

Shape Registration with Spherical Cross Correlation

Boris Gutman, Yalin Wang, Tony Chan, Paul M. Thompson, Arthur W. Toga

► **To cite this version:**

Boris Gutman, Yalin Wang, Tony Chan, Paul M. Thompson, Arthur W. Toga. Shape Registration with Spherical Cross Correlation. Xavier Pennec. 2nd MICCAI Workshop on Mathematical Foundations of Computational Anatomy, Oct 2008, New-York, United States. pp.56-67, 2008. <inria-00632874>

HAL Id: inria-00632874

<https://hal.inria.fr/inria-00632874>

Submitted on 16 Oct 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Shape Registration with Spherical Cross Correlation

Boris Gutman^{1,2} {bgutman@ucla.edu}, Yalin Wang^{1,2}, Tony Chan¹, Paul M. Thompson², and Arthur W. Toga²

¹ UCLA Department of Mathematics

² UCLA Laboratory of Neuro Imaging

Abstract. We present a framework for shape alignment that generalizes several existing methods. We assume that the shape is a closed genus zero surface. Our framework requires a diffeomorphic surface mapping to the 2-sphere which preserves rotation. Our similarity measure is a global spherical cross-correlation function of surface-intrinsic scalar attributes, weighted by the cross-correlation of the parameterization distortion. The final similarity measure may be customized according to the surface-intrinsic scalar functions used in the application.

1 Introduction

Problems of shape alignment are ubiquitous in medical imaging. While many problem-specific solutions exist for particularly common cases (e.g. cortex, hippocampus) [1–3], high quality general shape alignment remains very much an open problem. Our framework generalizes some existing methods without assuming the existence of any landmarks or data-specific features. We assume that the shape is a closed genus zero surface. Our framework requires a diffeomorphic surface mapping to the 2-sphere which preserves rotation. Our similarity measure is a global spherical cross-correlation function of surface-intrinsic scalar attributes, weighted by the cross-correlation of the parameterization distortion, sometimes known as the conformal factor. The final similarity measure may be customized according to the surface-intrinsic scalar functions used in the application. Higher order scalar functions such as mean and Gaussian curvature may be used in conjunction with low order ones, like distance to mass center, to incorporate more localized shape information in addition to global measures.

In this study we have used a global conformal mapping as the spherical parameterization, and only the distance to mass center as intrinsic scalar shape measure. Using our method, we created atlases and registered shapes from a population of hippocampi.

2 Previous work

There is doubtless a galaxy of existing general shape alignment methods, and due to space limitations we will only mention those most prevalent and those closest

to the present work. Davies [4] has used an information-theoretic framework to formulate shape alignment as a minimum description length (MDL) problem. Here, each point is treated as an independent variable, while the cost function is aimed at reducing the "code length" of each shape's representation in this shape space. The beauty of this statistical approach lies in its ability to register multiple shapes simultaneously without the need to select a "target" shape. Of course, this is also its limitation, as for example when a known atlas exists and all data shapes are to be registered to it. The method requires a spherical diffeomorphism like ours.

Variants of the well-known ICP algorithm are another variety of recent developments in rigid shape alignment. Granger [5] introduced the EM-ICP method. ICP's proclivity for terminating at suboptimal local minima is greatly reduced by treating the problem as a general expectation-maximization problem. A multiscale approach is used: at coarser scales, the blurring factor is sufficient to give crude but correct alignment, which is improved upon in later refined stages, where the algorithm approaches the original ICP. Though quite robust, the method depends on the scale of blurring factor being set correctly to avoid local minima. Thus, avoiding them is not guaranteed. Our method, by contrast, performs a global search non-iteratively. Thus, a global maximum correlation is guaranteed irrespective of the shape's original orientation without the need to tune any parameters.

Much like our algorithm, some previous methods have used spherical harmonics for rigid shape alignment. Among them are the first order ellipsoid (FOE) method, popularized by Brechbuhler [6] and used extensively in medical imaging [7, 8] applications and SHREC, a recent variant of the ICP algorithm. Like ours, these methods make use of rotational properties of spherical harmonics.

FOE alignment uses the fact that a shape reconstructed from only the first order spherical harmonics forms an ellipsoid in object space. The method works well when the ellipsoid's three axes have distinct lengths, which largely depends on the shape itself and the degree to which the spherical parameterization preserves area. The method gives only a crude alignment and fails when two or more axes have similar lengths. Even with a proper ellipsoid, there is a symmetry problem.

