Fast and Accurate Rotation Estimation on the 2-Sphere without Correspondences

Janis Fehr, Marco Reisert, and Hans Burkhardt

Chair of Pattern Recognition and Image Processing
Institute for Computer Science
Albert-Ludwigs-University, Freiburg, Germany
fehr@informatik.uni-freiburg.de

Abstract. We present a refined method for rotation estimation of signals on the 2-Sphere. Our approach utilizes a fast correlation in the harmonic domain to estimate rotation angles of arbitrary size and resolution. The method is able to achieve great accuracy even for very low spherical harmonic expansions of the input signals without using correspondences or any other kind of a priori information. The rotation parameters are computed analytically without additional iterative post-processing or "fine tuning".

The theoretical advances presented in this paper can be applied to a wide range of practical problems such as: shape description and shape retrieval, 3D rigid registration, robot positioning with omni-directional cameras or 3D invariant feature design.

1 Introduction

Methods involving functions on the 2-Sphere have recently received growing attention in a wide range of different computer vision and pattern recognition problems. For example in the area of shape description and retrieval, various invariant shape-features which have been proposed are based on spherical functions [1] [2] [3] [4]. For biomedical applications, many algorithms including invariant 3D features [5] or rigid 3D registration for 3D volume data analysis have their mathematical foundations in the rotation group $SO(3)$ which directly implies the use of functions on spheres. In all these cases, a fast and accurate rotation estimation would provide a helpful tool for further feature design and descriptors. Finally, in robotics, positioning algorithms which use images of omni-directional cameras [6] [7] [8] directly rely on a precise rotation estimation on the 2-Sphere. **Problem statement:** given any two real valued signals f_1 and f_2 on a 2-sphere which are considered to be equal or at least similar under some rotational invariant measure ($\sim_R$):

$$f_1 \sim_R f_2, R \in SO(3) \tag{1}$$

the goal is to estimate the parameters of an arbitrary rotation R as accurate as possible without any additional information other than f_1 and f_2.

Obviously, the computational cost of a direct matching approach - testing all possible rotations R - is way too high. Especially when we are considering arbitrary resolutions of the rotation parameters.

D. Forsyth, P. Torr, and A. Zisserman (Eds.): ECCV 2008, Part II, LNCS 5303, pp. 239–251, 2008.

1.1 Related Work

Recently, there have been proposals for several different methods which try to overcome the direct matching problem. Here, we are only considering methods which provide full rotational estimates (there are many methods covering only rotations around the z-axis) without correspondences.

A direct nonlinear estimation (DNE) which is able to retrieve the parameters for small rotations via iterative minimization techniques was introduced in [7]. However, this method fails for larger rotations and was proposed only for "fine tuning" of pre-aligned rotations. Most other methods use representations in the spherical harmonic domain to solve the problem. The possibility to recover the rotation parameters utilizing the spherical harmonic shift theorem (SHIFT) [9] has been shown in [10]. This approach also uses an iterative minimization and was later refined by [8]. Again, the estimation accuracy is limited to small rotations.

The basic method which we extend in this paper is a fast full correlation (FCOR) in the spherical harmonic domain. This approach was first suggested by [11], stating a fast correlation in two angles followed by a correlation in the third Euler angle in an iterative way. This method was later extended to a full correlation in all three angles by [12]. This approach allows the computation of the correlation directly from the harmonic coefficients via FFT, but was actually not intended to be used to recover the rotation parameters. Its angular resolution directly depends on the range of the harmonic expansion - making high angular resolutions rather expensive. But FCOR has been used by [7] to initialize the DNE and SHIFT "fine tuning" algorithms. The same authors used a variation of FCOR (using inverse Spherical Fourier Transform [13] in stead of FFT) in combination with SHIFT [6] to recover robot positions from omni-directional images via rotation parameter estimation.

1.2 Contributions

The contribution of this work is twofold: first, we introduce novel theoretical extensions (3.2, 3.3) of the previous methods which allow accurate estimations for even very low spherical harmonic expansions without any additional iterative post-processing or "fine tuning". Secondly, we show how the rather complex theoretical methods can be implemented efficiently - previous publications gave very little insight into many practical implementation problems.

The remainder of this paper is structured as follows: first we revise the mathematical foundations of our approach in section (2). In section (3) we show the actual rotation estimation in all theoretical details, followed by methods for the practical implementation in section (4). Finally, we show detailed experiments in section (5).

2 Mathematical Foundations

Let us start with a very brief introduction of the basic mathematical tools and conventions used throughout the rest of the paper. For more detailed introductions please refer to [14] and [9].

