

Proceedings of the MICCAI Workshop on Mesh Processing in Medical Image Analysis

In conjunction with MICCAI 2011
<http://www.imm.dtu.dk/MeshMed>

Sponsored by:

- NIH/NCRR Center for Integrative Biomedical Computing at the Scientific Computing and Imaging Institute, University of Utah
- Department of Mathematical Modeling at the Technical University of Denmark



Westin Harbor Castle, Toronto, Canada, September 18, 2011

Forward

MeshMed 2011, a workshop on mesh processing in medical image analysis, was held at Westin Harbor Castle in Toronto, Canada on September 18, 2011. This workshop, the first of its kind, was a joint effort supported from the Center for Integrated Biomedical Computing at the Scientific Computing and Imaging Institute and Department of Mathematical Modeling at the Technical University of Denmark. The generous support of DTU and the CIBC was integral in making this workshop possible.

The primary goal of this event was to bring together the geometry processing, computer graphics, and medical imaging communities to foster joint efforts and specific research. A fundamental theme, the image analysis pipeline and the significant role of geometric computation within it, was a central focus of the submitted works and discussions. Many of the technologies currently being developed by researchers in these communities are done so in an independent and isolated manner. We believe there is significant, fruitful research to be done that requires cross pollination of the geometry and imaging communities. Developing more sophisticated understanding of the interplay between surface and volume representations (meshes) and the clinical analysis and simulation is of utmost importance.

Of the nearly twenty submissions, we accepted eight for presentation along with our three plenary talks representative of the meshing, surfacing, and medical imaging communities. These works covered a broad range of topics, including integrating recent software tools, evaluation and comparisons of meshing techniques, and surface feature identification and deformation techniques. These great research ideas were then applied to image data acquired from cardiac, cerebral, and orthopedic domains.

We would like to thank the program and organizing committees for contributing both their time and ideas. Their efforts pushed this international project forward and enriched both its content and structure. Additionally, we would like to thank all authors who submitted works. Their provided material made this workshop a success, which we hope will initiate a new interest and appreciation on this topic. Finally, we are exceptionally grateful to the invited speakers for contributing their time and attracting external interest for the workshop.

Rasmus R. Paulsen and Joshua A. Levine
MeshMed Co-organizers

Organizing Committee

- Rasmus R. Paulsen. DTU Informatics, Technical University of Denmark. Denmark
- Joshua A. Levine. SCI Institute, University of Utah. USA
- Christian Barillot. CNRS / INRIA Research Unit, Rennes. France
- Nikos P. Chrisochoides. Old Dominion University. USA
- Hervé Delingette. INRIA - Team Asclepios, Sophia-Antipolis. France
- Ross T. Whitaker. SCI Institute, University of Utah. USA
- Yongjie Zhang. Carnegie Mellon University. USA

Program Committee

- Pierre Alliez. INRIA Sophia-Antipolis. France
- Jakob Andreas Bærentzen. Technical University of Denmark. Denmark
- Andrey Chernikov. Old Dominion University. USA
- Gary Christensen. University of Iowa. USA
- Tim Cootes. University of Manchester. UK
- Tron Darvann. University of Copenhagen. Denmark
- Tobias Heimann. German Cancer Research Center. Germany
- Luis Ibanez. Kitware. USA
- Michael Kazhdan. Johns Hopkins University. USA
- Steen Markvorsen. Technical University of Denmark. Denmark
- Mads Nielsen. University of Copenhagen. Denmark
- Sebastien Ourselin. University College London. UK
- Sylvain Prima. IRISA / INRIA Rennes. France
- Mauricio Reyes. University of Bern. Switzerland
- Daniel Rueckert. Imperial College London. UK
- Anuj Srivastava. Florida State University. USA
- Martin Styner. University of North Carolina at Chapel Hill. USA
- Guoliang Xu. Chinese Academy of Sciences. China
- Lilla Zöllei. Harvard Medical School. USA
- Lasse Riis Østergaard. Aalborg University. Denmark

Invited Speakers

Jonathan Shewchuk. UC Berkeley. USA Jonathan Shewchuk is an Associate Professor in the Department of Electrical Engineering and Computer Sciences at UC Berkeley. He is best known for his software Triangle for high-quality triangular mesh generation, which won the 2003 James Hardy Wilkinson Prize in Numerical Software, and his "Introduction to the Conjugate Gradient Method Without the Agonizing Pain."

Nicolas Smith. Kings College London. UK Nic Smith is Professor and Head of Department of Biomedical Engineering at Kings College London and is also Visiting Professor of Computational Physiology at the University of Oxford. His research is characterised by the development of integrated multi-scale and multi-physics models mainly of the heart, which provide the ability to link biophysically detailed experimental data to integrated function from sub-cellular to the whole organ level. Within the scope of this work he has developed computational techniques to enable specific model developments that have in turn been applied to provide insight into cardiac physiology. His research has been focused on cardiac electrophysiology and contraction at the cellular level, and the multi-scale translation of these models to enable coronary blood flow, cardiac electro-mechanics and coupled tissue mechanics-ventricular blood flow simulations at the organ level.

Nic is currently a central contributor to the Virtual Physiological Human (VPH) Project sponsored by the European Commission working to develop integrated multi-scale computational models of organ systems. He has taken a leading role in promoting the significance and optimization of the connection between data and model parameters. He is the scientific coordinator of euHeart an FP7 Integrated Project focused on developing multi-scale, multi-physics cardiac models which involves 18 academic, clinical and industrial partners across Europe. In addition to euHeart he holds grants from the EPSRC, BBSRC, MRC as Principle Investigator and is an EPSRC leadership fellow.

Michael Bronstein. University of Lugano. Switzerland Michael Bronstein is an assistant professor in the Institute of Computational Science, Faculty of Informatics, University of Lugano, Switzerland. Previously, he had a visiting appointment in the Department of Computer Science at Stanford university. Michael holds a B.Sc. from the Department of Electrical Engineering in 2002 and Ph.D. from the Department of Computer Science, Technion (Israel Institute of Technology) in 2007. His main research interests are theoretical and computational methods in metric geometry and their application to problems in computer vision, pattern recognition, shape analysis, computer graphics, image processing, and machine learning. He has authored over 70 publications in leading journals and conferences, over 20 patents, and the book "Numerical geometry of non-rigid shapes" (published by Springer Verlag). His research was recognized by numerous awards and prizes and was featured in CNN, SIAM News, Wired, and in the Abel lecture given in Oslo in honor of the 2009 Abel Prize laureate Mikhail Gromov. Michael Bronstein was the co-chair of the Workshop on Non-rigid shapes and deformable image alignment (NORDIA) in 2008-2011 and of the International Conference on n Scale Space and Variational Methods in Computer Vision (SSVM) in 2011, and has served on review and program committees of conferences in his field. In addition to academic activities, in 2004-2009 he served as a scientist and Vice President of video technology at the Silicon Valley start-up company Novafora Inc, leading a team of researchers and engineers developing Internet-scale computer vision and video analysis applications. Technology developed by Dr. Bronstein has been used in the foundation of several start-up companies, in which he is involved as advisor and co-founder.

Table of Contents

Session 1: Meshing

- Mesh-based vs. Image-based Statistical Appearance Model of the Human Femur: A Preliminary Comparison Study for the Creation of Finite Element Meshes. *Serena Bonaretti, Christof Seiler, Christelle Boichon, Philippe Büchler, and Mauricio Reyes.*
- High-Quality Multi-Tissue Mesh Generation for Finite Element Analysis. *Panagiotis Foteinos and Nikos Chrisochoides.*¹⁰

Session 2: Medical applications of surfacing and meshing

- MRI-free neuronavigation for transcranial magnetic stimulation in severe depression. *Benoît Combès, Charles Garraud, Xavier Morandi, Sylvain Prima, and Pierre Hellier.*
- Shape based Conditional Random Fields for Segmenting Intracranial Aneurysms. *Sajjad Baloch, Erkang Cheng, Ying Zhu, Ashraf Mohamed, Haibin Ling, and Tong Fang.*
- 3D Surface Realignment Tracking for Medical Imaging: A Phantom Study with PET Motion Correction. *Oline V. Olesen, Rasmus R. Paulsen, Rasmus R. Jensen, Sune H. Keller, Merence Sibomana, Liselotte Højgaard, Bjarne Roed, and Rasmus Larsen.*

Session 3: Surfaces

- Example based style classification. *Katarzyna Welnicka, Jakob Andreas Bærentzen, Henrik Aanæs, and Rasmus Larsen.*
- Template Deformation with User Constraints for Live 3D Interactive Surface Extraction. *Benoit Mory, Oudom Somphone, Raphael Prevost, and Roberto Ardon.*
- Automatic feature identification in dental meshes. *Yokesh Kumar, Ravi Janardan, Brent Larson, and Joe Moon.*

Poster Session

- 3D Shape Analysis of the Knee Extensor and Flexor Muscles in Patients with COPD using Mesh Projection-based Features. *Hengameh Mirzaalian, Ghassan Hamarneh, Bahareh HajGhanbari, and W. Darlene Reid.*
- Geometric calibration between PET scanner and structured light scanner. *Martin Kjer, Oline V. Olesen, Rasmus R. Paulsen, Liselotte Højgaard, Bjarne Roed, and Rasmus Larsen.*
- CLARCS, a C++ Library for Automated Registration and Comparison of Surfaces: Medical Applications. *Alexandre Abadie, Benoît Combès, Claire Haegelen, and Sylvain Prima.*
- Automatic Statistical Shape Analysis of Local Cerebral Asymmetry in 3D T1-weighted magnetic resonance images. *Antonietta Pepe, Lu Zhao, Jussi Tohka, Juha Koikkalainen, Jarmo Hietala, and Ulla Ruotsalainen.*
- High Quality Remeshing of Closed Genus-0 Surfaces. *Anton Semechko, Robert Dony, and Michele Oliver.*
- Development of Individualized Human Whole-body Thermoregulation Models Using Finite Element Analysis. *Tejash Patel, Gary P. Zientara, Xiaojiang Xu, and Reed W. Hoyt.*
- An Atlas-Based Geometry Pipeline for Cardiac Hermite Model Construction. *Yongjie Zhang, Xinghua Liang, Jun Ma, Yiming Jing, Matthew J. Gonzales, Adarsh Krishnamurthy, Paul Stark, Sanjiv M. Narayan, and Andrew D. McCulloch.*
- Ultra Fast Optical Sectioning: Signal preserving filtering and surface reconstruction. *Rasmus R. Jensen, Mike van der Poel, Rasmus Larsen, and Rasmus R. Paulsen.*
- Approach-guided controlled resolution brain meshing for FE-based interactive neurosurgery simulation. *Michel Audette, Denis Rivière, Matthew Ewend, and Sébastien Valette.*

Mesh-based vs. Image-based Statistical Appearance Model of the Human Femur: A Preliminary Comparison Study for the Creation of Finite Element Meshes

Serena Bonaretti¹, Christof Seiler^{1,2}, Christelle Boichon³, Philippe Büchler¹,
and Mauricio Reyes¹

¹ Institute for Surgical Technology and Biomechanics, University of Bern, Bern,
Switzerland

{serena.bonaretti,christof.seiler,philippe.buechler,mauricio.reyes}@istb.unibe.ch

² Asclepios Research Group, INRIA Sophia Antipolis, France

³ Ansys, Villeurbanne, France
christelle.boichon@ansys.com

Abstract. Statistical models have been recently introduced in computational orthopaedics to investigate the bone mechanical properties across several populations. A fundamental aspect for the construction of statistical models concerns the establishment of accurate anatomical correspondences among the objects of the training dataset. Various methods have been proposed to solve this problem such as mesh morphing or image registration algorithms. The objective of this study is to compare a mesh-based and an image-based statistical appearance model approaches for the creation of finite element (FE) meshes. A computer tomography (CT) dataset of 157 human left femurs was used for the comparison. For each approach, 30 finite element meshes were generated with the models. The quality of the obtained FE meshes was evaluated in terms of volume, size and shape of the elements. Results showed that the quality of the meshes obtained with the image-based approach was higher than the quality of the mesh-based approach. Future studies are required to evaluate the impact of this finding on the final mechanical simulations.

