direct approaches [1, 2] for computing spherical harmonics rotation matrices \mathbf{R}_l were already proposed in the 19th century but these are computationally expensive and numerically unstable [3]. Recurrence schemes [4, 5, 6, 3, 7] alleviate some of these problems but they remain slow and become numerically unstable for the large number of bands that are increasingly required in applications [8, 9]. In contrast to the literature, we do not compute spherical harmonics rotation matrices but employ a sampling formula for the l^{th} spherical harmonics band \mathcal{H}_l and exploit that the action of the rotation group $\text{SO}(3)$ is defined point-wise. A rotated function $\bar{f}_l = \mathbf{R}f_l$ is then obtained by evaluating the unrotated function f_l at rotated sampling locations $\bar{\lambda}_i = \mathbf{R}^{-1}\lambda_i$ and using a change of basis from the sampling basis to spherical harmonics. Extensive numerical experiments are presented, and to our knowledge this is the first thorough empirical validation of spherical harmonics rotation algorithms. The experiments demonstrate that our technique provides accuracy comparable to the best known methods. The fastest technique currently available is a variation of the algorithm by Pinchon and Hoggan [10], provided to us by these authors, which appears here for the first time. Our technique provides the additional advantage that it is easy to implement, overcoming a practical limitation of other algorithms.

The remainder of the paper is structured as follows. In Sec. 2 we develop our algorithm for rotating finite spherical harmonics expansions and discuss the trade-offs that are available by the choice of sampling locations. In Sec. 3 experimental results are presented, and in Sec. 4 we discuss the relevance of our algorithm and possible directions for future work. Our paper presents the rotation of real spherical harmonics but all results hold with the usual modifications in the complex setting.

2. Rotation of Finite Spherical Harmonics Expansions

In Sec. 2.1 we will derive our technique in the continuous setting, while its computer implementation is discussed in Sec. 2.2 when we formulate it using linear algebra. The choice of sampling points, which is important for the performance and accuracy of our technique, is discussed in Sec. 2.3.

2.1. A Rotation Algorithm for Spherical Harmonics: Continuous Theory

The space $(\mathcal{H}_{\leq L}, \langle \cdot, \cdot \rangle)$ with the standard L_2 inner product $\langle \cdot, \cdot \rangle$ over the sphere is spanned by (Legendre) spherical harmonics

$$y_{lm}(\omega) = y_{lm}(\theta, \phi) = \eta_{lm} P_{lm}(\cos \theta) \begin{cases} \sin(|m| \phi) & m < 0 \\ 1 & m = 0 \\ \cos(m \phi) & m > 0 \end{cases}, \quad (1)$$

where $\omega = (\theta, \phi)$ is a point on the sphere, η_{lm} is a normalization constant such that the y_{lm} form an orthonormal basis, P_{lm} is the associated Legendre polynomial of degree l and order m , and the indices satisfy $0 \leq l \leq L$ and $-l \leq m \leq l$. $\mathcal{H}_{\leq L}$ admits the orthogonal decomposition $\mathcal{H}_{\leq L} = \mathcal{H}_0 \oplus \dots \oplus \mathcal{H}_L$ and the bands \mathcal{H}_l are closed under the action of the rotation group $\text{SO}(3)$. For the rotation of a function $f \in \mathcal{H}_{\leq L}$ it is hence sufficient to consider the restrictions $f_l \in \mathcal{H}_l$ with $f = f_0 + \dots + f_L$. See for example the book by Freedman et al. [11] for a more detailed discussion of spherical harmonics.

Our algorithm for rotating finite spherical harmonic expansions employs a biorthogonal reproducing kernel basis for the space \mathcal{H}_l . A reproducing kernel $k_{\bar{\omega}}(\omega)$ is a function such that $\langle k_{\bar{\omega}}(\omega), f(\omega) \rangle = f(\bar{\omega})$ for arbitrary $\bar{\omega} \in \mathcal{S}^2$ and it follows from the general theory [12, 13] that for \mathcal{H}_l it is given by

$$k_{\bar{\omega}}(\omega) = \sum_{-l \leq m \leq l} y_{lm}(\omega) y_{lm}(\bar{\omega}) = \frac{2l+1}{4\pi} P_l(\bar{\omega} \cdot \omega),$$

where the last equality holds by the spherical harmonics addition theorem. With a sequence $\boldsymbol{\lambda} = \{\lambda_i\}$ of $2l+1$ points $\lambda_i \in \mathcal{S}^2$ such that the reproducing kernel functions $k_i(\omega) \equiv k_{\lambda_i}(\omega)$ at the λ_i are linearly independent, the functions $\{k_i(\omega)\}$ form a basis for \mathcal{H}_l and the associated dual basis $\{\tilde{k}_j(\omega)\}$ is uniquely defined by $\langle k_i(\omega), \tilde{k}_j(\omega) \rangle = \delta_{ij}$. Any function $f_l \in \mathcal{H}_l$ can then be written as

$$f_l(\bar{\omega}) = \sum_{i=1}^{2l+1} \langle f_l(\omega), k_i(\omega) \rangle \tilde{k}_i(\bar{\omega}) = \sum_{i=1}^{2l+1} f_l(\lambda_i) \tilde{k}_i(\bar{\omega}) \quad (2)$$

Algorithm 1: Rotation for general sampling points.