Spherical Harmonics. Spherical Harmonics (SH) [14] form an orthonormal base on the 2-sphere. Analogical to the Fourier Transform, any given real valued signal f on a sphere with its parameterization over the angles Θ, Φ (latitude and longitude of the sphere) can be represented by an expansion in its harmonic coefficients:

$$f(\Theta, \Phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{m=l} \hat{f}_{lm} Y_m^l(\Theta, \Phi) \tag{2}$$

where l denotes the band of expansion, m the order for the l-th band and $\hat{f}_{lm}$ the harmonic coefficients. The harmonic base functions $Y_m^l(\Theta, \Phi)$ are calculated as follows:

$$Y_m^l(\Theta, \Phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} \cdot P_m^l(\cos\Theta) \mathrm{e}^{im\Phi} \tag{3}$$

where P_m^l is the associated Legendre polynomial.

The harmonic expansion of a signal f will be denoted by $\hat{f}$ with corresponding coefficients $\hat{f}_{lm}$. In our case, where we are only considering signals on a discrete grid in $\mathbb{R}^3$, the $\hat{f}_{lm}$ can be computed via point-wise multiplication $(\cdot)$ of the 3D data grid with pre-computed discrete approximations of the harmonic base functions of fixed radii (more details in section (4)):

$$\hat{f}_{lm} = \sum_{\mathbb{R}^3} Y_m^l \cdot f \tag{4}$$

Rotations in SH. Throughout the rest of the paper we will use the Euler notation in zyz-convention denoted by the angles φ, θ, ψ with $\varphi, \psi \in [0, 2\pi[$ and $\theta \in [0, \pi[$ to parameterize the rotations $R \in SO(3)$ (short hand for $R(\varphi, \theta, \psi) \in SO(3)$). Rotations $R(\varphi, \theta, \psi) \cdot f$ in the Euclidean space find their equivalent representation in the harmonic domain in terms of the so called Wigner D-Matrices, which form an irreducible representation of the rotation group $SO(3)$. For each band l, $D^l(\varphi, \theta, \psi)$ (or short handed $D^l(R)$) defines a band-wise rotation in the SH coefficients. Hence, a rotation in the Euclidean space can be estimated in the harmonic domain (with a maximum expansion b), by

$$R \cdot f \approx \sum_{l=0}^{b} \sum_{m=-l}^{l} \sum_{n=-l}^{l} D_{mn}^l(R) \hat{f}_{lm} Y_m^l \tag{5}$$

In section (4.1), we further discuss practical issues concerning the computation of the D-matrices.

3 Estimating Rotation Parameters

We use several steps to recover the rotation parameters. First, we apply a fast correlation in the harmonic domain, gaining a similarity measure as stated in (1). Within the harmonic domain, we are then extending the angular resolution of the correlation function and minimize ambiguities of the Euler representation before we recover the full correlation matrix via FFT.

3.1 Fast Correlation in SH

We are following the fast correlation method which was introduced in [12] and later used in a similar way by [8]. The full correlation function $c : SO(3) \to \mathbb{R}$ of two signals f and g under the rotation $R \in SO(3)$ on a 2-sphere is given as:

$$c(R) := \int_{S^2} f \cdot (R \cdot g) \quad d\phi d\theta d\psi \tag{6}$$

Using the DFT Convolution Theorem and substituting f and g with their SH expansions (5, 4), leads to

$$c(R) = \sum_{lmn} \overline{D^l_{mn}(R)} \hat{f}_{lm} \overline{\hat{g}_{ln}} \tag{7}$$

The actual "trick" to obtain the fast correlation is now to factorize the original rotation $R(\varphi, \theta, \psi)$ into $R = R_1 \cdot R_2$, choosing $R_1(\xi, \pi/2, 0)$ and $R_2(\eta, \pi/2, \omega)$ with $\xi = \varphi - \pi/2, \eta = \pi - \theta, \omega = \psi - \pi/2$.