1 Introduction

Over the past years, statistical models have been introduced in computational orthopaedics as a useful investigation tool. They have been used for implant design evaluation [19,4], for fracture risk assessment [6,25] and for the discrimination of pathological from non pathological subjects [5]. The major advantage of statistical models consists in the ability of describing a dataset variations in terms of both shape and intensity. Moreover they allow creating new instances of the same dataset object, which belong to the same probability density function as the training dataset. This drives to the creation of a dense map of FE simulations and consequently a richer analysis of the bone mechanical behavior. These

characteristics have allowed overcoming one of the limitations of the studies in the field, which have been mainly performed on a small amount of data [16,13,23], limiting the generalization of these models to a complete population.

The definition of accurate correspondences is crucial for the creation of reliable statistical models [9]. For a correct modeling of both surface and volumetric information, correspondences have to be properly defined not only on the dataset object surface, but also inside the object volumes. Two different approaches exist to establish the anatomical correspondences: mesh morphing techniques and image registration techniques. In the first case, the dataset objects are represented as meshes and a chosen volume reference mesh is morphed to the other dataset [6,14]. In the second case, the dataset objects are represented by volume images and they are registered to the one chosen as reference [20,27].

Both approaches have strong features and weaknesses. The mesh-based approach provides an output that is directly compatible with finite element (FE) calculations. However, the smoothing required for the creation of well-shaped elements can penalize the accuracy of the anatomical correspondences. On the other hand, the image-based approach has no constraints related to node positioning since the FE mesh is created after the instantiation of synthetic new images. Therefore the accuracy of the correspondence establishment is less problematic for FE simulations. However, due to the nature of these techniques, the correspondence establishment at the interface background-object is more sensitive to errors as compared to pure surface-based registration techniques. Volume meshes are created separately for each instance, implying the impossibility of comparing simulation results on a element-wise basis. Finally deformation vector fields have to be invertible for the creation of new instances. To the authors' knowledge there is no study that compares the two different approaches used by the community in order to determine which is the most suitable for the creation of FE meshes.

The purpose of the paper is therefore to compare two existing methodologies which aim to create FE meshes from statistical models of appearance [10]. A comprehensive evaluation of surface or volumetric-based registration techniques is beyond the scope of this study. In the first part of this paper, we present the creation of two different pipelines, namely mesh-based and image-based. For each pipeline, starting from the same dataset, bone anatomical correspondences are determined, a statistical appearance model is built and new instances, which can be directly used for FE simulations, are created. In a second step, the quality of the FE meshes obtained with the two pipelines is compared based on common objective metrics.

2 Mesh-based and Image-based Pipelines for the Creation of FE Meshes

Two statistical appearance model approaches for the creation of finite element meshes are presented. The two approaches, namely mesh-based and image-based, had overlapping steps, as shown in Fig 1. The main difference between the two

approaches consisted in a different representation of the dataset objects. In the mesh-based approach, the bone objects were considered as meshes from the first step of the pipeline, whereas in the image-based approach they were represented by volumetric images for the whole pipeline. In this case, the images were transformed to FE meshes in the very final step. For completeness and reproducibility purpose, details about data and implementation can be found in Appendix A.

Segmented CT images were the input for both pipelines. Mesh morphing and image registration were used to find anatomical correspondences, respectively for the mesh-based and the image-based approaches. For both the mesh-based and the image-based pipelines, the same principals were used to calculate the statistical appearance models and to create new instances.

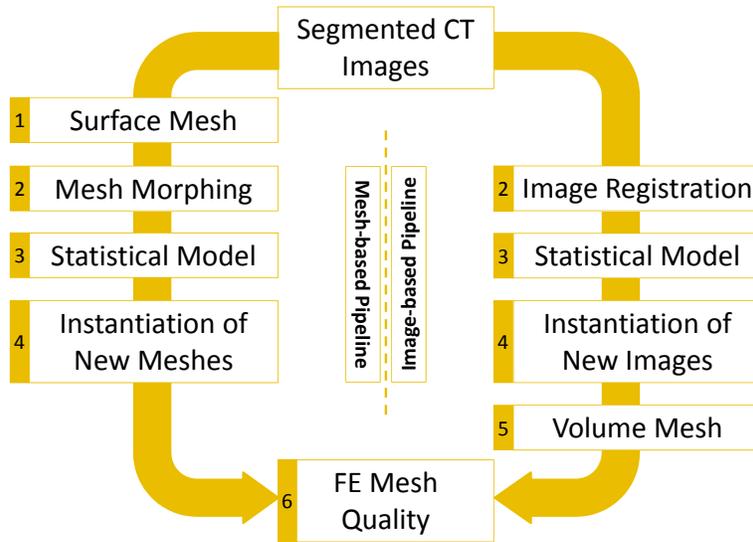


Fig. 1: Mesh-based and image-based statistical appearance model pipelines for the creation of femur finite element meshes. On the left side, the steps for the mesh-based approach are shown, whereas on the right side the steps for the image-based pipeline are depicted.

1. Surface Mesh Creation. In the mesh-based pipeline, for each bone of the dataset a surface mesh was created following the steps shown in Fig. 2. From each segmented CT image (Fig. 2a), a surface mesh was created using the marching cube algorithm [17] (Fig. 2b). The obtained mesh resulted too dense and rough and could not be directly used for the following processing. Therefore, in order to improve the mesh quality, each mesh was decimated and smoothed using the Laplacian operator. Since in many cases the node removal

caused the loss of mesh topology, node connections were reestablished using MRFSurface [18] (Fig. 2c).

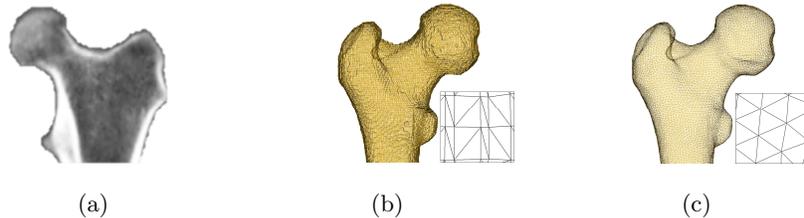


Fig. 2: Creation of femur surface mesh from CT image. (a) Sagittal view of a segmented femur head CT image. (b) Surface mesh obtained from the marching cube algorithm. (c) Surface mesh topology reconstruction.

2. Finding Correspondences: Mesh Morphing and Image Registration. In order to find anatomical correspondences among the dataset objects, mesh morphing was used for the mesh-based pipeline, whereas image registration was performed for the image-based approach. In both cases, correspondences were calculated with respect to the same reference bone. The choice of the reference bone was done in the image-based pipeline using an iterative procedure. In the first step, one bone of the dataset was randomly picked as the reference and all the remaining bones of the dataset were registered to it. The average transformation was then calculated and the bone whose transformation was the closest to the average transformation was considered as the new reference femur. These steps were repeated until convergence.

The mesh morphing was performed using the algorithm developed by [15] and extended by [3]. First, a volumetric mesh was created for the reference bone. From the surface mesh, the volume mesh was created using NETGEN [24]. The reference volumetric mesh was composed of 130000 quadratic tetrahedrons, in order to satisfy the necessary level of mesh refinement for femur FE simulations [28]. Then, the reference bone volumetric mesh was warped to all dataset surface meshes to create iso-topological tetrahedral FE meshes. To compute the mesh morphing, 10 landmarks were manually selected on the surface of both the reference volume mesh and the bone dataset surfaces meshes, as shown in Fig. 3. Four landmarks were selected in correspondence to the main anatomical features of the femur head (Fig. 3a) and 6 at easily detectable points in the condyle area (Fig. 3b). The landmarks were used as constraints during the morphing computations. The mesh morphing was executed first on the surface mesh and then on the volume mesh. From the reference volume mesh, the surface mesh was extracted. For each bone, both the reference surface mesh and the current bone surface mesh were projected on a disc of unitary radius. The position of the bone landmarks with respect to the reference bone landmarks was used as constraint for the moving of all bone surface nodes. The new position of the femur surface

nodes was then calculated using radial basis functions (RBF), where the center of the functions was represented by the landmarks. The bone surface nodes were then reprojected back from the parametric space to the physical space, and their positions smoothed. As second step, the reference volumetric mesh was morphed to the current bone new surface mesh. The position of the volumetric mesh inner nodes was calculated using RBF. In this case, the RBF center was represented by the just computed surface nodes.

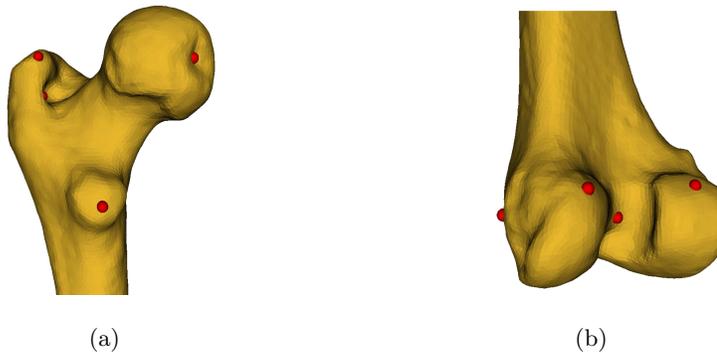


Fig. 3: Landmarks selection for mesh morphing. (a) Selection of four landmarks in the femur head. (b) Selection of six landmarks in the femur condyles (two landmarks not visible).

In the image-based pipeline, correspondences were detected using image registration. We used the Log-Domain Demons (LDD) registration algorithm regularized with a femur-specific polyaffine model [26]. The LDD finds Deformation Vector Fields (DVF) which are smooth and invertible [27]. This characteristic enables the creation of new instances, as explained in Section 2.4. In the LDD, DVF are generated through the exponential of Stationary Velocity Fields (SVF), which are the primary results of the registration process [2]. The LDD iteratively minimizes the energy functional composed of two terms, namely the ‘correspondence’ term, which calculates the SVF and the ‘regularization’ term, which imposes smoothness on the SVF [7]. In our case, the regularization term was replaced by a polyaffine model, which allowed us to capture the main anatomical variation of the femur. In the model, three regions were defined: femur head, shaft and condyle area. The parameters of the polyaffine model were jointly estimated using a closed form least square solution during each iteration step of the LDD.

3. Statistical Models. The statistical appearance model was created as a combination of the parameters obtained from the statistical shape model and the statistical intensity model [10].

In the mesh-based approach, the *statistical shape model* was computed on the volume mesh coordinates. In order to calculate the principal component

analysis (PCA), the volume meshes were aligned to the average mesh using the Procrustes method [11]. The model was then built on the aligned meshes [9]. In the image-based approach, instead, the statistical shape model was computed as a statistical deformation model [21]. Therefore PCA was directly calculated on the SVF obtained from the image registration process.