Input: $\mathbf{R}, \boldsymbol{\alpha}_l, \mathbf{C}_l, \boldsymbol{\lambda}_l$

Output: $\bar{\boldsymbol{\alpha}}_l$

- 1 Rotate sampling points, $\bar{\lambda}_i = \mathbf{R}^{-1}\lambda_i$.
 - 2 Evaluate $f_l \in \mathcal{H}_l$ at rotated sampling points, $\bar{\mathbf{f}}_l = \{\sum_{m=-l}^l \alpha_{lm} y_{lm}(\bar{\lambda}_i)\}$.
 - 3 Compute basis function coefficients of the rotated function, $\bar{\boldsymbol{\alpha}}_l = \mathbf{S}_l \bar{\mathbf{f}}_l$.
-

where the last equality follows from the reproducing property of the kernel functions. The basis function coefficients with respect to the kernel basis are thus given by the function values $f(\lambda_i)$ at the sampling locations. Eq. 2 is the analogue of the Shannon sampling theorem for the sphere and the lack of a regular sequence on \mathcal{S}^2 is a principal reason for the biorthogonality of the basis functions; in contrast, over the real line the set of sinc functions centered at the integers form an orthonormal basis for the space of bandlimited functions over the real line.

By Eq. 2, the representation of a rotated function \bar{f}_l in the kernel basis is

$$\bar{f}_l = \mathbf{R} f_l = \sum_{i=1}^{2l+1} \mathbf{R} f_l(\lambda_i) \tilde{k}_i = \sum_{i=1}^{2l+1} f_l(\mathbf{R}^{-1}\lambda_i) \tilde{k}_i = \sum_{i=1}^{2l+1} f_l(\bar{\lambda}_i) \tilde{k}_i \quad (3)$$

where we employed the point-wise definition of the rotation action. Since both the kernel basis and spherical harmonics with fixed l span \mathcal{H}_l , the spherical harmonics coefficients $\bar{\alpha}_{lm}$ of the rotated signal can be obtained by a change of basis

$$\bar{\alpha}_{lm} = \langle \bar{f}_l, y_{lm} \rangle = \left\langle \sum_{i=1}^{2l+1} f_l(\bar{\lambda}_i) \tilde{k}_i, y_{lm} \right\rangle. \quad (4)$$

2.2. A Rotation Algorithm for Spherical Harmonics: Discrete Formulation

With the isomorphism between \mathcal{H}_l and Euclidean space \mathcal{R}^{2l+1} provided by spherical harmonics y_{lm} , a representation of the reproducing kernel basis functions is given by their basis function coefficients. These can be arranged in matrix form to yield

$$\mathbf{K}_l \equiv \mathbf{K}_l(\boldsymbol{\lambda}) = \begin{bmatrix} y_{l,-l}(\lambda_1) & \dots & y_{l,l}(\lambda_1) \\ \vdots & \ddots & \vdots \\ y_{l,-l}(\lambda_n) & \dots & y_{l,l}(\lambda_n) \end{bmatrix} \in \mathcal{R}^{n \times n} \quad (5)$$

Algorithm 2: Rotation for equi-latitude sampling points at ζ .

Input: $\mathbf{R}, \boldsymbol{\alpha}_l, \zeta$

Output: $\bar{\boldsymbol{\alpha}}_l$

- 1 Construct and rotate sample points,
 $\bar{\lambda}_i = \mathbf{R}^{-1} \lambda_i$ with $\lambda_i = (\zeta, \phi_i)$ and $\phi_i = (2\pi i)/(2l + 1)$.
 - 2 Evaluate $f_l \in \mathcal{H}_l$ at rotated sample points, $\bar{\mathbf{f}}_l = \{\sum_{m=-l}^l \alpha_{lm} y_{lm}(\bar{\lambda}_i)\}$.
 - 3 Compute basis function coefficients of the rotated function,
 $\alpha_{lm} = \text{DFT}\{\bar{\mathbf{f}}_l\}_m / (\eta_{lm} P_{lm}(\cos \zeta))$.¹
-

where $n = 2l + 1$. \mathbf{K}_l provides the change of basis from spherical harmonics to the kernel basis so that $f_l(\boldsymbol{\lambda}) = \mathbf{K}_l \boldsymbol{\alpha}_l$, where $f_l(\boldsymbol{\lambda}) = (f(\lambda_1), \dots, f(\lambda_n))^T$ and $\boldsymbol{\alpha}_l = (\alpha_{l,-l}, \dots, \alpha_{l,m}, \dots, \alpha_{l,l})^T$. Conversely, $\boldsymbol{\alpha}_l = \mathbf{K}_l^{-1} f_l(\boldsymbol{\lambda}) = \mathbf{S}_l f_l(\boldsymbol{\lambda})$. The desired spherical harmonics coefficients $\bar{\boldsymbol{\alpha}}_l$ of the rotated function are thus obtained as

$$\bar{\boldsymbol{\alpha}}_l = \mathbf{S}_l f_l(\bar{\boldsymbol{\lambda}}) \quad (6)$$

where $f_l(\bar{\boldsymbol{\lambda}}) = (f_l(\bar{\lambda}_1), \dots, f_l(\bar{\lambda}_n))^T$ are the function values at the rotated sampling locations $\bar{\lambda}_i = \mathbf{R}^{-1} \lambda_i$; it is easily verified that Eq. 6 is equivalent to Eq. 4. Our rotation algorithm for finite spherical harmonics expansions is summarized in Algo. 1. Note that the matrix \mathbf{S}_l depends only on the sampling locations but not on the rotation. It can hence be precomputed.