Using the fact that

$$D^l_{mn}(\varphi, \theta, \psi) = \mathrm{e}^{-im\varphi} d^l_{mn}(\theta) \mathrm{e}^{-in\psi} \tag{8}$$

where d^l is a real valued and so called "Wigner (small) d-matrix" (see section (4.1)), and

$$D^l_{mn}(R_1 \cdot R_2) = \sum_{h=-l}^{l} D^l_{nh}(R_1) D^l_{hm}(R_2) \tag{9}$$

we can rewrite

$$D^l_{mn}(R) = \sum_{h=-l}^{l} d^l_{nh}(\pi/2) d^l_{hm}(\pi/2) \mathrm{e}^{-i(n\xi + h\eta + m\omega)} \tag{10}$$

Substituting (10) into (7) provides the final formulation for the correlation function regarding the new angles ξ, η and ω:

$$c(\xi, \eta, \omega) = \sum_{lmhm'} d^l_{mh}(\pi/2) d^l_{hm'}(\pi/2) \hat{f}_{lm} \overline{\hat{g}_{lm'}} \mathrm{e}^{-i(m\xi + h\eta + m'\omega)} \tag{11}$$

The direct evaluation of this correlation function is of course not possible - but it is rather straight forward to obtain the Fourier transform of (11), hence eliminating the missing angle parameters:

$$\hat{c}(m, h, m') = \sum_{l} d^l_{mh}(\pi/2) d^l_{hm'}(\pi/2) \hat{f}_{lm} \overline{\hat{g}_{lm'}} \tag{12}$$

Finally, the correlation $c(\xi, \eta, \omega)$ can be retrieved via inverse FFT of $\hat{c}$

$$c(\xi, \eta, \omega) = FFT^{-1}(\hat{c}(m, h, m')), \tag{13}$$

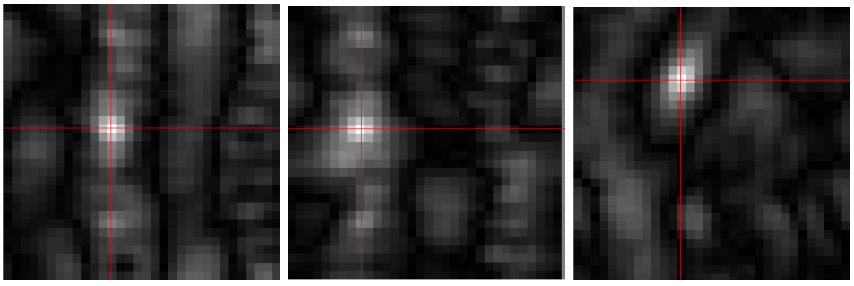

Fig. 1. Ortho view of a resulting 3D correlation grid in the (ξ, η, ω)-space with a maximum spherical harmonic expansion to the 16th band, $\phi = \pi/4, \theta = \pi/8, \psi = \pi/2$. From left to right: xy-plane, zy-plane, xz-plane.

revealing the correlation values on a sparse grid in a three dimensional (ξ, η, ω)-space. Figure (1) shows a resulting correlation grid for a sample rotation of random input data. Given our fast correlation, the basic idea is now to simply search the 3D grid for the position of the maximum correlation and compute the rotation parameters (φ, θ, ψ) from the (ξ, η, ω)-angles associated with this position.

3.2 Increasing the Angular Resolution

Let us take a closer look at Figure (1): first of all, it appears (and our experiments clearly support this assumption) that the fast correlation function has a clear and stable maximum in a point on the grid. This is a very nice property, and we could now simply recover the corresponding rotation parameters which are associated with this maximum position. But there are still some major problems: The image in Figure (1) appears to be quite coarse - and in fact, the parameter grids for expansions up to the 16th band have the size of $33 \times 33 \times 33$ since the parameters m, m', h in (12) are running from $-l, \ldots, l$. Given rotations up to 360°, this leaves us in the worst case with a overall estimation accuracy of less than 15°.

In general, even if our fast correlation function (13) would perfectly estimate the maximum position in all cases, we would have to expect a worst case accuracy of

$$Err_{corr} = 2 \cdot \frac{180^\circ}{2l} + \frac{90^\circ}{2l} \tag{14}$$

Hence, if we would like to achieve an accuracy of 1°, we would have to take the harmonic expansion roughly beyond the 180th band. This would be computationally expensive. Even worse, since we are considering discrete data, the signals on the sphere are band-limited. So for smaller radii, higher bands of the expansion are actually not carrying any valuable information.

Due to this resolution problem, the fast correlation has so far only been used to initialize iterative algorithms [6][7].

Sinc Interpolation. Now, instead of increasing the sampling rate of our input signal by expanding the harmonic transform, we have found an alternative way to increase the correlation accuracy: interpolation in the frequency domain.