For the mesh-based approach, the *statistical intensity model* was created on the image intensity values at the corresponding mesh node positions. For the creation of valid FE meshes, the assignment of correct grey levels to the mesh nodes is crucial. In fact CT information and bone mechanical properties are strictly linked [22]: CT image grey levels represent the bone mineral density [8], which is related to the Young’s modulus through empirical relationships [12]. One critical aspect affecting the Young’s modulus assignment is the partial volume effect that occurs on the bone surface [13]. Therefore the bone outer layer was first eroded to delete the partial volume effect area and then dilated to rebuild the canceled outer cortical bone. At each node of the mesh, intensities were calculated as a 26-voxel connectivity linear interpolation. The computed intensities were finally used for the PCA calculation. On the other side, for the image-based approach, the statistical intensity model was created through the calculation of the PCA on the original images warped to the reference bone in order to have spatial correspondence among the gray levels of the dataset images.

For both the mesh-based and the image-based approaches, the *statistical appearance model* was built on the combination of shape and intensity parameters, as shown in Eq. 1.

$$b = \begin{pmatrix} W_s b_s \\ b_g \end{pmatrix} = \begin{pmatrix} W_s \Phi_s^T (x - \bar{x}) \\ \Phi_g^T (g - \bar{g}) \end{pmatrix}, \quad (1)$$

where b represents the combined parameters, composed by the shape parameters b_s and the intensity parameters b_g ; Φ_s^T and Φ_g^T represent respectively the transposition of the modes calculated from the statistical shape model and the statistical intensity model; x and g are the initial mesh coordinate and intensity dataset; \bar{x} and \bar{g} are the average shape and intensity, respectively. Since the shape and the intensity parameters were represented by different units, the shape parameters were multiplied by the matrix W_s in order to make the parameters homogeneous. W_s was calculated as $W_s = rI$, where r is the square root of the ratio between the total variation obtained from the statistical intensity model and the total variation calculated from the statistical shape model, and I is the identity matrix.

4. Instantiation of New Samples. For both the mesh-based and the image-based pipelines, new instances were created from the statistical appearance model. The new shape \tilde{x} was created using Eq. 2 and the corresponding intensity distribution \tilde{g} using Eq. 3:

$$\tilde{x} = \bar{x} + \Phi_s W_s^{-1} \Phi_{c,s} c \quad (2)$$

$$\tilde{g} = \bar{g} + \Phi_g \Phi_{c,g} c, \quad (3)$$

where

$$\Phi_c = \begin{pmatrix} \Phi_{c,s} \\ \Phi_{c,g} \end{pmatrix}, \quad (4)$$

represents the combined eigenvectors derived from the calculation of PCA on the combined parameters b , divided in its shape component $\Phi_{c,s}$ and intensity component $\Phi_{c,g}$. The parameter c was calculated as $-2\sqrt{\lambda_i} \leq c \leq +2\sqrt{\lambda_i}$, where λ_i is the current eigenvalue, whereas the weight between -2 and $+2$ was calculated using the latin hypercube sampling method.

In the mesh-based approach, the new shape was built assigning to the reference topology the new node coordinates. The calculated gray levels were then associated to the mesh nodes. In the image-based pipeline, the new image intensity were calculated in the reference shape. The obtained image was then warped to the new shape calculated, thanks to the invertibility of the Demons DVF.

5. Volume Mesh Creation. At the end of the image-based approach, for each new instance, a FE mesh was created. Similar to the processing performed in steps 1 and 2, from each new image, a surface mesh was created using the marching cubes algorithm. To ensure good mesh quality, the mesh was then simplified and smoothed, and its topology reconstructed using MRFSurface. Finally from the obtained surface mesh, a tetrahedral quadratic FE mesh of about 130000 elements was created using NETGEN.

6. FE Mesh Quality Assessment. Three different criteria were used to evaluate the quality of the finite element meshes. We evaluated mesh elements in terms of volume, size and shape, using the following metrics:

- a Jacobian. It is a volume metric which describes the distortion of the element from the ideal shape. At its extremes, $+1$ and -1 , the element shape was considered perfect and distorted, respectively; at 0 the element had null volume.
- b Edge ratio. It is a size metric which is calculated as the ratio between the longest and the shortest edge of the element.
- c Minimum angle. It is a shape metrics that evaluates the smallest angle of the element sides.

3 Experiments on Femur CTs

To compare the two pipelines, we conducted calculations on left femur bone CT images. In the following paragraphs we show the results that we obtained for the steps of the pipeline. The numbering of the steps refers to Fig 1.

Finding Correspondences: Mesh Morphing and Image Registration (step 2). A total of 196 segmented CT images were used in this study. According to a visual evaluation, mesh morphing failed for 25 meshes, whereas image registration failed for 24 images. Both mesh morphing and image registration

succeeded in 157 cases, which were considered as input dataset for the statistical appearance models.

Statistical Models (step 3). The results of the statistical appearance models are shown in Table 1. The variations of the dataset were described in a more compact way by the model obtained with the mesh-based approach than with the image-based approach (Table 1a). The evaluation of the computational time showed that the mesh-based approach was less time and memory consuming than the image-based one, even on a less powerful machine (Table 1b). In the mesh-based pipeline, computations were done in a few minutes, whereas hours of calculations were required in the image-based pipeline.

Table 1: Statistical appearance model computation results. (a) Model compactness for the mesh-based and the image-based pipelines. (b) Computational costs for the calculations.

(a)		
model compactness	mesh-based pipeline (# modes)	image-based pipeline (# modes)
50%	2	1
75%	6	6
80%	10	9
90%	40	40
95%	75	85
100%	157	157

(b)		
	mesh-based pipeline	image-based pipeline
Shape Model	5 min ¹	5 hrs ²
Intensity Model	2 min ¹	1.5 hrs ²
Appearance Model	9 min ¹	6.5 hrs ¹

¹ Processor: Intel Core Duo, E8500 @ 3.16GHz. RAM: 8GB

² Processor: Intel Xeon CPU, X5550 @ 2.67GHz. RAM: 48GB

Instantiation of New Samples (step 4). A total of 30 new instances were created for each pipeline. In order to create samples that represented the 90% of variation of the population, 40 modes were used for the mesh-based approaches, whereas 86 modes were used for the image-based approach.

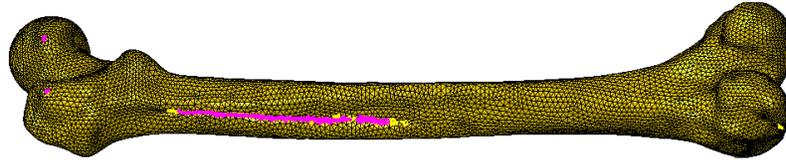
FE Mesh Quality Assessment (step 6). The evaluation of the mesh element quality resulted as follows:

- a Jacobian. Two of the 30 meshes created with the mesh-based pipeline had respectively 1 and 64 elements with a zero or negative Jacobian (Fig. 4a). None of the 30 volume meshes created with the image-based pipeline had distorted elements.
- b Edge ratio. For the meshes created with the mesh-based pipeline, both the average and the standard deviation of the edge ratio were greater and spread in a wider range than for the image-based approach (Fig. 4b). One mesh had 2 elements whose edge ratio was greater than 10 [1], therefore the mesh could not be used for FE calculation. For the image-based pipeline, meshes had a smaller edge ratio with less variations.
- c Minimum angle. As for the previous metric, for the meshes created with the mesh-based pipeline, the minimum angle results were more spread than for the meshes created with the image-based approach (Fig. 4c). Moreover, the meshes created with the mesh-based approach had a smaller minimum angle with larger variations. Six meshes had from 1 to 37 elements whose minimum angle was less than 10 degrees [1], compromising the meshes usability in mechanical simulations.

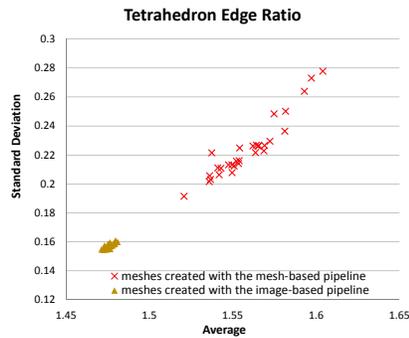
4 Conclusion

In this study we presented a preliminary comparison between two existing different statistical appearance model pipelines for the creation of FE meshes, namely a mesh-based and an image-based approaches. The mesh-based pipeline allowed the calculation of a more compact statistical model of appearance and the direct creation of iso-topological meshes as output of the statistical model. On the other hand, in the image-based approach the correspondence among mesh elements is lost since each mesh is created separately. Moreover the image-based approach is computational expensive and has the issue of the inversion of the DVF for new instances creation. In both pipelines, the ability of the models at describing the training dataset variation resulted to be similar. However the quality of the mesh tetrahedrons created with the image-based pipeline was higher. Image-based pipeline meshes performed better in terms of element distortion, size and shape.

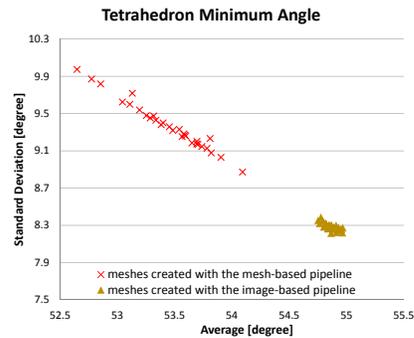
Future work is required to merge the strengths of both pipelines, to compare the mechanical properties assignment in the two approaches and to evaluate the implication for finite element simulations.



(a)



(b)



(c)

Fig. 4: FE mesh quality assessment. (a) Mesh with distorted elements (tetrahedrons with zero or negative Jacobian are coloured in pink). (b) Edge ratio evaluation. For each mesh the average of the edge ratio of its elements vs. the standard deviation is plotted. The image-based pipeline meshes showed a better edge ratio. (c) For each mesh, the average of the minimum angle of each element vs. the standard deviation is plotted. The meshes created with the image-based pipeline showed higher minimum angle.