An alternative interpretation of our algorithm is to consider it as a factorization of the classical spherical harmonics rotation matrix \mathbf{R}_l , similar in spirit to the work by Pinchon and Hoggan [10]. The rotated signal at the unrotated sampling locations is given by $\bar{f}_l(\boldsymbol{\lambda}) = \mathbf{K}_l(\boldsymbol{\lambda}) \bar{\boldsymbol{\alpha}}_l$ and by definition $\bar{f}_l(\boldsymbol{\lambda}) = f_l(\mathbf{R}^{-1} \boldsymbol{\lambda})$. Additionally, the unrotated signal at the rotated sampling locations is $f_l(\mathbf{R}^{-1} \boldsymbol{\lambda}) = \mathbf{K}_l(\mathbf{R}^{-1} \boldsymbol{\lambda}) \boldsymbol{\alpha}_l$. One therefore has

$$\begin{aligned} \mathbf{K}_l(\boldsymbol{\lambda}) \bar{\boldsymbol{\alpha}} &= \bar{f}_l(\boldsymbol{\lambda}) \\ \mathbf{K}_l(\boldsymbol{\lambda}) \bar{\boldsymbol{\alpha}} &= \mathbf{K}_l(\mathbf{R}^{-1} \boldsymbol{\lambda}) \boldsymbol{\alpha} \\ \bar{\boldsymbol{\alpha}} &= \mathbf{K}_l^{-1}(\boldsymbol{\lambda}) \mathbf{K}_l(\mathbf{R}^{-1} \boldsymbol{\lambda}) \boldsymbol{\alpha} \end{aligned}$$

and the kernel matrix \mathbf{K}_l factors the classical spherical harmonics rotation matrix \mathbf{R}_l as $\mathbf{R}_l = \mathbf{K}_l^{-1}(\boldsymbol{\lambda}) \mathbf{K}_l(\mathbf{R}^{-1} \boldsymbol{\lambda})$, and analogously when $m > n$.

¹DFT $\{\bar{\mathbf{f}}_l\}_m$ refers to the m^{th} coefficients of the Discrete Fourier Transform of $\bar{\mathbf{f}}_l$ with $-l \leq m \leq l$ where negative m correspond to $\sin(|m|x)$ and positive m to $\cos(mx)$.

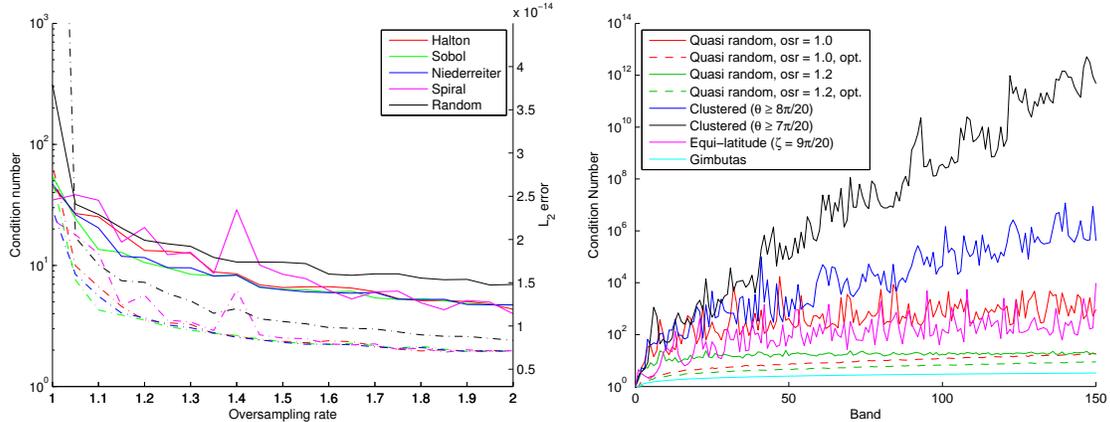


Figure 1: Left: Condition number (full) and average L_2 rotation error (dashed) as a function of the oversampling rate for $l = 20$. Right: Condition numbers for different strategies to obtain sampling sequences with different oversampling rates (osr).

2.3. Sampling Locations

Previously we posited the existence of sampling sequences λ such that the sampling points λ_i yield linearly independent kernel functions $k_i(\omega) = k_{\lambda_i}(\omega)$ forming a basis for \mathcal{H}_l . The existence of sampling sequences is guaranteed by a theorem due to Müller [14, p. 13][11, p. 51] and it can be shown that up to a set of measure zero any set of $2l + 1$ points on the sphere can be employed. However, in applications not only the existence of sampling sequences but also their quality is of critical importance, and their choice enables balancing accuracy and performance.