In general, considering the Sampling Theorem and given appropriate discrete samples a_n with step size Δ_x of some continuous 1D signal $a(x)$, we can reconstruct the original signal via sinc interpolation [15]:

$$a(x) = \sum_{n=-\infty}^{\infty} a_n \mathrm{sinc}(\pi(x/\Delta_x - n)) \tag{15}$$

with

$$\mathrm{sinc}(x) = \frac{\sin(x)}{x} \tag{16}$$

For a finite number of samples (15) changes to:

$$a(x) = \sum_{k=0}^{N} a_k \frac{\sin(\pi(x/\Delta_x - k))}{N \sin(\pi(x/\Delta_x - k)/N)} \tag{17}$$

This sinc interpolation features two nice properties[15]: it entirely avoids aliasing errors and it can easily be applied in the discrete Fourier space. Given the DFT coefficients α_r of the discrete signal $a_n, n = 0, 1, \ldots, N-1$, sinc interpolation is simply a zero padding of the spectrum between $a_{N/2}$ and $a_{(N/2)-1}$.

Returning to our original correlation problem, it is easy to see that the (m, h, m')-space in (12) actually is nothing else but a discrete 3D Fourier spectrum. So we can directly apply the 3D extension of (17) and add a zero padding into the (m, h, m')-space. This way, we are able to drastically increase the resolution of our correlation function at very low additional cost (see section (4) for implementation issues as well as suitable pad sizes). Figure (2) shows the effect of the interpolation on the correlation matrix for different pad sizes p.

It has to be noted, that even though the sinc interpolation implies some smoothing characteristics to the correlation matrix, the maxima remain fixed to singular positions in the grid.

Finally, we are now theoretically able to reduce the worst case accuracy to arbitrarily small angles for any given band.

$$Err_{corr}^{pad} = 2 \cdot \frac{180^\circ}{2l + p} + \frac{90^\circ}{2l + p} \tag{18}$$

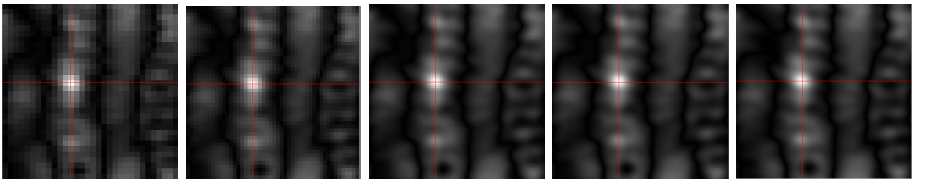

Fig. 2. Same experiment as in Fig. (1) but with increasing size of the sinc interpolation padding. From left to right: $p = 0, p = 16, p = 64, p = 128, p = 256$.

Of course, the padding approach has practical limitations - inverse FFTs are becoming computationally expensive at some point. But as our experiments show, resolutions below one degree are possible even for very low expansions.

3.3 Euler Ambiguities

After applying the sinc interpolation and the inverse FFT, (13) provides us with a fine parameter grid. The final obstacle towards the recovery of the rotation parameters inherits from the Euler parameterization used in the correlation function. Unfortunately, Euler angle formulations cause various ambiguities and cyclic shift problems.

One minor problem is caused by the fact that our parameter grid range is from $0, \ldots, 2\pi$ in all dimensions, while the angle θ is only defined $\theta \in [0, \pi[$. This causes two correlation peaks at $\theta = \beta$ and $\theta = 2\pi - \beta$ for an actual rotation of $\theta = \beta$. We avoid this problem by restricting the maximum search to $\theta \in [0, \pi[$, hence neglecting half of the correlation grid.

The formulation of the correlation function also causes further cyclic shifts in the grid representation of the Euler angles. This way, the zero rotation $R(\phi = 0, \theta = 0, \psi = 0)$ does not have its peak at the zero position $c(0, 0, 0)$ of the parameter grid as one would expect. For a more intuitive handling of the parameter extraction from the grid, we extend the original formulation of (12) and use a shift in the frequency space in order to normalize the mapping of $R(\pi, 0, \pi)$ to $c(0, 0, 0)$:

$$\hat{c}(m, h, m') = \sum_{l} d^l_{mh}(\pi/2) d^l_{hm'}(\pi/2) \hat{f}_{lm} \overline{\hat{g}_{lm'}} \cdot i^{m+2h+m'} \tag{19}$$