SHREC [9] is another variation of the ICP. As in our case, the correspondence search is done iteratively on the sphere via Euler's rotation formula and icosahedral subdivision. The mapping satisfies our conditions, while the similarity measure is the RMSD. Since RMSD depends on the position of the object in space, the algorithm requires an initial pre-alignment in both spaces. Rigid Quaternion transform (RQT) is used to align shapes in \mathbb{R}^3 after each iteration of parametric alignment. The main limitation of this algorithm is that it is not guaranteed to converge to the optimal solution (i.e. a correspondence which, when applied to RQT, minimizes RMSD). This is because the optimization in parameter space depends on the object's position in native space. Thus, though the search is more global than in the original ICP, the parameter space search is still locally biased. The original ICP suffers the same problem.

The last algorithm uses a brute force correlation: it computes a cost function anew for each rotation. To mitigate the cost of this, a hierarchical approach is used. Instead, we reduce parameter space alignment to a global refined search via the FFT. Our numerical scheme separates the effects on computation time of the level of detail used for alignment and the number of rotation samples. In SHREC, these two are tied together owing to their brute-force nature. This means that we can refine rotation space tessellation while maintaining the same level of surface detail without significantly affecting computation time. Our use of orientation-invariant shape attributes in conjunction with scale invariant cross-correlation makes our approach completely independent of changes in object position and size. SHREC, by contrast, requires volume normalization as a pre-processing step.

3 Mathematical preliminaries

3.1 Spherical harmonics

Spherical harmonics are functions $f : \mathbb{S}^2 \rightarrow \mathbb{C}$ which are simultaneously eigenfunctions of the Laplace-Beltrami and the angular momentum operators; they are expressed explicitly as

$$Y_l^m(\theta, \phi) = \sqrt{\frac{(2l+1)(l-m)!}{4(l+m)!}} P_l^m(\cos\theta) e^{im\phi} \quad (1)$$

for degree and order $m, l \in \mathbb{Z}$, $|m| \leq l$, where $P_l^m(x)$ is the associated Legendre polynomial. Spherical harmonics form a countable orthonormal basis for square-integrable functions on the sphere. A projection of a function $f \in L_2(\mathbb{S}^2)$ onto this basis yields the SPH coefficients

$$\hat{f}(l, m) = \langle f, Y_l^m \rangle \quad (2)$$

where $\langle f, g \rangle$ is the usual L_2 inner product.

A key property of spherical harmonics is their behavior under a shift on the sphere. Given an element of the rotation group $R \in SO(3)$, a rotated spherical harmonic is expressed as

$$Y_l^m(\omega) = \sum_{n=-l}^l Y_l^n(R^{-1}\omega) D_{m,n}^l(R) \quad (3)$$

where

$$D_{m,n}^l(R) = e^{-i(m\alpha+n\gamma)} d_{m,n}^l(\beta), \quad (4)$$

α, β, γ are the Euler angles of R and $d_{m,n}^l$ are irreducible representations of $SO(3)$ [10],

$$d_{m,n}^l(\beta) = \sum_t (-1)^t \times \frac{\sqrt{(l+n)!(l-n)!(l+m)!(l-m)!}}{(l+n-t)!(l-m-t)!(t+m-n)!t!} \quad (5)$$

$$\times \left(\cos \frac{\beta}{2}\right)^{2l+n-m-2t} \left(\sin \frac{\beta}{2}\right)^{2t+m-n}.$$

In particular, (3) implies that

$$f(\omega) = g(R^{-1}\omega) \implies \widehat{f(l,m)} = \sum_{n=-l}^l \widehat{g(l,n)} D_{m,n}^l(R). \quad (6)$$

3.2 Discrete spherical cross correlation

Given two functions $f, g \in L_2(\mathbb{S}^2)$, their spherical cross-correlation is defined as

$$C_{f,g}(R) = \int_{\mathbb{S}^2} f(\omega) g(R^{-1}\omega) d\omega. \quad (7)$$