By Eq. 6 and $\mathbf{S}_l = \mathbf{K}_l^{-1}$, with our technique the rotation of finite spherical harmonics expansions is equivalent to the solution of a linear system. The condition number $\text{cond}(\mathbf{K}_l)$ hence provides a quality measure for sampling locations. A condition number of unity corresponds to an optimal sampling sequence with an orthogonal kernel basis while a set of locations for which the kernel functions are not linearly independent has an infinite condition number. As one would expect, well-distributed points on the sphere yield low condition numbers and are well-suited for our rotation algorithm. We obtain such locations by mapping quasi-random sequences from the unit square to the upper hemisphere [15]; the restriction to the hemisphere is advantageous since antipodal points would yield co-linear basis functions. Additionally, we also employ the spiral points that were proposed by Saff and Kuijlaars [16]

as well-distributed sequences on \mathcal{S}^2 . Fig. 1 shows that the performance of different well-distributed sequences is qualitatively equivalent and that they outperform sampling points obtained with a (pseudo) random number generator. The graphs also verify the close correlation between condition number and rotation error, a connection that is only violated when a rotated sampling location is close to a pole and the accurate evaluation of spherical harmonics is difficult.

We investigated two strategies to improve the condition numbers obtained with well-distributed point sets: increasing the number of sampling points and optimizing their location. Previously we assumed that $n = 2l + 1$ sampling points are employed. When $m > n$ points are used one obtains an overcomplete basis (or a frame) for which the dual basis functions are defined by a left pseudo-inverse of the kernel matrix.² It is known from the signal processing literature that oversampling often improves accuracy and robustness [17, 18], and in Appendix C we show that the condition number approaches unity as the number of sampling points goes to infinity. Fig. 1 demonstrates empirically that overcomplete representations improve the condition number and the rotation error, and that a small oversampling rate is sufficient to obtain close to optimal accuracy. Oversampling requires more function evaluations than critical sampling, making it computationally more expensive. To improve the accuracy of our rotation algorithm without increasing the computational costs at runtime, we employed numerical optimization of the sampling locations to improve the condition number of the kernel matrix. We were not able to derive an analytic expression for the gradient of $\text{cond}(\mathbf{K}_l)$ but a BFGS optimizer [19] with finite differences and restarts based on different, well-distributed point sets performed well for our purposes. Fig. 1 shows the considerable improvements in the condition number that can be obtained by optimization. The main limitation of optimization are the high computational costs, although these occur only once and during precomputation.

An interesting alternative to the sequences obtained with the above constructions is provided by sampling sequences $\lambda_i = (\zeta, \phi_i)$ with a fixed latitude ζ . By Eq. 1, the kernel matrix can then be factored as $\mathbf{K}_l(\zeta) = \mathbf{P}_l(\zeta) \mathbf{F}$ where \mathbf{F} is the discrete Fourier transform matrix and the nonzero

²For our experiments we employed the Moore-Penrose pseudo-inverse that yields the minimal norm dual.

elements of the diagonal matrix $\mathbf{P}_l(\zeta)$ are given by $p_{mm} = P_{lm}(\cos \zeta)$ with $-l \leq m \leq l$. The basis function coefficients $\bar{\alpha}_{lm}$ of the rotated function can then be recovered with Algo. 2. From the orthogonality of the discrete Fourier transform it follows that $\text{cond}(\mathbf{K}_l(\zeta))$ is given by $\text{cond}(\mathbf{P}_l(\zeta)) = \max(|P_{lm}(\cos \zeta)|) / \min(|P_{lm}(\cos \zeta)|)$ and the accuracy of the algorithm is determined by ζ . Experiments show that beyond a critical latitude the condition number deteriorates as the distance from the equator increases and for $l \leq 150$ latitudes with $75^\circ \leq \zeta < 90^\circ$ should be employed. When other latitudes are used or for very large l the accuracy is insufficient only for a very small number of rotated coefficients $\bar{\alpha}_{lm}$ where $P_{lm}(\zeta)$ is very small. A practical work-around is hence to compute these coefficients with Algo. 1 using a small number of additional sampling points with $\theta \neq \zeta$. An interesting choice for the latitude is $\zeta = \pi/2$ which locates the sampling sequence on the equator and leads to the recent algorithm by Gimbutas and Greengard [9]; see Appendix A for details.

3. Experimental Evaluation

We evaluated our algorithm for the rotation of finite spherical harmonics expansions with different choices for the sampling points and compared its accuracy and performance to various techniques in the literature. As sampling sequences we employed optimized point sets with different oversampling rates, nested sampling points where the sampling locations for band l are a subset of those for band $l + 1$, and equi-latitude points with $\zeta = 9\pi/20$ for which the fast Fourier transform was used to speed up computations. With nested sampling points the recurrent structure of spherical harmonics evaluation can be exploited, enabling faster evaluation at the sampling locations. Even with nesting, in particular for low bands, a large fraction of the computation time is spent on spherical harmonics evaluation. We therefore determined the sample values $f_i(\bar{\lambda}_i)$ also from a representation of f_l in the reproducing kernel basis, with primary and dual basis functions interchanged, which is slightly more efficient than spherical harmonics when the spherical harmonics addition theorem is exploited and only Legendre polynomials have to be evaluated. Our technique has a computational complexity of $O(sL^3)$ for nested sampling points and $O(sL^4)$ for non-nested points, where s is the oversampling rate. A detailed analysis of the instruction count for our technique is available in Appendix B.