5 Appendix A

Supplementary data associated with this study can be found at <https://sites.google.com/site/serenabonaretti/>

References

1. Abaqus: ABAQUS/CAE user's manual : version 6.9.
2. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A Log-Euclidean Framework for Statistics on Diffeomorphisms. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI'06). vol. 9, pp. 924–31 (2006)
3. Boichon, C., Rochette, M., Schileo, E., Grassi, L., Taddei, F., Viceconti, M.: Shape Indexing of Human Femora Using Morphing and Principal Component Analysis. In: 1st Virtual Physiological Human Conference - VPH2010. No. 1 (2001)

4. Bonaretti, S., Reimers, N., Reyes, M., Nikitsin, A., Joensson, A., Nolte, L., Büchler, P.: Assessment of Peri-Articular Implant Fitting Based on Statistical Finite Element Modeling. In: MICCAI Workshop on Computational Biomechanics for Medicine III (2008)
5. Bredbenner, T.L., Eliason, T.D., Potter, R.S., Mason, R.L., Havill, L.M., Nicoletta, D.P.: Statistical Shape Modeling Describes Variation in Tibia and Femur Surface Geometry between Control and Incidence Groups from the Osteoarthritis Initiative Database. *Journal of Biomechanics* 43(9), 1780–6 (2010)
6. Bryan, R., Mohan, P.S., Hopkins, A., Galloway, F., Taylor, M., Nair, P.B.: Statistical Modelling of the Whole Human Femur Incorporating Geometric and Material Properties. *Medical Engineering & Physics* 32(1), 57–65 (2010)
7. Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N.: Iconic Feature Based Nonrigid Registration: the PASHA Algorithm. *Computer Vision and Image Understanding* 89, 272–298 (2003)
8. Ciarelli, M.J., Goldstein, S.A., Kuhn, J.L., Cody, D.D., Brown, M.B.: Evaluation of Orthogonal Mechanical Properties and Density of Human Trabecular Bone from the Major Metaphyseal Regions with Materials Testing and Computed Tomography. *Journal of Orthopaedic Research* 9(5), 674–82 (1991)
9. Cootes, T., Taylor, C., Cooper, D., Graham, J.: Active Shape Models - Their Training and Application. *Computer Vision and Image Understanding* 61(1), 38–59 (1995)
10. Cootes, T.F., Taylor, C.J.: Statistical Models of Appearance for Medical Image Analysis and Computer Vision. In: *SPIE Medical Imaging*. pp. 236–248. Spie (2001)
11. Goodall, C.: Procrustes Methods in the Statistical Analysis of Shape. *Journal of the Royal Statistic society. Series B (Methodological)* 53(2), 285–339 (1991)
12. Helgason, B., Perilli, E., Schileo, E., Taddei, F., Brynjólfsson, S., Viceconti, M.: Mathematical Relationships between Bone Density and Mechanical Properties: A Literature Review. *Clinical Biomechanics* 23(2), 135–46 (2008)
13. Helgason, B., Taddei, F., Pálsson, H., Schileo, E., Cristofolini, L., Viceconti, M., Brynjólfsson, S.: A Modified Method for Assigning Material Properties to FE Models of Bones. *Medical Engineering & Physics* 30(4), 444–53 (2008)
14. Hraiech, N.: Mesh Morphing and Shape Indexing in Human Femur Modeling Applications. Ph.D. thesis (2010)
15. Hraiech, N., Taddei, F., Malvesin, E., Rochette, M., Viceconti, M.: Fast 3D Mesh Generation of Femur Based on Planar Parameterization and Morphing. In: *5th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. pp. 1561–1564. No. 3 (2008)
16. Keyak, J., Falkinstein, Y.: Comparison of In Situ and In Vitro CT Scan-based Finite Element Model Predictions of Proximal Femoral Fracture Load. *Medical Engineering & Physics* 25(9), 781–787 (2003)
17. Lorensen, W.E., Cline, H.E.: Marching Cubes: A High Resolution 3D Surface Construction Algorithm. *Computer Graphics* 21(4), 163–169 (1987)
18. Paulsen, R.R., Baerentzen, J.A., Larsen, R.: Markov Random Field Surface Reconstruction. *IEEE Transactions on Visualization and Computer Graphics* 16(4), 636–46 (2010)
19. Querol, L., Büchler, P., Rueckert, D., Nolte, L.: Statistical Finite Element Model for Bone Shape and Biomechanical Properties. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI'06)*. pp. 405 – 411 (2006)
20. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid Registration Using Free-Form Deformations: Application to Breast MR Images. *IEEE Transactions on Medical Imaging* 18(8), 712–21 (1999)

21. Rueckert, D., Frangi, A.F., Schnabel, J.a.: Automatic Construction of 3-D Statistical Deformation Models of the Brain Using Nonrigid Registration. *IEEE Transactions on Medical Imaging* 22(8), 1014–25 (2003)
22. Schileo, E., Dall’Ara, E., Taddei, F., Malandrino, A., Schotkamp, T., Baleani, M., Viceconti, M.: An Accurate Estimation of Bone Density Improves the Accuracy of Subject-Specific Finite Element Models. *Journal of Biomechanics* 41(11), 2483–91 (2008)
23. Schileo, E., Taddei, F., Cristofolini, L., Viceconti, M.: Subject-Specific Finite Element Models Implementing a Maximum Principal Strain Criterion Are Able to Estimate Failure Risk and Fracture Location on Human Femurs Tested In Vitro. *Journal of Biomechanics* 41(2), 356–67 (2008)
24. Schöberl, J.: NETGEN An Advancing Front 2D/3D-Mesh Generator Based on Abstract Rules. *Computing and Visualization in Science* 1, 41–52 (1997)
25. Schuler, B., Fritscher, K.D., Kuhn, V., Eckstein, F., Link, T.M., Schubert, R.: Assessment of the Individual Fracture Risk of the Proximal Femur by Using Statistical Appearance Models. *Medical Physics* 37(6), 2560–2571 (2010)
26. Seiler, C., Pennec, X., Ritacco, L., Reyes, M.: Femur Specific Polyaffine Model to Regularize the Log-Domain Demons Registration. In: *In SPIE Medical Imaging*. vol. 7962 (2011)
27. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric Log-Domain Diffeomorphic Registration: A Demons-based Approach. *Medical Image Computing and Computer-Assisted Intervention (MICCAI’08)* 11, 754–61 (2008)
28. Viceconti, M., Davinelli, M., Taddei, F., Cappello, A.: Automatic Generation of Accurate Subject-Specific Bone Finite Element Models to Be Used in Clinical Studies. *Journal of Biomechanics* 37(10), 1597–605 (2004)

High-Quality Multi-Tissue Mesh Generation for Finite Element Analysis

Panagiotis Foteinos^{1,2} and Nikos Chrisochoides²

¹ Computer Science Department, College of William and Mary, Virginia, USA,
pfot@cs.wm.edu

² Computer Science Department, Old Dominion University, Virginia, USA,
nikos@cs.odu.edu

Abstract. Mesh generation on 3D segmented images is a fundamental step for the construction of realistic biomechanical models. Mesh elements with low or large dihedral angles are undesirable, since they are known to underpin the speed and accuracy of the subsequent finite element analysis. In this paper, we present an algorithm for meshing 3D multi-label images. A notable feature of our method is its ability to produce tetrahedra with very good dihedral angles respecting, at the same time, the interfaces created by two or more adjoining tissues. Our method employs a Delaunay refinement scheme orchestrated by special point rejection strategies which remove poorly shaped elements without deteriorating the representation of the objects' anatomical boundaries. Experimental evaluation on CT and MRI atlases have shown that our algorithm produces watertight meshes consisting of elements of very good quality (all the dihedral angles were between 19 and 150 degrees) which makes our method suitable for finite element simulations.

1 Introduction

Meshing multi-labelled medical images (like those obtained by segmenting MRI or CT images) provides the means for constructing accurate bio-mechanical models for subsequent finite element analysis. Multi-material mesh generation imposes challenges, since it should meet two conflicting requirements: *fidelity* and *quality*.

Fidelity measures the capability of the mesher to preserve the boundaries formed by two or more adjoining tissues. Quality regards the shape of the elements: tetrahedra with small or large dihedral angles (i.e., low quality tetrahedra) result in interpolation errors and in ill-conditioned stiffness matrices undermining in this way the accuracy and speed of the associated finite element analysis [1].

The difficulty in mesh generation is that the need to preserve high-curvature creases of the object's surface (i.e., high fidelity) deteriorates the quality of the meshes; on the other hand, the quality of mesh elements should be as high as possible when dealing with isotropic materials [2].

In this paper, we propose a Delaunay meshing algorithm able to respect the interfaces of multi-material domains and produce tetrahedra with very good dihedral angles and radius-edge ratios (and therefore very good aspect ratios), offering at the same time control over the size of the mesh.

1.1 Previous work

In the literature, there has been work on multi-tissue meshing but the issue of high quality has not been adequately addressed.

Meyer *et al.* [3] employ a particle-based scheme producing watertight meshes that respect the interfaces formed by adjoining tissues. However, elements with practically zero dihedral angles (slivers) do appear in the final meshes. Furthermore, the execution times reported range from 3 to 12 hours even for small datasets.

Liu *et al.* [4] compress a body-centered cubic lattice (BCC) using a point-based registration method. The dihedral angles, however, can be as low as 4° . Also, the uniform lattice results in an unnecessary large number of elements in the interior of the objects.

Chentanez *et al.* [5] model the insertion of needles into soft tissues. The resulting conforming meshes are observed to consist of elements of angles more than 10.3° and less than 160° . It is worth noting that their goal is to represent a 1-dimensional curvilinear object (the needle) as a subset of a single-tissue mesh, which is a goal quite different from ours.

Goksel and Salcudean [6] present a variational meshing technique which combines both meshing and segmentation. They report minimum angles as large as 20° . The synthetic data they used for the evaluation is a sphere, that is, a 2-manifold. Usually, multi-tissue domains consists of complicated geometries, i.e., non-manifold parts which intersect with more than one tissues. These domains impose challenges to any meshing technique and are the focus of this work.

Zhang *et al.* [7] develop an octree-based meshing algorithm. Although edge-contraction and smoothing schemes are employed for quality improvement, the authors do not report the dihedral angles observed in their meshes.

Hu *et al.* [8] and Hartmann and Kruggel [9] develop uniform meshes for multi-material domains achieving dihedral angles more than 10° , albeit without good fidelity: their meshes suffer from the “staircase” effect.

Based on previous work on single material Delaunay surface [10] and volume meshing [11], Pons *et al.* [12] present a meshing algorithm for multi-tissue domains. Recently, Boltcheva *et al.* [13] extend the work of Pons *et al.* [12], so that 0- and 1-junctions are preserved in the final meshes. Both these methods apply *sliver exudation* [14] in order to improve the quality of the mesh. Edelsbrunner and Guoy [15], however, have shown that in most cases sliver exudation does not remove all poor tetrahedra: elements with dihedral angles less than 5° survive. Indeed, Pons *et al.* [12] and Boltcheva *et al.* [13] report dihedral angles as low as 4° .

1.2 Our contribution

In this paper, we present a Delaunay refinement algorithm for meshing multi-tissue medical data so that the boundaries between neighboring tissues are conforming. It works directly on segmented images meshing both the surface and the volume of the tissues.

A notable feature of our method is its ability to produce tetrahedra with very good dihedral angles: in all the experiments on synthetic and real images we ran, our algorithm produces watertight meshes consisting of tetrahedra with dihedral angles larger than 19° and smaller than 150° .

The technique we employ for quality improvement is inspired by the work of Shewchuk [16]. Therein, poor tetrahedra are eliminated by inserting the center of their circumball, giving priority to tetrahedra with larger radius-edge ratio. Shewchuk, however, meshes input domains bounded by polyhedral surfaces. In this paper, we extend this technique to deal with multi-tissue domains bounded by curved surfaces. The main difficulty is that vertices near the surface might be inserted during quality improvement. This fact in turn hurts fidelity: edges that cross interfaces or holes appear. To overcome this problem, we propose special *point rejection strategies*. They improve the quality of elements preventing the insertion of points near the surface; rather, carefully chosen points are inserted precisely on the surface. This allows to achieve both good quality and good fidelity meshes.

The rest of the paper is organized as follows: Section 2 outlines the concept of Delaunay refinement. The multi-tissue capability of our algorithm and the point rejection strategies are described in Section 3. Lastly, Section 4 presents results on CT and MR multi-label images and Section 5 concludes the paper.

2 Background

Delaunay meshes have been shown to successfully approximate the surface of both manifold and non-manifold surfaces [10], due to the properties of the *restricted Delaunay triangulations*, first introduced by Amenta and Bern [17].