Rotation Parameters. Finally, we are able to retrieve the original rotation parameters. For a given correlation peak at the grid position $c(x, y, z)$, with maximum harmonic expansion b and padding p the rotation angles are:

$$\phi = \begin{cases} \pi + (2\pi - x\Delta) & \text{for } x\Delta > \pi \\ \pi - x\Delta & \text{otherwise} \end{cases} \tag{20}$$

$$\theta = \begin{cases} (2\pi - y\Delta) & \text{for } y\Delta > \pi \\ y\Delta & \text{otherwise} \end{cases} \tag{21}$$

$$\psi = \begin{cases} \pi + (2\pi - z\Delta) & \text{for } z\Delta > \pi \\ \pi - z\Delta & \text{otherwise} \end{cases} \tag{22}$$

$$\text{with} \quad \Delta = 2\pi/(b + p)$$

The resulting rotation estimates return very precise and unique parameter sets. Only one ambiguous setting has to be noted: for $\theta = 0, \pi$ all zyz-Euler formulations which hold $\phi + \psi = 2\pi$ encode the very same rotation (see Figure (3)). This is actually not a problem for our rotation estimation task, but might be quite confusing especially in the case of numerical evaluation of the estimation accuracy.

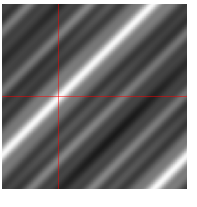

Fig. 3. $\phi\psi$-plane for of the correlation matrix with $\theta = 0$

3.4 Normalized Cross-Correlation

In many cases, especially when one tries to estimate the rotation parameters between non identical objects, it is favorable to normalize the (cross-)correlation results. We follow an approach which is widely known from the normalized cross-correlation of images: First, we subtract the mean from both functions prior to the correlation and then divide the results by the variances:

$$c_{norm}(R) := \int_{S^2} \frac{(f - \bar{f}) \cdot (R \cdot (g - \bar{g}))}{\sigma_f \sigma_g} \quad d\phi d\theta d\psi \tag{23}$$

Analogous to Fourier transform, we obtain the expected values $\bar{f}$ and $\bar{g}$ directly from the 0th SH coefficient. The variances σ_f and σ_g can be estimated directly from the band-wise energies:

$$\sigma_f \approx \sum_l |\hat{f}_l| \tag{24}$$

4 Implementation

An actual implementation of our algorithm can be quite challenging. In this section, we give some hints on how the methods can be implemented very efficiently and what pitfalls one should avoid. The main problem is, that even small implementation errors can easily produce correlation matrices which appear to be correct at first sight, but are in fact containing wrong estimates. Then, the search for errors in the $\hat{c}(m, h, m')$-space can be very tedious.

Harmonic Base Functions. The Spherical Harmonic base functions can be pre-computed off-line for a fixed radius and be stored in a file. The actual computation is straight forward (see [14]), but often not necessary because most common math tools like MATLAB or the GNU-scientific-library are already providing SH implementations.

Sinc Interpolation Padding. The implementation of the inverse FFT in (13) combined with the frequency space padding requires some care: we need an inverse complex to real FFT with an in-place mapping (the grid in the frequency space has the same size as the resulting grid in $\mathbb{R}^3$). Most FFT implementations are not providing such an operation. Due to the symmetries in the frequency space not all complex coefficients need to be stored, hence most implementations are using reduced grid sizes. We can avoid the tedious construction of such a reduced grid from $\hat{c}$ by using an inverse complex to complex FFT and taking only the real part of the result. In this case, we only have to shuffle the coefficients of $\hat{c}$, which can be done via simple modulo operations while simultaneously applying the padding. We rewrite (19) to:

$$\hat{c}(a, b, c) = \sum_l d^l_{mh}(\pi/2) d^l_{hm'}(\pi/2) \hat{f}_{lm} \overline{\hat{g}_{lm'}} \cdot i^{m+2h+m'} \tag{25}$$

where

$$s \coloneqq 2bp, \quad a \coloneqq (m+s+1) \bmod s, \quad b \coloneqq (h+s+1) \bmod s, \quad c \coloneqq (m'+s+1) \bmod s$$

Concerning the pad size: due to the nature of the FFT, most implementations achieve notable speedups for certain grid sizes. So it is very useful to choose the padding in such a way, that the overall grid size has, e.g., prime factor decompositions of mostly small primes [16].