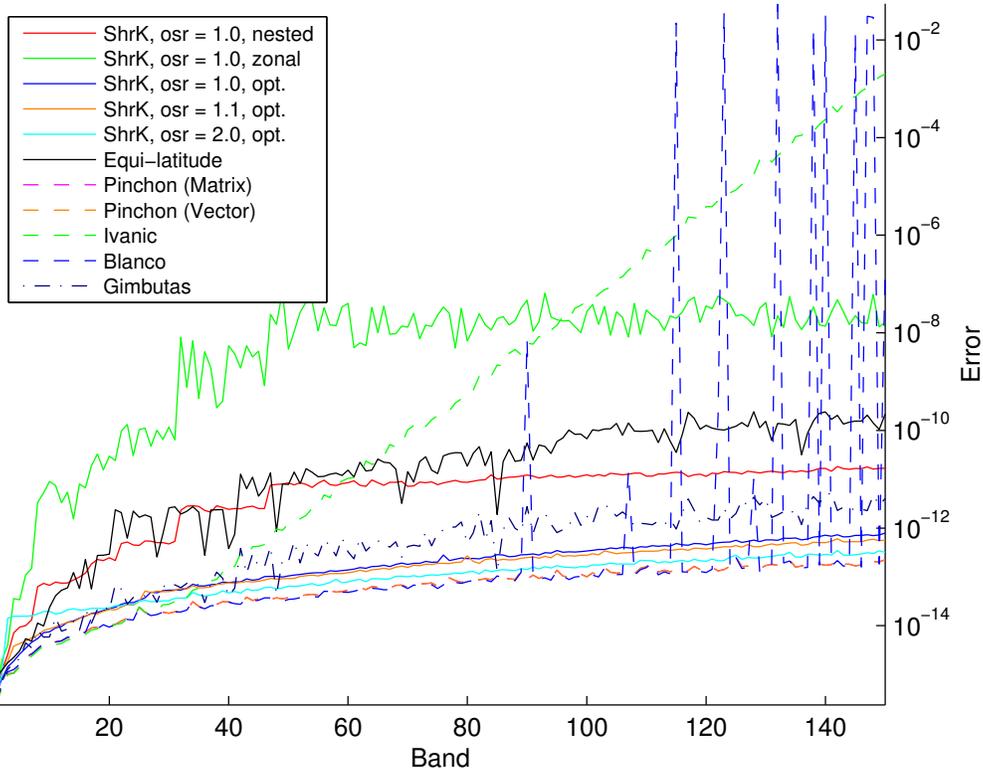


Figure 2: Average L_2 error of experiments (osr = oversampling rate).

From the literature we employed the algorithms by Ivanic and Ruedenberg [3], Blanco et al. [7], Pinchon and Hoggan [10], and Gimbutas and Greengard [9], implemented in C/C++ using double precision and, when available, based on implementations provided by the original authors. For the algorithm proposed by Pinchon and Hoggan [10] we employed two variants: the first obtains the spherical harmonics rotation matrix and then applies it to a coefficient vector as proposed in the original publication (‘Pinchon Matrix’), while the second is a variation suggested to us by the authors that avoids the explicit computation of the rotation matrix and applies the coefficient vector incrementally to the highly sparse factorization employed in the work (‘Pinchon Vector’), yielding a computational complexity of only $O(L^2)$ instead of $O(L^3)$ for the original technique.

Average L_2 errors per band and execution times are reported in Fig. 2 and Fig. 3. Other error norms are qualitatively equivalent to the presented results

and have been omitted. All graphs have been obtained by averaging over a large number of random rotations and reference solutions were obtained as in previous work [9].

The experimental results in Fig. 3 show that nested sampling points (ShrK, $\text{osr} = 1.0$, nested) are more efficient than optimized sampling sequences without this structure, as expected from the lower computational complexity. For lower bands, a representation of f_l in the kernel basis (ShrK, $\text{osr} = 1.0$, zonal) provides an additional performance advantage. For higher bands, the fast Fourier transform makes equi-latitude points almost as efficient as nested sampling points, and we expect even better performance when the nested structure of the equi-latitude points is also exploited. In all cases, Pinchon Vector is the fastest algorithm, due to the lower computational complexity. As shown in Fig. 2, optimized sampling locations (ShrK, $\text{osr} = *$, opt.) provide significantly higher accuracy than well-distributed sequences (ShrK, $\text{osr} = 1.0$, nested) and equi-latitude points, in particular when combined with oversampling. In this case only little accuracy is lost and our algorithm performs comparably to the technique by Pinchon and Hoggan [10] which is currently the most accurate one in the literature.

4. Discussion

With our algorithm the spherical harmonics coefficients of a rotated signal are obtained using a sampling formula for the sphere and by exploiting the point-wise definition of the action of $\text{SO}(3)$ on functions. In contrast to most techniques in the literature that construct spherical harmonics rotation matrices, our algorithm is simple to implement and enables one to trade-off accuracy and performance, making it well-suited for a wide range of applications. We presented an empirical validation of our technique and to our knowledge this is the first extensive comparison of spherical harmonics rotation algorithms available in the literature. The experimental results demonstrate that our technique attains accuracy comparable to the best methods. The fastest technique currently available is those by Pinchon and Hoggan [10] when the rotation matrix is not constructed explicitly, as was suggested and provided to us by these authors, and it has been evaluated for the first time in this paper. A possible disadvantage of our technique is that the sampling matrices \mathbf{S}_l have to be precomputed and stored. However, the technique by Pinchon and Hoggan [10] suffers from similar drawbacks and for our technique it can be avoided with equi-latitude sampling locations.