Let $V \subset \mathcal{R}^3$ be a set of vertices and $\mathcal{D}(V)$ their Delaunay triangulation. Any Delaunay triangulation satisfies the *empty ball* property: the circumscribing open ball (a.k.a *circumball*) of each tetrahedron in $\mathcal{D}(V)$ does not contain any vertex.

The *voronoi point* of a tetrahedron $t \in \mathcal{D}(V)$ is defined as the center (a.k.a *circumcenter*) of t 's circumball. The *voronoi edge* of a triangle $f \in \mathcal{D}(V)$ is the segment containing those points of \mathcal{R}^3 such that (a) they are equidistant from f 's vertices and (b) they are closer to f 's vertices than to any other vertex in V .

Let \mathcal{O} be the multi-label domain to be meshed. We denote \mathcal{O} 's surface with $\partial\mathcal{O}$. The *restriction* of $\mathcal{D}(V)$ to \mathcal{O} (denoted with $\mathcal{D}_{|\mathcal{O}}(V)$) is defined as the set of tetrahedra in the triangulation whose circumcenter lies inside \mathcal{O} .

It can be shown [10, 11] that if V samples $\partial\mathcal{O}$ sufficiently densely, then the set of boundary triangles (a.k.a *restricted facets*) of $\mathcal{D}_{|\mathcal{O}}(V)$ is a good approximation of $\partial\mathcal{O}$ in a both topological and geometric sense. The approximation guarantees

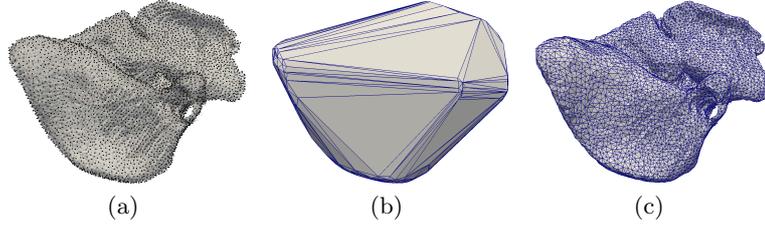


Fig. 1. (a) Sample set V of a liver's surface. (b) The Delaunay triangulation $\mathcal{D}(V)$ of the samples. (c) The restricted triangulation $\mathcal{D}|_{\mathcal{O}}(V)$.

hold as long as $\partial\mathcal{O}$ does not have sharp corners. This is a reasonable assumption, since biological tissues do not exhibit sharp features on their surface. See Figure 1 for a single-tissue example. The same idea extends to more than one tissues as well.

As an interesting consequence of the way $\mathcal{D}|_{\mathcal{O}}(V)$ is defined, only the voronoi edges of the restricted facets intersect the surface $\partial\mathcal{O}$, a property that we will exploit in Section 3 to improve quality.

3 Our method

The input of our algorithm is an image \mathcal{I} containing the multi-material object \mathcal{O} . Image \mathcal{I} can be seen as a function $f : \mathcal{R}^3 \mapsto \{0, 1, 2, \dots, n\}$, such that $f(p)$ is the label that point $p \in \mathcal{R}^3$ belongs to. More precisely, $f(p)$ is the label of the voxel that p lies in. Usually, a label of 0 denotes voxels outside \mathcal{O} .

Points on the surface $\partial\mathcal{O}$ of object \mathcal{O} are classified as those points lying in a voxel of label i which is incident to at least one other voxel of label j , such that $i < j$. In this way, surface $\partial\mathcal{O}$ contains not only the portions of the image that separate \mathcal{O} from the background, but it also contains the interfaces that separate any adjoining tissues. The goal is to recover $\partial\mathcal{O}$ and mesh the volume (induced by $\partial\mathcal{O}$) at the same time.

Our algorithm first creates a box by inserting its 8 corners. The box contains \mathcal{O} such that the (shortest) distance between the box and $\partial\mathcal{O}$ is larger than $2\delta\sqrt{2}$. Parameter δ is the only parameter that the users have to specify. This parameter determines how densely $\partial\mathcal{O}$ will be sampled: lower values indicate a denser sampling which in turn implies a better surface approximation. Notice that the calculation of the corners of the box is a quite trivial task, since it requires just one image traversal.

Next, the Delaunay triangulation of these corners is computed. This triangulation is the initial mesh (consisting of 12 tetrahedra) where the actual refinement starts from.

The refinement is governed by 2 steps, namely, *mesh conformity* and *point rejection quality improvement*. Upon termination, the tetrahedra whose circum-

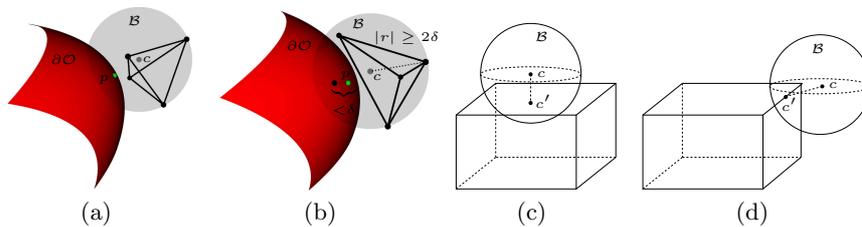


Fig. 2. (a) The closest surface point p to circumcenter c is inserted, (b) c is inserted but p is not, (c)-(d) c does not lie inside the box and therefore it is not inserted. Its projection c' is inserted and the vertices closer than δ to c' are deleted from the mesh.

center belongs to label i constitute the mesh representing the i^{th} tissue. Below, we outline each step separately.

3.1 Mesh conformity

As noted in Section 2, vertices on $\partial\mathcal{O}$ have to be inserted in order for the mesh boundary (i.e., triangles incident to 2 or more tetrahedra of different labels) to be a good approximation of $\partial\mathcal{O}$. For this reason, we keep track of the tetrahedra whose circumball \mathcal{B} intersects the surface $\partial\mathcal{O}$. We call such elements *intersecting tetrahedra*.

Suppose that an intersecting tetrahedron t is found. We compute the closest surface point —say p — to the center c of t 's circumball \mathcal{B} . To facilitate the computation of such a point, we make use of an image euclidean distance transformation [18]. If p is not closer than δ to any other surface vertex (already inserted in the mesh), then p is inserted (see Figure 2(a)). Otherwise, and if the radius of \mathcal{B} is larger than 2δ , c is inserted instead (see Figure 2(b)). In this way, we can show that this step does not cause the insertion of infinite number of vertices and therefore, termination is not compromised.

For the same reason, we also require that no vertex is ever inserted outside the box. When the circumcenter c of an intersecting tetrahedron is chosen for insertion, however, c might lie outside the box. To prevent such cases, c is rejected and its projection on the box is inserted instead. See Figure 2(c) and Figure 2(d) for a couple of examples.

At the end of this step, all the vertices that do not lie on $\partial\mathcal{O}$ are deleted from the triangulation. At this moment, the restricted facets of the mesh are a good approximation of $\partial\mathcal{O}$, because the vertices remained in the triangulation form a dense sample of $\partial\mathcal{O}$ (see Section 2). Also, we can show that no 2 vertices are closer than δ and this is why δ controls the size of the mesh.

3.2 Point rejection quality improvement

Our algorithm keeps track of poor tetrahedra, i.e., tetrahedra with small or large dihedral angles. Poor tetrahedra are eliminated by inserting their circumcenter.

Priority is given to the tetrahedra with higher radius-edge ratio as in [16]. The radius-edge ratio of a tetrahedron t is defined as the length of t 's circumball radius divided by the length of t 's shortest edge.

Problems arise, however, when the circumcenter of a poor tetrahedron (about to be eliminated) lies close to the surface. If this is the case, the restricted facets in the triangulation are not any more a good approximation of $\partial\mathcal{O}$. See Figure 3(a) for an example: the boundary facets have vertices that do not lie precisely on the surface.

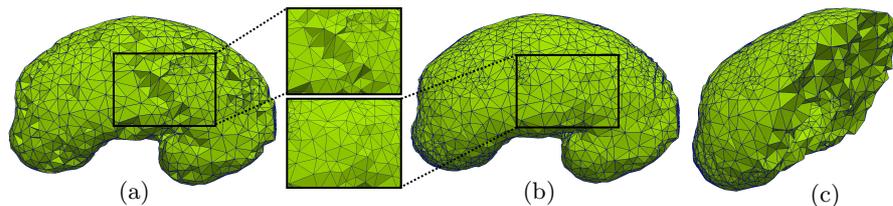


Fig. 3. Meshes for a kidney. All dihedral angles are between 19° and 150° . (a) No extra care has been taken to preserve fidelity and holes appear. (b) The point rejection strategies prevented the creation of holes. (c) A cross section of the mesh in (b).

To overcome this issue, we propose special *point rejection strategies*. Their goal is to make sure that all poor tetrahedra are eliminated without inserting points close to the surface.

Our algorithm first tries to convert *illegal* facets to *legal* ones. We define legal facets to be those restricted facets whose three vertices lie precisely on $\partial\mathcal{O}$. Conversely, a restricted facet with at least one vertex not lying on $\partial\mathcal{O}$ is called an illegal facet.

Let t be an illegal facet and e its voronoi edge (see Figure 4(a) for an illustration). Recall that e has to intersect $\partial\mathcal{O}$ (see Section 2) at a point p . Any vertex v of t which does not lie precisely on $\partial\mathcal{O}$ is deleted from the triangulation, while point p is inserted. Note that since only non-surface vertices are deleted from the triangulation and since p is inserted on $\partial\mathcal{O}$, this step does not introduce an infinite loop: points that are inserted are never deleted.

In addition, the algorithm tries to keep in the Delaunay triangulation as many legal facets as possible. Let c be the circumcenter of a poor tetrahedron considered for insertion. If the insertion of c eliminates a legal facet t (see Figure 4(b)), then c is not inserted. Instead, a point p on the intersection of $\partial\mathcal{O}$ and t 's voronoi edge e is inserted.

Figure 3(b) and Figure 3(c) show how our algorithm meshed a kidney; observe that now the boundary facets have vertices lying precisely on $\partial\mathcal{O}$. In the next section, we will demonstrate that our point rejection strategies work also very well on multi-material domains.

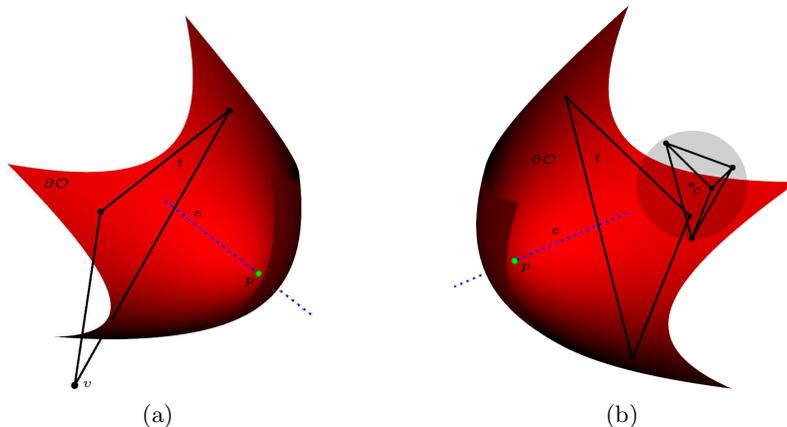


Fig. 4. The point rejection strategies. (a) t is an illegal facet. (b) t is a legal facet.