4.1 Computation of the Wigner d-Matrices

The greatest implementation challenge is the actual computation of the Wigner d-Matrices. Even though, the d-Matrices for (25) can be pre-computed (we are always only considering $d^l_{mn}(\theta)$ with a fixed $\theta = \pi/2$), we still need an efficient implementation. In a direct approach, the d-Matrices can be computed by the sum

$$\begin{aligned} d^l_{mn}(\theta) = & \sum_t (-1)^t \frac{\sqrt{(l+m)!(l-m)!(l+n)!(l-n)!}}{(l+m-t)!(l-n-t)!t!(t+n-m)!} \\ & \cdot \cos(\theta/2)^{2l+m-n-2t} \cdot \sin(\theta/2)^{2t+n-m} \end{aligned} \tag{26}$$

over all t which lead to non-negative factorials [9]. It is easy to see that the constraints on t are causing the computational complexity to grow with the band of expansion. To overcome this problem, [17] introduced a recursive method for the d-Matrix computation. We are applying a closely related approach inspired by [18], where we retrieve d-matrices from recursively computed D-matrices.

Recursive Computation of the Wigner D-Matrices. Given D^l for the first two bands $l = 0$ and $l = 1$,

$$D^0(\phi,\theta,\psi) := 0$$

$$D^1(\phi,\theta,\psi) := \begin{pmatrix} \mathrm{e}^{-i\psi}\frac{1+\cos(\theta)}{2}\mathrm{e}^{-i\phi} & \frac{-\sin(\theta)}{\sqrt{2}}\mathrm{e}^{-i\phi} & \mathrm{e}^{i\psi}\frac{1-\cos(\theta)}{2}\mathrm{e}^{-i\phi} \\ \mathrm{e}^{-i\psi}\frac{\sin(\theta)}{\sqrt{2}} & \cos(\theta) & -\mathrm{e}^{i\psi}\frac{\sin(\theta)}{\sqrt{2}} \\ \mathrm{e}^{-i\psi}\frac{1-\cos(\theta)}{2}\mathrm{e}^{i\phi} & \frac{\sin(\theta)}{\sqrt{2}}\mathrm{e}^{i\phi} & \mathrm{e}^{i\psi}\frac{1+\cos(\theta)}{2}\mathrm{e}^{i\phi} \end{pmatrix}$$

we can compute D^l via band-wise recursion:

$$\begin{aligned} D^l_{mn} = & \sum_{m,m',n,n'=-l}^{l} D^1_{m'n'} D^{l-1}_{(m-m')(n-n')} \\ & \cdot < l-1, m | 1, m', l, m-m' > \\ & \cdot < l-1, n | 1, n', l, n-n' > \end{aligned} \tag{27}$$

where $< l, m | l', m', l'', m'' >$ denotes Clebsh-Gordan coefficients [9] known from angular momentum theory. Using (10), we finally obtain:

$$d^l_{mn}(\theta) = D^l_{mn}(0, \theta, 0) \tag{28}$$

5 Experiments

Unlike previous publications [6][7][8], which only performed a small set of experiments with a fixed number of predefined example rotations, we evaluate our methods with a series of large scale experiments on real word data.

Fig. 4. Volume rendering of an "ant"-object from the PSB

We use the "Princeton Shape Benchmark" (PSB) [1] dataset (which contains about 1800 3D objects) for our experiments. Figure (4) shows the volume rendering of a sample object. If not mentioned otherwise, all experiments have the same basic setup: for each parameter set, we evaluate the error statistics of 100 random rotations of random objects. We generate the rotations over all possible angles $\varphi, \psi \in [0, 2\pi[$ and $\theta \in [0, \pi[$ with a resolution of $0.001 \approx 0.1°$. Note that an error of $1° \approx 0.017$. All given error rates are the combined errors of all three angles.

5.1 Rotating Objects in the Harmonic Domain

In this first series of experiments, we extract a harmonic expansion with a fixed radius around the object center and then rotate this expansion using (5).

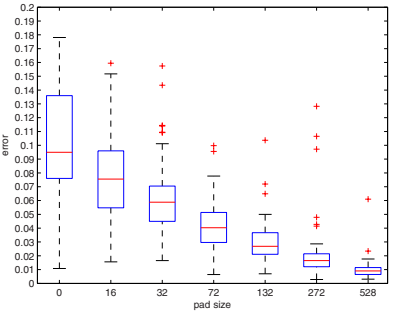

Fig. 5. Estimation errors with $b = 24$ and increasing pad size p

Pad Size. In a first experiment, we are able to show the effect of our padding method on the estimation accuracy. Figure (5) clearly shows the correlation of the pad size and the expected error. It is also evident, that we are able to achieve a precision of less than 1 degree while the experimental errors are found to be well within the theoretical bounds given in (18).

Maximum Band. The next two experiments investigate the practical influence of the maximum expansion band on the estimation errors. Figure (6) strongly supports our initial assumption, that the original formulation is not able to achieve accurate estimates for low expansions. Our method on the other hand, achieves very low error rates even for extremely low expansions with $b = 2$.