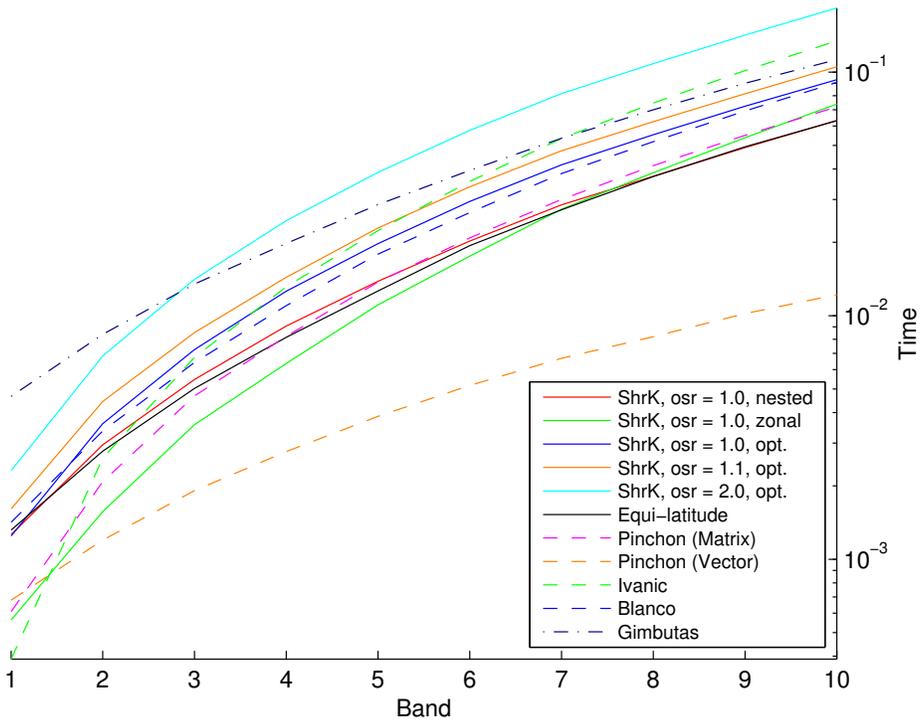
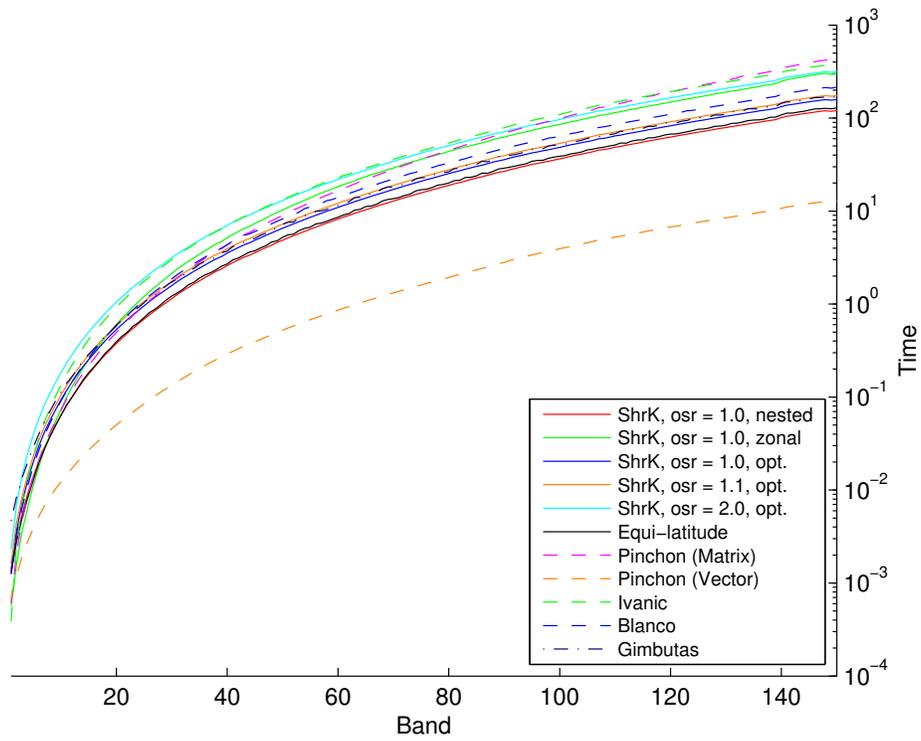


Figure 3: Execution times of experiments (osr = oversampling rate).