4 Results

We ran our experiments on a 64 bit machine equipped with a 2.80GHz quad-core Intel i7 processor and 8GB of memory. Our algorithm was built on top of the *Computational Geometry Algorithms Library* (CGAL, <http://www.cgal.org>). We used the *Insight Toolkit* (ITK, <http://www.itk.org>) for image processing. Lastly, the *Visualization Toolkit* (VTK, <http://www.vtk.org>) rendered the meshes.

Figure 5(a) illustrates the output mesh obtained for a segmented CT image taken from IRCAD (<http://www.ircad.fr>). Similarly, Figure 5(b) depicts the output mesh obtained for the MR brodmann atlas (<http://www.sph.sc.edu/comd/rorden/micro.html>). Observe that the mesh elements are of excellent quality. Although we do not give guarantees on the minimum and maximum angles achieved by our method, we observed that the point rejection strategies are able to remove elements with angles less than 19° and more than 150° (in any image input we tried) without creating an edge smaller than $\frac{\delta}{10}$. It would be interesting to theoretically investigate why elements of worse quality are eliminated so easily without introducing very small edges. We leave that exploration as a future work. See the columns “ircad” and “brodmann” of Table 1 for some statistical results.

The last row of Table 1 shows the largest tetrahedron aspect ratio. Aspect ratio is defined as the ratio of a tetrahedron’s circumradius to its inradius. The reported aspect ratio is normalized such that the best aspect ratio equals 1. Therefore, the aspect ratio ranges from 1 to $+\infty$. A high aspect ratio is an indication of bad quality.

Lastly, we show that the size of the mesh can be controlled directly by parameter δ . For this reason, we ran our mesher on the same CT image (obtained by IRCAD), but this time we set the value of δ at 8. That is, we double the value of δ used to obtain the mesh of Figure 5(a). See Figure 5(c) for an illustration

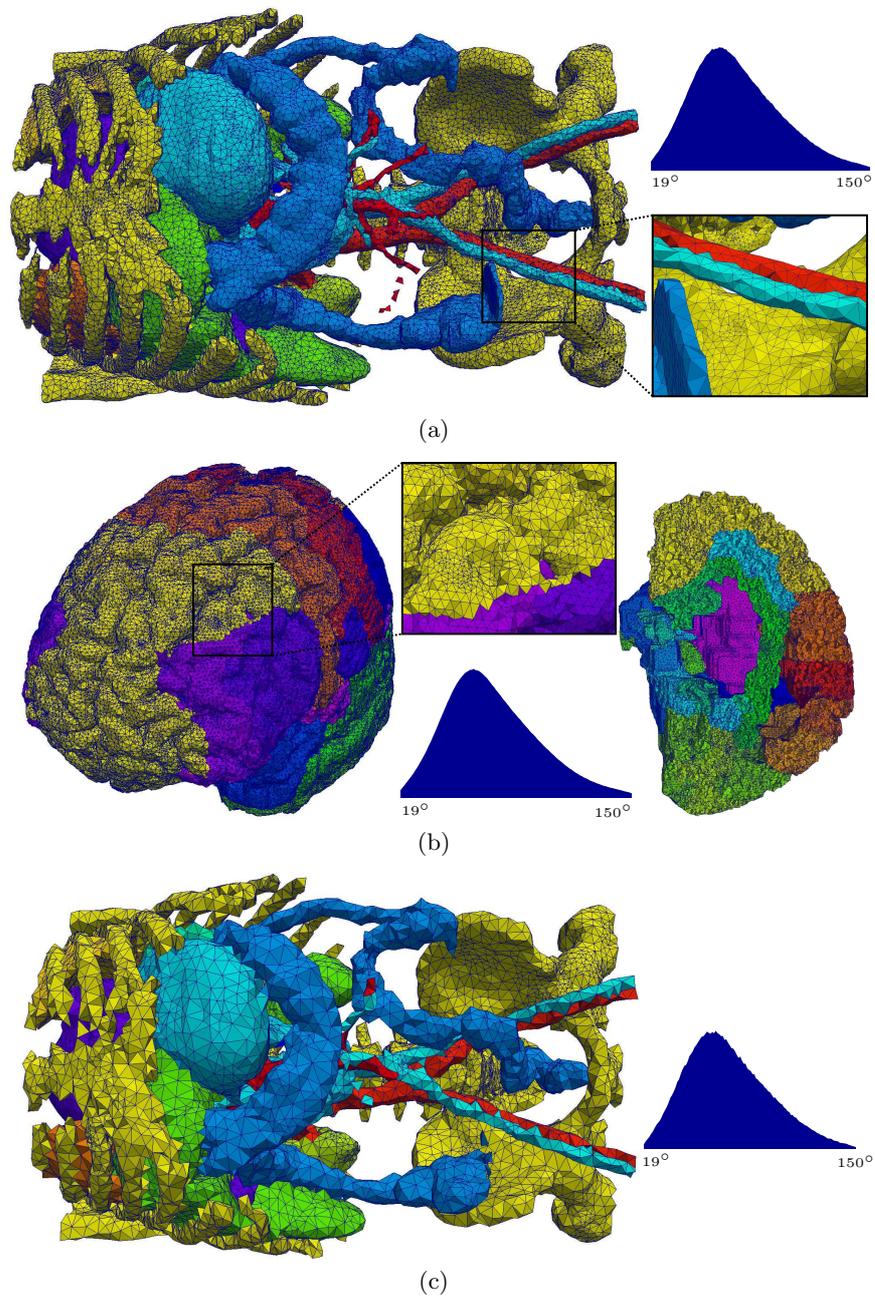


Fig. 5. Whole meshes, zoomed views, cross sections, and distributions of the dihedral angles for (a) the CT multi-label image and (b) the MR brain atlas. In (c), we show a coarser mesh on the same input image that was used in (a).

Table 1. Information about the images used for the evaluation, the chosen value for parameter δ , and some quantitative results of the final meshes produced by our algorithm.

Experiment	ircad	brodmann	ircad(coarse)
Image size	$512 \times 512 \times 219$	$181 \times 217 \times 181$	$512 \times 512 \times 219$
Image resolution (<i>mm</i>)	$0.961 \times 0.961 \times 2.4$	$1 \times 1 \times 1$	$0.961 \times 0.961 \times 2.4$
#Labels	20	41	20
δ (<i>mm</i>)	4	2	8
Time (<i>sec</i>)	421	1,066	96
#Vertices	139,740	473,994	41,097
#Tetrahedra	783,445	2,575,220	173,575
Dihedral angles (degrees)	19 – 150	19 – 150	19 – 150
Max. aspect ratio (normalized)	4.67	6.22	4.55

and the last column of Table 1 for some statistical results. Observe that the number of elements, the number of vertices, and the execution time are greatly reduced, in the expense of worse fidelity. This is an expected trade-off: the fewer elements a mesh has, the less likely it is to represent complex surface creases accurately.

To evaluate our method, we compare it with CGAL (<http://www.cgal.org>). A comparison with other popular meshing techniques like Tetgen (<http://tetgen.berlios.de/>) or Netgen (<http://www.hpfem.jku.at/netgen/>) is omitted in this paper, because they do not operate directly on images. Rather, they require that the surface is already meshed as a piecewise linear complex. In contrast, both our algorithm and CGAL mesh the surface and the volume at the same time.

We run CGAL on the ircad CT image and report the achieved quality. We set CGAL’s sizing parameters to values that gave output meshes with size similar to the size of our mesh depicted in Figure 5(a). Furthermore, we set the quality parameters to their best theoretical values as described in [11]. Quantitative results for CGAL’s output mesh are shown in Table 2. Compare it with the first column of Table 1. Observe that the quality of the CGAL mesh is lower than ours in terms of dihedral angles and aspect ratios. Another popular quality metric is the minimum scaled Jacobian value [19, 20]. It ranges from -1 to 1 with 1 being the best value. A negative value means that some elements are inverted. Both our algorithm and CGAL report positive scaled Jacobian values. In fact, the minimum Jacobian value achieved by our algorithm is 30 times larger than that achieved by CGAL. In terms of absolute numbers, we feel that the Jacobian values of our method is low, an issue we are looking into as future work.

5 Conclusions and future work

In conclusion, we have shown that Delaunay refinement techniques are able to mesh multi-material domains with tetrahedra of very good angles, which makes our method suitable for subsequent finite element analysis. The point rejection strategies proposed in this work maintain mesh conformity and high quality offering, at the same time, control over the mesh size.

Table 2. Quantitative results of the final mesh produced by CGAL.

Experiment	ircad
#Vertices	156,902
#Tetrahedra	756,462
Dihedral angles (degrees)	3 – 174
Max. aspect ratio (normalized)	16,823

Note that surface patches of high curvature need to be meshed with more elements than patches that are not sharp. In its current state, our method meshes the surfaces uniformly. In the future, we plan to extend our method to produce graded triangular surfaces and, therefore, smaller meshes.

Acknowledgments. We thank the anonymous referees for their comments and insight that helped improve the presentation and completeness of the paper. This work is supported in part by NSF grants: CCF-1139864, CCF-1136538, and CSI-1136536 and by the John Simon Guggenheim Foundation and the Richard T. Cheng Endowment.

References

1. Shewchuk, J.R.: What is a Good Linear Element? - Interpolation, Conditioning, and Quality Measures. In: Proceedings of the 11th International Meshing Roundtable, Sandia National Laboratories (2002) 115–126
2. Goksel, O., Salcudean, S.E.: High-quality model generation for finite element simulation of tissue deformation. In: 12th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). MICCAI '09, Berlin, Heidelberg, Springer-Verlag (2009) 248–256
3. Meyer, M., Whitaker, R., Kirby, R.M., Ledergerber, C., Pfister, H.: Particle-based sampling and meshing of surfaces in multimaterial volumes. *IEEE Transactions on Visualization and Computer Graphics* **14(6)** (2008) 1372–1379
4. Liu, Y., Foteinos, P., Chernikov, A., Chrisochoides, N.: Multi-tissue mesh generation for brain images. In: Proceedings of the 19th International Meshing Roundtable, Chattanooga, Tennessee, USA (2010) To appear.
5. Chentanez, N., Alterovitz, R., Ritchie, D., Cho, L., Hauser, K.K., Goldberg, K., Shewchuk, J.R., O'Brien, J.F.: Interactive simulation of surgical needle insertion and steering. In: Proceedings of ACM SIGGRAPH 2009. (2009) 88:1–10
6. Goksel, O., Salcudean, S.E.: Image-based variational meshing. *IEEE Transactions on Medical Imaging* **30** (2011) 11–21
7. Zhang, Y., Hughes, T.J., Bajaj, C.L.: An automatic 3d mesh generation method for domains with multiple materials. *Computer Methods in Applied Mechanics and Engineering* **199** (2010) 405 – 415 *Computational Geometry and Analysis*.
8. Hu, P., Chen, H., Wu, W., Heng, P.A.: Multi-tissue tetrahedral mesh generation from medical images. In: International Conference on Bioinformatics and Biomedical Engineering (iCBBE). (2010) 1–4