In the special case where $f(\omega) = g(R'^{-1}\omega)$, $C_{f,g}(R)$ is maximized when $R = R'$, assuming that f is not spherically symmetric. For bandlimited functions, i.e. for those functions whose spherical harmonic coefficients vanish for all $l \geq B$ for some bandwidth B , the correlation becomes

$$C_{f,g}(R) = \sum_{l=0}^{B-1} \sum_{m=-l}^l \widehat{f}(l,m) \overline{\widehat{g}(l,m)} \Lambda(R). \quad (8)$$

Here, $\Lambda(R)$ is the operator associated with the rotation matrix. The expression for shifted spherical harmonic coefficients (6) implies that

$$C_{f,g}(R) = \sum_{l,m,n} \widehat{f}(l,m) \overline{\widehat{g}(l,n)} D_{m,n}^l(R). \quad (9)$$

This expression forms the basis of our similarity measure.

3.3 Fast cross correlation via FFT

The material presented so far has been used in the prior works we mentioned. Now, we present a simple lemma which leads to a great speed up in computing the correlation (9). It suffices to make the observation that any rotation $R(\alpha, \beta, \gamma)$ may be expressed as a product of two rotations:

$$R(\alpha, \beta, \gamma) = R_1(\alpha + \pi/2, \pi/2, 0) R_2(\beta + \pi, \pi/2, \gamma + \pi/2). \quad (10)$$

Now using the fact that

$$D_{m,n}^l(R_1 \cdot R_2) = \sum_k D_{m,k}^l(R_1) D_{k,n}^l(R_2), \quad (11)$$




Fig. 1. Two randomly selected hippo surfaces. The red circle is homologous to the north pole on the sphere; the blue circle, to the south pole; and the blue line, to $\phi = 0$. The initial spherical mappings do not align these two very well.

we substitute (10) into (9) to obtain

$$\begin{aligned}
 C_{f,g}(R(\alpha, \beta, \gamma)) &= \sum_{l,m,n,k} \widehat{f}(l, m) \overline{\widehat{g}(l, n)} \quad (12) \\
 &\times D_{m,k}^l(\alpha + \pi/2, \pi/2, 0) D_{k,n}^l(\beta + \pi, \pi/2, \gamma + \pi/2) \\
 &= \sum_{l,m,n,k} \widehat{f}(l, m) \overline{\widehat{g}(l, n)} d_{m,k}^l(\pi/2) d_{k,n}^l(\pi/2) \times e^{i(m(\alpha+\pi/2)+k(\beta+\pi)+n(\gamma+\pi/2))} \\
 &= \text{IFFT} \left\{ \sum_l \widehat{f}(l, m) \overline{\widehat{g}(l, n)} d_{m,k}^l(\pi/2) d_{k,n}^l(\pi/2) \right\} (\alpha + \pi/2, \beta + \pi, \gamma + \pi/2).
 \end{aligned}$$

This simple result has been shown elsewhere [11], but to the best of the authors' knowledge this is the first time it has been used for shape registration.

4 Shape registration with cross correlation

4.1 Similarity measure

Given a 2-manifold $M \subset \mathbb{R}^3$, a diffeomorphic spherical parameterization $f : \mathbb{S}^2 \mapsto M$ and a family of rotation-invariant shape attributes $s_i : M \mapsto \mathbb{R}$, $0 < i \leq N$, let

$$S_i = s_i \circ f. \quad (13)$$

Then, given two manifolds M_1, M_2 and their corresponding shape attribute maps $\{S_{1,i}\}_{i=0}^N, \{S_{2,i}\}_{i=1}^N$, we define our shape similarity measure as




Fig. 2. Top surface from figure 1 is correlated to the bottom one, and its signal is shifted on the sphere without changing the shape's position in object space. The homologous points now appear to be in good correspondence between the surfaces.

$$\mathcal{S}_{M_1, M_2}(R) = C_{\lambda_1, \lambda_2}(R) \sum_{i=1}^N \kappa_i C_{S_{1,i}, S_{2,i}}(R), \quad (14)$$

where κ_i are user-defined shape attribute weights, and λ_1, λ_2 are spherical maps of the conformal factor of each manifold. These last two are used to mitigate the fact that scalar functions which appear similar on the sphere may in fact represent vastly differently-sized regions on the original surfaces due to varying area distortion of the spherical map.