Our algorithm was inspired by work by Higgins and Kempinski [20, 21] and Freedman and co-workers [11] who proposed sampling theorems for the sphere similar to Eq. 2. However, these authors did not consider biorthogonal kernel bases that are an important ingredient to our technique. Algorithms similar to ours were proposed previously by Stern [22], see also the paper by James [5], and Gimbutas and Greengard [9]. Unfortunately, Stern’s work received only very little attention after its inception and he also did not investigate different sampling locations or the mathematics underlying the technique. Gimbutas and Greengard [9] recently proposed an interesting variation of our technique with sampling points on the equator, cf. Appendix A. Our work provides a general framework for their algorithm and clarifies its theoretical underpinning.

An interesting question for future work is the existence of optimal sampling sequences although the connection to other point distribution problems on the sphere [16, 23] makes us believe that the problem is very hard. Additionally, orthogonal kernel bases, which are particularly desirable from a computational point of view, cannot exist [24] due to the non-existence of tight spherical designs [25, 26]. A better theoretical understanding of suitable latitudes for equi-latitude sampling points, in particular for very large bands, is also desirable.

A Matlab implementation of our algorithm as well as the C++ framework used for the experiments, including optimized sampling sequences, are available at <http://www.dgp.toronto.edu/people/lessig/shrk/>. Parts of the Matlab code are based on Frederik Simons repository [27].

Acknowledgement

We thank Leslie Greengard, Bos Kempinski, Joe Ivanic, and in particular Frederik Simons and Didier Pinchon for helpful discussions and providing source code. The comments of the anonymous reviewers also helped to substantially improve the paper. Jos Stam and Mathieu Desbrun helped with comments on drafts of the manuscript.

Appendix A. Spectral Spherical Harmonics Rotation

In this appendix we will briefly outline the connection between the recent work by Gimbutas and Greengard [9], which they referred to as spectral method, and our algorithm. Assume equi-latitude sampling points on the

equator are employed so that $\zeta = \pi/2$. At first sight, this choice seems unfortunate since one-half of the $P_{lm}(0)$ vanish and the corresponding basis function coefficients cannot be recovered. However, the remaining coefficients $\alpha_{l\tilde{m}}$ with $\tilde{m} = \{-l \leq m \leq l \mid P_{lm}(0) \neq 0\}$ can be computed to very high accuracy. The condition number for this sub-problem is less than ten even for band 1000, cf. Fig. 1. By exploiting that the derivative $\partial P_{lm}(t)/\partial t|_{t=0}$ of the associated Legendre polynomials on the equator is non-zero exactly for those $\hat{m} = \{m\} \setminus \{\tilde{m}\}$ where the function values vanish, we can determine the remaining $\alpha_{l\hat{m}}$ by derivative sampling, an approach which is in fact well known in the sampling literature [28]. The closure of the derivative under rotation follows from the equivalence with one of the vector Spherical Harmonics basis functions [11, Lemma 12.7.2].

Appendix B. Instruction Count

The instruction count for our technique for nested sampling points for an oversampling rate s and a maximum spherical harmonics band l are

mult / div	$12.66 s L^3 + 79.5 s L^2 + 62.83 s L + 9 s$
sqrt	$0.66 s L^3 + 2.5 s L^2 + 1.83 s L$
sin / cos	$8 s L + 4 s$

and for non-nested sampling points one has

mult / div	$4.25 s L^4 + 3.1 L^3 + 49.5 s L^2 + 41.54 s L + 9 s$
sqrt	$0.25 s L^4 + 1.66 s L^3 + 1.55 s L^2 + 0.23 s L$
sin / cos	$4 s L^2 + 8 s L$

The complete analysis is available online at http://www.dgp.toronto.edu/people/lessig/shrk/data/shrk_instruction_count.pdf.

Appendix C. Proofs

Lemma 1. *The condition number of the kernel matrix converges to unity as the oversampling rate approaches infinity.*

Proof. The kernel matrix has unit condition number if and only if its columns are orthonormal. The scalar product of two columns is an unnormalized Monte Carlo estimator for the inner product of the corresponding Spherical Harmonics basis functions. Since these are orthonormal and the condition number is invariant under scalar scaling this immediately yields the desired result. \square

References

- [1] A. Schmidt, Formeln zur Transformation der Kugelfunktionen, *Z. Math. Phys.* 44 (1899) 327.
- [2] E. P. Wigner, *Gruppentheorie und ihre Anwendungen auf die Quantenmechanik der Atomspektren*, Vieweg Verlag, Braunschweig, 1931.
- [3] J. Ivanic, K. Ruedenberg, Rotation Matrices for Real Spherical Harmonics. Direct Determination by Recursion, *The Journal of Physical Chemistry* 100 (15) (1996) 6342–6347. doi:10.1021/jp953350u.
URL <http://pubs.acs.org/doi/abs/10.1021/jp953350u>
- [4] A. R. Edmonds, *Angular Momentum in Quantum Mechanics*, Princeton University Press, Princeton, NJ, 1957.
- [5] R. W. James, Transformation of Spherical Harmonics Under Change of Reference Frame, *Geophysical Journal International* 17 (3) (1969) 305–316. doi:10.1111/j.1365-246X.1969.tb00239.x.
URL <http://www.blackwell-synergy.com/doi/abs/10.1111/j.1365-246X.1969.tb00239.x>
- [6] C. A. White, M. Head-Gordon, Rotating around the quartic angular momentum barrier in fast multipole method calculations, *The Journal of Chemical Physics* 105 (12) (1996) 5061–5067.
URL <http://link.aip.org/link/?JCP/105/5061/1>
- [7] M. A. Blanco, M. Flórez, M. Bermejo, Evaluation of the rotation matrices in the basis of real spherical harmonics, *Journal of Molecular Structure* 419 (1997) 19–27.
- [8] S. Kenyon, J. Factor, N. Pavlis, S. Holmes, Towards the Next Earth Gravitational Model, in: *Society of Exploration Geophysicists 77th Annual Meeting*, San Antonio, Texas, USA, 2007.
URL http://earth-info.nga.mil/GandG/wgs84/gravitymod/new_egm/EGM08_papers/EGM-2007-final.pdf
- [9] Z. Gimbutas, L. Greengard, A fast and stable method for rotating spherical harmonic expansions, *Journal of Computational Physics* 228 (16) (2009) 5621–5627. doi:10.1016/j.jcp.2009.05.014.