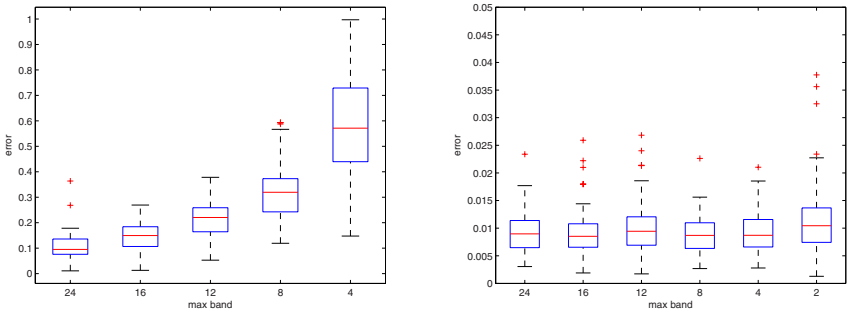

Fig. 6. Estimation errors with increasing maximum expansions. Left: $p = 0$. Right: $p > 512$ (p is not fix due to the padding to optimal FFT sizes). Note that the experiment with $p = 0, b = 2$ is left out because the result was so poor that it did not fit into the the chosen error scale.

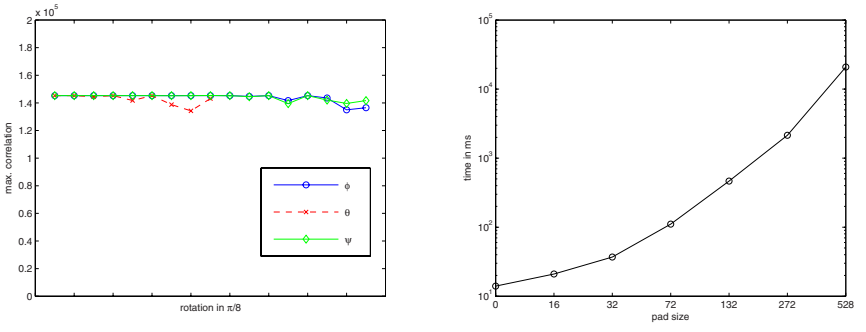

Fig. 7. Left: Maximum correlation for separate rotations in each angle. Right: Computational costs in ms on a standard PC.

Rotational Invariance and Computational Costs. are investigated in the last two experiments (figure (7)) of the first series. We rotate the object in $\pi/8$ steps in every angle to show that the correlation maximum is stable and indeed independent of the rotation. The computational complexity is largely dominated by the costs for the inverse FFT, hence growing with the pad size. So accuracy comes at some cost but reasonable accuracy can still be achieved well within 1 second.

5.2 Rotating Objects in $\mathbb{R}^3$

The results of figure (6) suggest, that the maximum expansion band has no influence on the quality of the rotation estimation - of course, this is only true if we are considering input signals which are limited to the very same maximum band. This is very unlikely for very low bands in the case of real data.

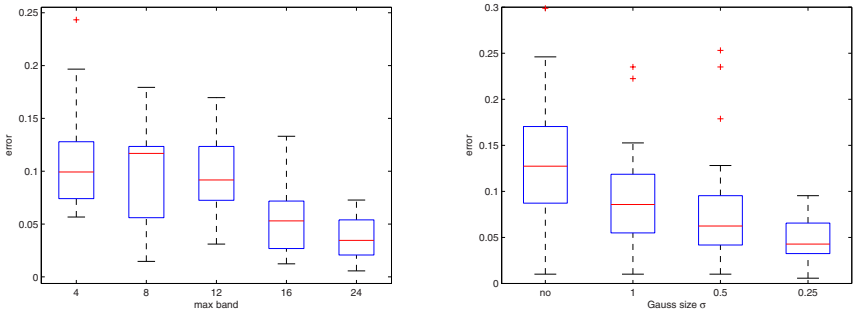

Fig. 8. Accuracy for rotations in $\mathbb{R}^3$. Left: Influence of the maximum band b, $p > 512, \sigma = 0.25$. Right: Correct sampling does matter! Without Gaussian smoothing and with different values for σ, $p > 512, b = 24$.