Because we recover shifts in object space with shifts on \mathbb{S}^2 , we require that the spherical parameterization preserve rotation in the following sense. Suppose $M_2 = R \circ M_1$, and $f_1, f_2 : \mathbb{S}^2 \mapsto M_1, M_2$ are spherical maps. Then

$$f_2(\omega) = R \circ f_1(R^{-1}\omega). \quad (15)$$

Many existing parameterizations satisfy this requirement, e.g. [6, 12].

4.2 Previous methods as special cases

SHREC and FOE are special cases of our method. FOE simply takes the conformal factor λ to be constant and uses spherical harmonics up to order one only. This is equivalent to setting the bandwidth B to 2. The single shape attribute used is the Euclidian distance to the surface average value. This is roughly the same as distance to mass center, especially for area-preserving spherical maps with which FOE is typically used.

SHREC minimizes RMSD, which can be reduced to spherical cross-correlation since $4\pi \text{RMSD}_{M_1, M_2}^2(R) =$

$$\begin{aligned} & \sum_{i \in \{x, y, z\}} \int_{\mathbb{S}^2} \|S_{1,i} - \Lambda(R)S_{2,i}\|^2 \\ &= \sum_{i \in \{x, y, z\}} \int_{\mathbb{S}^2} [S_{1,i}^2 + (\Lambda(R)S_{2,i})^2 - 2S_{1,i}\Lambda(R)S_{2,i}] \end{aligned} \quad (16)$$

$$= \sum_{i \in \{x,y,z\}} \|S_{1,i}\|_2^2 + \|S_{2,i}\|_2^2 - 2C_{S_{1,i},S_{2,i}}(R)$$

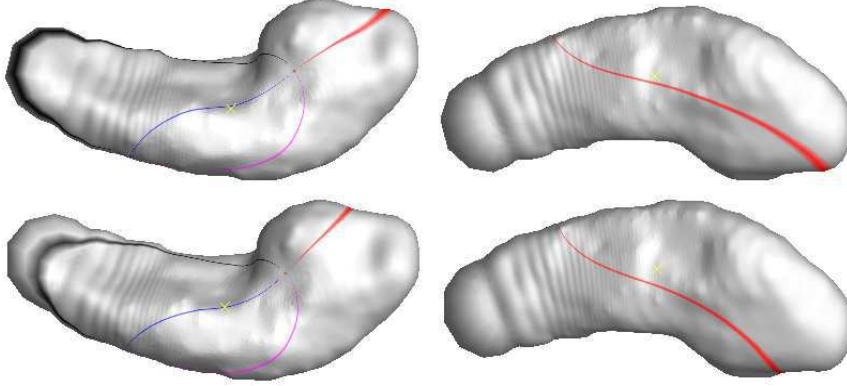


Fig. 3. Top surface from figure 1 is initially aligned using FOE (top row) and subsequently aligned using a modified version of SHREC (bottom row). Though this is slightly different from the original SHREC, in principle the two algorithms are the same. Here, we see the local minimum problem suffered by SHREC, typical of an ICP-type algorithm: there is very little change in alignment after initialization.

Thus, minimizing RMSD is equivalent to maximizing the correlation of the two shapes' spatial coordinates. These are, of course, not quite the scalar shape attributes we intend to use in our similarity measure. SHREC's dependence on iterative RQT refinement for correspondence optimization makes it less robust. This is the price of using orientation-dependent features. Further, it is not clear whether the correspondence which, when applied to RQT, minimizes RMSD is truly the best correspondence. One can conceive two shapes with some patches quite similar and other very different. One may then like to align the two objects to get the best correspondence between the similar patches without regard to the different ones. In such a case, cross correlation of invariant features will achieve a better alignment. Still, SHREC could be made faster with the use of FFT-based correlation rather than a brute-force approach.

SHREC requires $O(B^3 N_R)$ operations, where B is again the bandwidth and N_R the number of rotation samples. This is because recomputing $D_{m,n}^l(\alpha, \beta, \gamma)$ and the corresponding shifted spherical harmonic coefficients requires $O(B^3)$ operations. Our method requires $O(B^4 + N_R \log(N_R))$. N_R is roughly also of the order B^3 even with hierarchical sampling; this means we have effectively reduced the order of operations from $O(B^6)$ to $O(B^4)$. This allows us to sample rotation space more finely without a significant change in execution time. It also makes using higher order coefficients and hence greater level of detail for alignment purposes feasible. We see this in experiments below.