URL <http://linkinghub.elsevier.com/retrieve/pii/S0021999109002691>

- [10] D. Pinchon, P. E. Hoggan, Rotation Matrices for Real Spherical Harmonics: General Rotations of Atomic Orbitals in Space-Fixed Axes, *Journal of Physics A: Mathematical and Theoretical* 40 (2007) 1597–1610.
URL <http://stacks.iop.org/1751-8121/40/1597>
- [11] W. Freeden, T. Gervens, M. Schreiner, *Constructive Approximation on the Sphere (With Applications to Geomathematics)*, Oxford Sciences Publication. Clarendon Press, Oxford University, 1998.
- [12] N. Aronszajn, Theory of Reproducing Kernels, *Transactions of the American Mathematical Society* 68 (3) (1950) 337–404.
URL <http://www.jstor.org/stable/1990404>
- [13] S. Saitoh, *Integral Transforms, Reproducing Kernels and their Applications*, Pitman Research Notes in Mathematics, Longman Scientific & Technical, 1997.
- [14] C. Müller, *Spherical Harmonics*, Lecture Notes in Mathematics, Springer, 1966.
- [15] J. Cui, W. Freeden, Equidistribution on the Sphere, *SIAM J. Sci. Comput.* 18 (1997) 595–609.
doi:<http://dx.doi.org/10.1137/S1064827595281344>.
- [16] E. Saff, A. Kuijlaars, Distributing many points on a sphere, *The Mathematical Intelligencer* 19 (1) (1997) 5–11. doi:10.1007/BF03024331.
URL <http://www.springerlink.com/index/10.1007/BF03024331>
- [17] J. Kovacevic, A. Chebira, Life Beyond Bases: The Advent of Frames (Part I), *Signal Processing Magazine, IEEE* 24 (4) (2007) 86 – 104. doi:10.1109/MSP.2007.4286567.
URL http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4286567
- [18] S. G. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, third ed. Edition, Academic Press, 2009.

- [19] C. Zhu, R. H. Byrd, P. Lu, J. Nocedal, Algorithm 778: L-BFGS-B:Fortran subroutines for large-scale bound-constrained optimization, *ACM Transactions on Mathematical Software (TOMS)* 23 (4).
URL <http://portal.acm.org/citation.cfm?id=279236>
- [20] B. L. Kempster, Extension of the Whittaker-Shannon Sampling Series Aided by Symbolic Computation, M.Sc. thesis, Anglia Polytechnique University, Cambridge (1995).
- [21] J. R. Higgins, *Sampling Theory in Fourier and Signal Analysis: Foundations*, Oxford University Press, Oxford, 1996.
- [22] D. Stern, Classification Of Magnetic Shells, *J. Geophys. Res.* 70 (15) (1965) 3629–3634.
URL <http://www.agu.org/journals/ABS/1965/JZ070i015p03629.shtml>
- [23] D. Armentano, C. Beltrán, M. Shub, Minimizing the discrete logarithmic energy on the sphere: The role of random polynomials, *Trans. Amer. Math. Soc.* 363 (2011) 2955–2965.
URL <http://www.ams.org/journals/tran/2011-363-06/S0002-9947-2011-05243-8/home.html>
- [24] I. H. Sloan, R. Womersley, Extremal Systems of Points and Numerical Integration on the Sphere, *Advances in Computational Mathematics* 21 (1) (2004) 107–125. doi:10.1023/B:ACOM.0000016428.25905.da.
URL <http://www.springerlink.com/content/w2304u80w0542578/>
- [25] E. Bannai, R. M. Damerell, Tight spherical designs, I, *Journal of the Mathematical Society of Japan* 31 (1) (1979) 199–207.
URL <http://projecteuclid.org/euclid.jmsj/1240319488>
- [26] E. Bannai, R. M. Damerell, Tight Spherical Designs, II, *J. London Math. Soc.* s2-21 (1) (1980) 13–30.
- [27] F. J. Simons, F. A. Dahlen, M. A. Wieczorek, Spatiospectral Concentration on a Sphere, *SIAM Review* 48 (3) (2006) 504–536.
URL <http://link.aip.org/link/?SIR/48/504/1>
- [28] P. P. Vaidyanathan, *Sampling Theorems for Nonbandlimited Signals*, Birkhäuser, Boston, 2004, Ch. 5, pp. 115–136.