In order to evaluate the actual influence of the maximum expansion band, we need to rotate the objects in $\mathbb{R}^3$ and extract a second harmonic expansion after the rotation. As mentioned before, the usability of our sinc interpolation approach is limited to correctly sampled (concerning the Sampling Theorem) input signals. Hence, one must not expect to obtain precise rotation estimates for low band expansions, which act as a low pass filter, of high frequent input signals. Luckily, for most input data, we are not depending on the high frequent components in order to find the maximum correlation. Hence, we can apply a low pass filter (Gaussian) on the input data prior to the harmonic expansion. Figure (8) shows the impact of the maximum band and smoothing for rotations in $\mathbb{R}^3$. Overall, the estimation results are slightly worse than before, but are still quite reasonable.

6 Conclusions

We presented a fast and accurate rotation estimation on the 2-Sphere. We were able to show that our method is capable to estimate arbitrary rotations directly in the harmonic domain without any further iterative "fine tuning". We also introduced several implementation techniques, which allow a robust and efficient realization of the estimation algorithm.

Despite the theoretic nature of this paper, we emphasize the practical relevance of the rotation estimation for e.g. robotics or feature design.

References

1. Shilane, P., Min, P., Kazhdan, M., Funkhouser, T.: The princeton shape benchmark. In: Shape Modeling International, Genova, Italy (2004)
2. Kazhdan, M., Funkhouser, T., Rusinkiewicz, S.: Rotation invariant spherical harmonic representation of 3d shape descriptors. In: Symposium on Geometry Processing (2003)

3. Saupe, D., Vranic, D.V.: 3d model retrieval with spherical harmonics and moments. In: Radig, B., Florczyk, S. (eds.) DAGM 2001. LNCS, vol. 2191, pp. 392–397. Springer, Heidelberg (2001)
4. Vranic, D.V.: An improvement of rotation invariant 3d shape descriptor based on functions on concentric spheres. In: IEEE International Conference on Image Processing (ICIP 2003), vol. 3, pp. 757–760 (2003)
5. Ronneberger, O., Fehr, J.: Voxel-wise gray scale invariants for simultaneous segmentation and classification. In: Kropatsch, W.G., Sablatnig, R., Hanbury, A. (eds.) DAGM 2005. LNCS, vol. 3663, pp. 85–92. Springer, Heidelberg (2005)
6. Makadia, A., Daniilidis, K.: Rotation recovery from spherical images without correspondencesameesh. IEEE Transactions on Pattern Analysis and Machine Intelligence 28 (2006)
7. Makadia, A., Sorgi, L., Daniilidis, K.: Rotation estimation from spherical images. In: International Conference on Pattern Recognition, Cambridge (2004)
8. Makadia, A., Daniilidis, K.: Direct 3d-rotation estimation from spherical images via a generalized shift theorem. In: IEEE Conference on Computer Vision and Pattern Recognition, Madison (2003)
9. Brink, D., Satchler, G.: Angular Momentum, 2nd edn. Clarendon Press, Oxford (1968)
10. Burel, G., Henoco, H.: Determination of the orientation of 3d objects using spherical harmonics. Graph. Models Image Process. 57, 400–408 (1995)
11. Crowther, R.A.: The molecular replacement method. In: Rossmann, M. (ed.) Gordon and Breach, New York, pp. 173–178 (1972)
12. Kovacs, J.A., Wriggers, W.: Fast rotational matching. Acta Crystallogr (58), 1282–1286 (2002)
13. Healy Jr., D., Rockmore, D., Kostelec, P., Moore, S.: Ffts for the 2-sphere - improvements and variations. The Journal of Fourier Analysis and Applications 9, 341–385 (2003)
14. Groemer, H.: Geometric Applications of Fourier Series and Spherical Harmonics. Cambridge University Press, Cambridge (1996)
15. Yaroslavsky, L.: Boundary effect free and adaptive discrete signal sinc-interpolation algorithms for signal and image resampling. Appl. Opt. 42, 4166–4175 (2003)
16. Frigo, M., Johnson, S.G.: The design and implementation of FFTW3. Proceedings of the IEEE 93, 216–231 (2005); special issue on Program Generation, Optimization, and Platform Adaptation
17. Trapani, S., Navaza, J.: Calculation of spherical harmonics and wigner d functions by fft. applications to fast rotational matching in molecular replacement and implementation into amore. Acta Crystallogr, 262–269 (2006)
18. Reisert, M., Burkhardt, H.: Irreducible group representation for 3d shape description. In: Franke, K., Müller, K.-R., Nickolay, B., Schäfer, R. (eds.) DAGM 2006. LNCS, vol. 4174, pp. 132–142. Springer, Heidelberg (2006)