Fig. 4. Average of two shapes from figure 1 before correlation.




Fig. 5. Average of two shapes from figure 1 after correlation.

5 Results

We used a population of 45 right hippocampal surfaces extracted from healthy elderly subjects. Our spherical parameterization was the global conformal map of Gu. et al. [12]. Spherical harmonics were computed with the spherical FFT of [13] and cross correlation was computed with the help of the fftw library [14]. $SO(3)$ was sampled at $200 \times 100 \times 200$ samples, which yielded an angle frequency of ~ 1.8 degrees. A bandwidth of 64 was used throughout the testing. Only the distance to mass center was used as an invariant shape attribute for both populations.

As a preliminary experiment, we applied the cross correlation algorithm to a pair of hippocampal surfaces shown in figure 1. Here we see that the initial spherical maps do not provide a very good correspondence. Figure 2 shows the result of a spherical shift based on cross correlation. Figures 4 and 5 show two point-wise averages of the shapes, before and after cross correlation. The improvement is obvious. For $B = 64$, the average running time for above experiment was 44.6 ± 3.6 seconds on a Gateway 7426GX Laptop with a 2.41 GHz AMD processor and 1 GB RAM, tested with 45 hippocampal surfaces. This includes computation of spherical harmonics, rotation matrices $d_{m,n}^l$ and correlation. Shen [9] reports an average of 23.5 seconds running time on a common laptop for SHREC, while using $B = 12$. One would expect SHREC to take $(64/12)^3$ times longer for our bandwidth (see above), or on the order of 60 minutes. As a preliminary comparison, we implemented a modified version of SHREC. The only difference with the original is that at each iteration, our parameter alignment was initially done with cross correlation as described above, and subsequently refined according to the scheme outlined in [9]. Again, a bandwidth of 64 was used. Due to time limitations, we could not run this program on our whole dataset. We only present the results of one subject in figure 3. Here, we can see that the initial alignment

determines the final result to a great degree. This example converged after only 4 iterations. Execution time was close to 7 minutes. Since the execution time reported in [9] was for a MATLAB implementation, while ours is in C++, and because we use a fast cross-correlation, this time is significantly lower than the 60 minute estimate.

To test the effect of higher coefficients on correlation quality, we increased the bandwidth to 128, while keeping the same rotation tessellation and correlated 6 of the 45 subjects in our HP population. Running time increased to 404 ± 10.4 seconds, while the shape distance to the target hippocampus decreased on average only 3.9 ± 9.4 %. To test for the effect of rotation sampling frequency, we also decreased N_R to $100 \times 50 \times 100$ samples, while keeping the bandwidth at 64. Running time was 38.2 ± 3.5 seconds, an insignificant improvement in speed. In another experiment, we limited our bandwidth to 16 and reran cross correlation based alignment while keeping angle resolution at $200 \times 100 \times 200$. Execution time was reduced to about 3 seconds.




Fig. 6. Average of 45 right hippocampi with cross correlation, bandwidth = 64.




Fig. 7. Average of 45 right hippocampi with FOE.

We constructed hippocampal averages using FOE and two versions of cross correlation ($B = 16$ and $B = 64$) and compared results. First, a shape was selected, and all remaining shapes registered to it using each of the methods. Then, the shapes were averaged, normalized for volume (after registration, only




Fig. 8. Average of 45 right hippocampi with cross correlation, bandwidth = 16.

for a fair distance comparison) and a rigid quaternion transform applied to each shape to align it both to the target subject and to the volume-normalized average shape. Our shape distance was defined as vertex-wise distance between the surfaces, weighted by product of the sum of areas of adjacent triangles in each mesh. The results are in table 1. The table does not show a notable fact: distance to the target subject was improved by our method for every subject compared to FOE. Minimal improvement was 8%, and maximal 117%. All but one subject registered with $B = 64$ had superior alignment to the result of using $B = 16$. Compared to FOE, $B = 16$ reduced shape distance for all but 5 subjects. Table 2 shows a summary of intra-subject differences by registration method. Using a bandwidth of 16 gives a significant improvement compared to FOE, but the results are still much improved by using a bandwidth of 64. Figures 6 - 8 illustrate the hippocampal averages achieved with the three methods. Note that these averages were computed without spatially aligning the subjects to the target. Doing so would have likely given a more detailed shape.