9. Hartmann, U., Kruggel, F.: A Fast Algorithm for Generating Large Tetrahedral 3D Finite Element Meshes from Magnetic Resonance Tomograms. In: Proceedings of the IEEE Workshop on Biomedical Image Analysis. WBIA, Washington, DC, USA, IEEE Computer Society (1998) 184–192
10. Boissonnat, J.D., Oudot, S.: Provably good sampling and meshing of surfaces. *Graphical Models* **67** (2005) 405–451
11. Oudot, S., Rineau, L., Yvinec, M.: Meshing volumes bounded by smooth surfaces. In: Proceedings of the International Meshing Roundtable, San Diego, California, USA, Springer-Verlag (2005) 203–219
12. Pons, J.P., Ségonne, F., Boissonnat, J.D., Rineau, L., Yvinec, M., Keriven, R.: High-Quality Consistent Meshing of Multi-label Datasets. In: *Information Processing in Medical Imaging*. (2007) 198–210
13. Dobrina Boltcheva, Mariette Yvinec, Jean-Daniel Boissonnat: Mesh Generation from 3D Multi-material Images. In: *Medical Image Computing and Computer-Assisted Intervention*, Springer (2009) 283–290
14. Cheng, S.W., Dey, T.K., Edelsbrunner, H., Facello, M.A., Teng, S.H.: Sliver exudation. *Journal of the ACM* **47** (2000) 883–904
15. Edelsbrunner, H., Guoy, D.: An experimental study of sliver exudation. *Engineering with Computers* **18** (2002) 229–240
16. Shewchuk, J.R.: Tetrahedral mesh generation by delaunay refinement. In: Proceedings of the 14th ACM Symposium on Computational Geometry, Minneapolis, MN (1998) 86–95
17. Amenta, N., Bern, M.: Surface reconstruction by Voronoi filtering. In: *SCG '98: Proceedings of the fourteenth annual symposium on Computational geometry*, New York, NY, USA, ACM (1998) 39–48
18. Danielsson, P.E.: Euclidean Distance Mapping. *Computer Graphics and Image Processing* **14** (1980) 227–248
19. Knupp, P.M.: Achieving finite element mesh quality via optimization of the jacobian matrix norm and associated quantities. Part II-A framework for volume mesh optimization and the condition number of the jacobian matrix. *International Journal for Numerical Methods in Engineering* **48** (2000) 1165–1185
20. Shepherd, J.F., Tuttle, C.J., Silva, C.T., Zhang, Y.: Quality improvement and feature capture in hexahedral meshes. *SCI institute technical report*, University of Utah (2006)

MRI-free neuronavigation for transcranial magnetic stimulation in severe depression

Benoît Combès^{a,b,c}, Charles Garraud^{a,e}, Xavier Morandi^{a,b,c,d}, Sylvain Prima^{a,b} and Pierre Hellier^{a,e}

^aINRIA, VisAGeS Project-Team, F-35042 Rennes, France

^bINSERM, U746, F-35042 Rennes, France

^cUniversity of Rennes I, F-35043 Rennes, France

^dPontchaillou Univ. Hospital, Department of Neurosurgery, F-35033 Rennes, France

^eSyneika, F-35000 Rennes, France

Abstract. This paper presents a MRI-free neuronavigation technique for repetitive transcranial magnetic stimulation (rTMS). This method is composed of three steps: 1) surface sampling of the subject’s scalp and face thanks to a 3D tracker, followed by Poisson surface reconstruction, 2) non-linear surface registration of an atlas surface to the extracted surface using an efficient modified non-linear EM-ICP algorithm [2], and 3) extrapolation of the transformation to a cortical stimulation target. Results have been obtained on a database of 10 subjects and have shown an accuracy of 10.2 mm. Although clearly less accurate than neuronavigation on the subject’s MRI, we advocate that this neuronavigation method is reproducible and acceptable for routine application of rTMS in severe depression.

1 Introduction

Repetitive transcranial magnetic stimulation (rTMS) is a cortical stimulation technique where a strong electromagnetic field is generated by a coil. The field modifies the neuronal activity beneath the coil and since the stimulation is focal, rTMS has a broad range of potential applications in psychiatry and neurology [13]. It has been recently shown that there is some correlation between the stimulation accuracy and the therapeutic efficacy [4]. Therefore, the stimulation is preferably conducted using a neuronavigation system [11] that allows one to visualize the actual stimulation locus on the subject’s MRI; thus the clinician can move the coil so that the stimulation is at the planned anatomical target defined on the MRI prior to stimulation. In this paper, we propose a new method that allows neuronavigation without MRI in the context of rTMS for severe, drug-resistant depression. It is generally admitted that a good stimulation target for depression is the dorsolateral prefrontal cortex (DLPFC), defined as Brodmann areas 9-46 in [12]. There is a huge number of patients with severe depression that could benefit from rTMS but the cost of the MRI scan, together with the waiting time (approximately between 5 and 15 weeks depending on countries), is a bottleneck for widespread MRI-based neuronavigated rTMS. We

argue that a system that would allow a fast, MRI-free navigation with a slightly degraded accuracy but an excellent reproducibility between sessions would be valuable in clinical routine. Hence, in this paper we propose a surface-based neuronavigation on atlas for rTMS. Firstly, the subject’s skin surface (scalp and face) is sampled thanks to a pointer localized by a 3D tracker of the neuronavigation system and reconstructed using the Poisson method. Secondly, the atlas surface, on which the stimulation target has been localized beforehand, is registered to the subject’s surface thanks to an EM-ICP non-linear registration technique. Thirdly, the non-linear transformation computed on the surface is extrapolated to the cortical stimulation target to predict the patient’s stimulation target for treatment-resistant depression. Accuracy was evaluated on a dataset of 10 subjects and was found to be 10.2 mm on average. The closest work to our method was presented in [8], where the authors used a small number of surface points (22) to register a database of 56 subjects so as to create a functional probabilistic atlas. In our paper, the atlas is deformed toward the subject using a large number of points (generally, several thousands) and an anatomical target is transformed and used in the neuronavigation system to guide a rTMS coil.

2 Material and method

2.1 Method overview

An overview of the method is presented in Figure 1. For a given subject, a surface sampling is first performed, using a pointer tracked by a 3D localizer of the neuronavigation system. Then, the segmented surface of the atlas X is registered to the subject’s surface Y and the resulting transformation T is used to map the coordinates of the target into the subject’s space. The data are then used by the neuronavigation system to guide the rTMS coil. The next sections will describe the following steps: surface acquisition and reconstruction of Y in Section 2.2 and surface registration of X to Y in Section 2.3. Visual results are presented in Section 3.

2.2 Surface acquisition and reconstruction

Anatomical sampling of the subject’s scalp and face is performed using an optical tracking system (Claron Technology Inc., Toronto) composed of the Micron-Tracker camera and recognized markers. Sampling of points is achieved by covering the subject’s head with the calibrated pointer tooltip. A software module was developed to record 4 anatomical landmarks denoted as L^Y on the subject’s head (inion, nasion, left and right tragus) used for an initial landmark-based registration (coined the LDM method hereafter) and then at least 100 surface points. The average number of samples was 3500 points, and the maximal sampling time was 10 minutes.

Then, the set of unorganized points is processed and a closed surface is reconstructed: firstly, the point set is processed through functions dedicated to

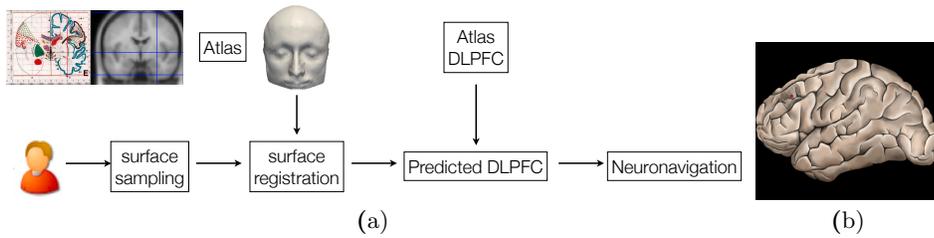


Fig. 1. (a): overview of the surface-based neuronavigation on atlas. (b): a surface rendering of the cortex with the Brodmann area 9/46, or DLPFC, dashed on the surface. To compute Euclidean errors, the area is identified as the centroid, represented here as the red point.

smoothing, outliers removal, normal estimation using the CGAL library [6]. The output is an oriented, smoothed point set suitable for surface reconstruction. Secondly, a surface reconstruction algorithm is applied to generate an iso-surface from the oriented pre-processed point set. 3D surface reconstruction from points set is a well-studied issue. Commonly used surface reconstruction techniques are based on Delaunay triangulation, implicit functions or parametric models. Given that the output of the pre-processing stage is a non-uniform, outlier reduced and smoothed point set, the reconstruction method used in our study should (1) infer the topology of the unknown surface, (2) fit the data, and (3) fill holes. For these reasons, the Poisson reconstruction technique is used [7]. This method allows reconstruction with greater details than other techniques. An implementation of the Poisson surface reconstruction algorithm based on the VTK framework [3] was used. This process produces the point set Y describing the subject's scalp and face, along with the landmark set L^Y .

2.3 Surface registration: efficient EM-ICP algorithm for linear and non-linear registration

In this section, we show how to compute the transformation T that best superposes the surface of the atlas X on the surface of the subject Y (using additional landmarks $L^X = (l_k^X)$ and $L^Y = (l_k^Y)$) so as to compute the coordinates of the target in the subject's space.

General scheme The EM-ICP algorithm [5] is an efficient and elegant solution for rigid registration of point sets. It relies on a probabilistic modeling of the point-to-point correspondences that allows i) a pragmatic definition of the superimposition of two point sets and ii) to deal with a relatively smooth cost function to minimize (in contrast to the classical ICP algorithm). Moreover, it is mathematically well-grounded (monotonic convergence), generic (no assumption on tessellation/topology/number of points) and can be specialized in many ways (*e.g.* to deal with non-linear deformations or to estimate shape models [1]).

Most importantly, the EM-ICP algorithm [5] allows dealing very efficiently and robustly with large point sets. This algorithm can be shown to be equivalent to the alternated iterative minimization of an energetic criterion. In this work, we propose a slight modification of the underlying criterion that allows one i) to use pairs of landmarks to constrain the registration, ii) to symmetrize the estimation of point-to-point correspondences and iii) to add a regularization on T to deal with non-linear deformations:

$$\begin{aligned} \mathcal{E}(T(X), Y, A, B) = & \sum_{j,k} (A_{jk} + B_{jk}) \rho_{\delta}(\|y_j - T(x_k)\|^2) + 2\sigma^2 \beta \sum_k \|l_k^Y - T(l_k^X)\|^2 \\ & + 2\sigma^2 \sum_{j,k} A_{jk} \log(A_{jk}) + 2\sigma^2 \sum_{j,k} B_{jk} \log(B_{jk}) + 2\sigma^2 \alpha L(T), \end{aligned}$$

with $\forall j, \sum_k A_{jk} = 1$ and $\forall k, \sum_j B_{jk} = 1$ (1)

where

- if we drop matrix B for a moment, for the sake of clarity, $A = (A_{jk})$ is the unknown *match* matrix encoding the *a posteriori* probabilities of correspondence between points of X and Y . This probabilistic interpretation of A is made possible thanks to the barrier function $\sum_{j,k} A_{jk} \log(A_{jk})$ [1]. In essence, the greater A_{jk} , the more likely the point $x_k \in X$ to be the correspondent of the point $y_j \in Y$. σ^2 is the Gaussian noise variance of X . This fuzzy control on the correspondences allows one to handle problems due to differences of sampling/number of points between X and Y : we do not look for one-to-one correspondences between points of each surface but instead for “fuzzy” correspondences linking each point of Y to each point of X . This match matrix is a row stochastic matrix, which leads to many-to-one correspondences. This asymmetric formulation makes the algorithm unable to achieve a good matching in specific cases and can make the choice of the source and target sets critical. To tackle this problem, our criterion also contains a second match matrix B , that is column stochastic, in addition to the row stochastic matrix A . The *a posteriori* probabilities of correspondences between points of X and Y is encoded in both matrices A and B , making the point-to-point matching much more symmetrical compared to [5].
- T is a transformation (to be later defined) superposing X on Y . This transformation is subject to a regularizer L .
- α and β are two positive parameters weighing respectively the regularization $L(T)$ term and the landmark-to-landmark discrepancy term $\sum_k \|l_k^Y - T(l_k^X)\|^2$ over the other terms.