Method	Distance to Target	Distance to Average
FOE	4.14 ± 1.33	2.89 ± 1.05
Cross cor.	2.64 ± 0.92	2.36 ± 0.75
$B = 16$	3.38 ± 1.06	2.33 ± 0.85

Table 1. Weighted point-wise distance of 45 (44) hippocampal surfaces by registration method.

We should note here that the poor result of the FOE average is due in part to the large area distortion of the conformal map. The tail of the hippocampus is mapped to such a small region on \mathbb{S}^2 that it is very hard to match well. We see in figure 7 that the tail suffers the most. The resulting spherical harmonic representation contains much redundant information; hence, the area distortion partially cancels out the benefits of using high order coefficients. The mapping used in [6, 9] is by contrast area-preserving, and hence better suited for alignment. The results of [6, 8] look closer to the one achieved here with cross correlation. This is, however, indicative of the potential our method has when applied to area-

preserving spherical maps, and we intend to experiment with this idea in the immediate future.

common shape	FOE vs. B = 64	FOE vs. B = 16	B = 16 vs. B = 64
target	61 % \pm 26 %	25 % \pm 23 %	30 % \pm 18 %
average	24 % \pm 29 %	26 % \pm 22 %	-6 % \pm 29 %

Table 2. Intra-subject improvement by registration method, in percentage of the second method’s result.

6 Conclusion

We have presented a framework for shape alignment which generalizes several existing methods. Our method is robust, fast and allows for use of greater detail in alignment than was possible before. The correspondence search is performed globally and no pre-alignment is required; thus, the result and computation time are independent of the shape’s size and initial orientation. Reaching a global maximum is always guaranteed. Our method can be tailored to suit a particular application by selecting the appropriate shape features for a particular data type. We intend to experiment with various shape attributes, apply area-preserving spherical maps to our method and extend the technique to automated patch selection and matching. Lastly, all software used here is available through the LONI Pipeline environment. Please contact the authors for more details.

References

1. Shi, Y., Thompson, P., et al.: Direct mapping of hippocampal surfaces with intrinsic shape context. *Neuroimage* **37**(3) (2007) 792–807
2. Tosun, D., Prince, J.: Cortical surface alignment using geometry driven multispectral optical flow. In: *IPMI 2005*. (2005)
3. Fischl, B., et al.: High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping* **8** (1999) 272–284
4. Davies, R., et al.: A minimum description length approach to statistical shape modeling. *IEEE Transactions on Medical Imaging* **21**(5) (2002) 525–537
5. Granger, S., Pennec, X.: Multi-scale em-icp: A fast and robust approach for surface registration. In: *Lecture Notes in Computer Science*. (2002)
6. Brechbuhler, C., Gerig, G., Kubler, O.: Parameterization of closed surfaces for 3d shape description. *Computer Vision and Image Understanding* **61**(2) (1995) 154–170
7. Gerig, G., Styner, M.: Shape versus size: Improved understanding of the morphology of brain structures. In: *MICCAI 2001*. (2001)
8. Gerig, G., Styner, M., et al.: Shape analysis of brain ventricles using spharm. In: *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MM-BIA’01)*. (2001)

9. Shen, L., Huang, H., Makedon, F., Saykin, A.: Efficient registration of 3d spharm surfaces. In: IEEE 4th Canadian Conference on Computer and Robot Vision. (2007)
10. Talman, J.: Special functions: A group-theoretic approach. W.A.Benjamin, Inc. (1968)
11. Sorgi, L., Daniilidis, K.: Template gradient matching in spherical images. Image Processing: Algorithms and Systems III. Proceedings of the SPIE **5298** (2004) 88–98
12. Gu, X., Wang, Y., et al.: Genus zero surface conformal mapping and its application to brain surface mapping. IEEE Transactions on Medical Imaging **23**(8) (2004) 949
13. Healy, D., Rockmore, D., et al.: Ffts for the 2-sphere-improvements and variations. J. Fourier Anal. Applicat. **9**(4) (2003) 341–385
14. Frigo, M.: A fast fourier transform compiler. In: ACM SIGPLAN Conference on Programming Language Design and Implementation. (1999)