- the function $(\rho_\delta : r \mapsto r \text{ if } r < \delta \text{ and } \delta \text{ else})$ is a robust function allowing points of X (resp. Y) having no homologue in Y (resp. X) to be discarded from the estimation of T .

It can be shown that the criterion \mathcal{E} can be minimized over the three unknown parameters T , A and B using the following two-step algorithm:

Algo NL-Sym-EM-ICP: Symmetric robust non-linear EM-ICP with landmarks

Step 1:

initialise \tilde{A} and \tilde{B} to the null matrix

$\forall x_k \in X;$

$\mathcal{S} = \{y_j \in Y \text{ such that } \|y_j - \tilde{T}(x_k)\|^2 < \delta\}$ (using a *kd-tree*)

$\forall y_j \in \mathcal{S};$

$\forall y_j \in \mathcal{S}; \tilde{A}_{jk} = \exp(-(\|y_j - \tilde{T}(x_k)\|^2 / (2\sigma^2)))$

$\forall y_j \notin \mathcal{S}; \tilde{A}_{jk}$ is left equal to 0

$\tilde{B} = \tilde{A}$

normalise \tilde{A} in rows and \tilde{B} in columns

compute the vectors $(\tilde{p}_j) : \forall j, \tilde{p}_j = 0$ if $\sum_i \tilde{A}_{ji} = 0$ and 1 else

compute the vectors $(\tilde{q}_k) : \forall k, \tilde{q}_k = 0$ if $\sum_i \tilde{B}_{ik} = 0$ and 1 else

Step 2: solve the approximation problem:

$\arg \min_T \sum_{j,k} (\tilde{p}_j \tilde{A}_{jk} + \tilde{q}_k \tilde{B}_{jk}) \|y_j - T(x_k)\|^2 + 2\sigma^2 \beta \sum_k \|T(l_k^X) - l_k^Y\|^2 + 2\sigma^2 \alpha L(T)$

Specifying T/L In practice, we design a coarse-to-fine approach to estimate T by first computing the rigid transformation best superposing L^X on L^Y and then using the NL-Sym-EM-ICP algorithm by modeling T successively as a rigid R , an affine F and a non-linear transformation W . For rigid and affine transformations T , we take $L(T) = 0$ and Step 2 has a closed-form solution. For the non-linear transformation, we parametrize W as a deformation field (*i.e.* $W(x) = x + t(x)$) and design L as a scalar Fourier-based regularizer over t :

$$L(t = (t_1, t_2, t_3)^T) = L(t_1) + L(t_2) + L(t_3), \text{ with } L(t_i) = \frac{1}{(2\pi)^3} \int_{-\infty}^{\infty} \frac{|t_i^*(\omega)|^2}{\phi^*(|\omega|/b)} d\omega,$$

where $*$ is the Fourier transform operator, $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is an integrable function and b is a real strictly positive rescaling factor. We choose ϕ as a Wu compactly supported positive definite kernel. The advantage of such a regularization (over the TPS regularizer for example) is that it i) provides efficient solutions for Step 2 based on sparse linear algebra [14] and ii) introduces a scaling parameter b . The larger the b value, the more drastic the penalization of the high frequencies, *i.e.* by choosing a large b value, we focus on capturing the global/high-scale deformation superposing X and Y .

Interpolating the target position Each estimated transformation R , F and W is defined over \mathbb{R}^3 and has a closed-form expression. As a result, any point y of Y (thus those representing the DLPFC) can be mapped in the space of X as $y' = W(F(R(y)))$.

2.4 Material

Phantom To assess the accuracy of the point acquisition and reconstruction steps, a human-head phantom was used, shown in Figure 2-(a). The phantom was scanned with an isotropic 0.5 mm resolution CT (size $512 \times 512 \times 841$). A set of 4099 points was acquired as described in Section 2.2. The surface of the phantom was extracted using heuristic thresholding and mathematical morphology operators.

Subjects The experiments were performed using anatomical brain MRI of 10 healthy volunteers. Each subject had a brain MRI scan using a 3D T1-weighted sequence with 1 mm isotropic resolution, and the skin surface was extracted using thresholding and mathematical morphology operators. On the subjects, the surface representing the scalp and face was sampled using the process described in Section 2.2.

MRI Atlas In the experiments, the MRI atlas described in [9] was used. The skin surface was extracted using thresholding and mathematical morphology operators, and the four landmark points L^X were manually localized. This atlas was used instead of the Colin27 atlas since on the latter, the surface extraction process is difficult due to an antenna visible on the top of the atlas.

DLPFC localization Manual positioning of the DLPFC on the subjects' MRI and on the atlas was performed by an expert in neuroanatomy. The DLPFC was defined as the second third, *i.e.* middle part, of the middle frontal gyrus along the antero-posterior axis, corresponding to Brodmann areas 9/46 as described in [12], and as shown in Figure 1-(b). We have chosen to characterize the position of the DLPFC by its centroid, shown by the red dot. Therefore, hereafter the DLPFC is considered as a point, what enables to compute Euclidean distances between estimated targets and reference targets.

3 Validation and results

3.1 Poisson reconstruction accuracy

The point set acquired on the physical phantom was registered to the phantom's skin surface extracted from the CT using a rigid EM-ICP. The computed transformation was then applied to the reconstructed Poisson surface. The mean point-to-point error between the Poisson surface and the surface extracted from

the CT scan was 1.6 mm (standard deviation 1.08 mm). Results are presented in Figure 2. These results demonstrate the accuracy of the point sampling and show that the reconstructed surface is valid to represent the subject’s scalp and face.

3.2 Distance between predicted and reference DLPFC

The skin surface extracted on the MRI was registered toward the Poisson reconstructed surface using rigid EM-ICP. On the other side, the Poisson reconstructed surface was registered toward the atlas skin surface using the non-linear EM-ICP. Therefore, the predicted DLPFC can be compared to the reference DLPFC. Various degrees of freedom of the transformation were tested: landmark-based registration that corresponds to a Procrustes alignment using the 4 anatomical landmarks (LDM), LDM followed by a rigid registration (R), R followed by an affine registration (RA) and RA followed by a non-linear registration (NL). Table 1 shows the obtained errors for the 10 subjects for all transformations. Firstly, it can be observed that the mean error decreases when the degrees of freedom of the transformation increase. Secondly, the error obtained with the fully non-linear method is 10.2 mm on average. These results are discussed hereafter.

4 Conclusion

In this paper, a new method was proposed to perform neuronavigation without MRI for transcranial magnetic stimulation. The method is based on a subject’s surface sampling and reconstruction, followed by a non-linear registration of an atlas surface to the obtained surface. The transformation is then used to map the coordinates of the DLPFC in the atlas to the subject’s space. Firstly, we found the surface sampling and reconstruction on the phantom to be accurate (mean error 1.6 mm, standard deviation 1.08 mm). Secondly, we compared the surface-based DLPFC localization to a ground truth localization performed by an expert in neuroanatomy. Results have shown on 10 subjects that the mean error was 10.2 mm.

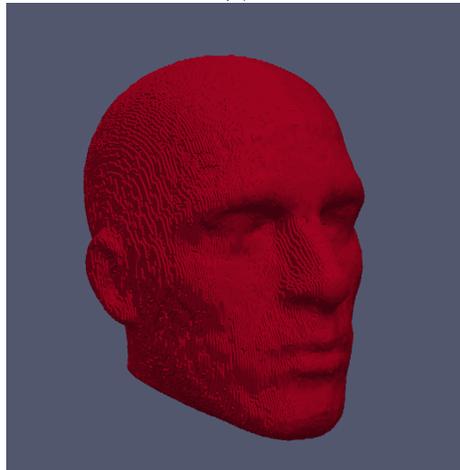
In a previous study [10], we have assessed the performance of 3 clinicians on 25 subjects to localize the DLPFC on MRI and found a large discrepancy between raters: accuracy compared to reference localization varied between 8 mm and 14 mm (standard deviation 5.7 mm). In comparison, the mean error of 10.2 mm can be considered acceptable. In addition, we focus on rTMS application where the electromagnetic field has some dispersion. Therefore, we argue that, although less accurate than neuronavigation using the subject’s MRI, this MRI-free neuronavigation technique is reproducible between sessions and is valuable for a routine clinical use of rTMS. Further work should 1) enlarge the database of subjects and extend this technique to other stimulation targets, such as the temporoparietal cortex to treat auditory hallucination in schizophrenia for instance and 2) evaluate the therapeutic efficiency on patients.



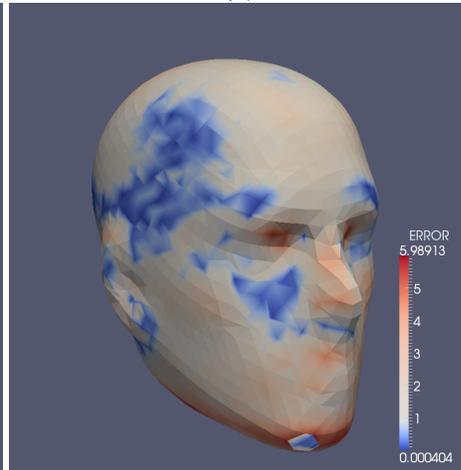
(a)



(b)



(c)



(d)

Fig. 2. This figure shows the validation of the acquisition and reconstruction of a point set on a physical-head phantom. (a) shows a picture of the physical phantom used to assess the surface acquisition and reconstruction method. (b) shows the acquired point set and the reconstructed Poisson surface. (c) shows the skin surface extracted from CT. (d) shows the Poisson surface registered to the CT surface. The color represents the spatial distribution of errors after reconstruction and registration. The mean point-to-point error was 1.6 mm with standard deviation was 1.08 mm.

	Landmark-based	Rigid	Affine	Non-linear
Subject 1	11.3	5.7	3.5	3.3
Subject 2	13.1	13.5	12.1	11.9
Subject 3	23.8	15.1	13.6	12.3
Subject 4	27.7	10.5	8.8	9.2
Subject 5	16.8	4.7	3.2	3.1
Subject 6	22.5	9.1	7.6	6.8
Subject 7	32.9	21	20.3	19.5
Subject 8	22.7	18.7	14.2	13.6
Subject 9	21.3	19.2	15.6	15.4
Subject 10	23.3	7.7	6.4	6.7
Mean (mm)	21.54	12.52	10.53	10.18
Std.dev. (mm)	6.46	5.86	5.56	5.33

Table 1. Results of the surface-based localization of the DLPFC when compared to the reference DLPFC. Various parametrization of the surface-based method were tested. The fully non-linear registration provided a mean error of 10.2 mm.

References

1. H. Chui, A. Rangarajan, Jie Zhang, and C. Morison Leonard. Unsupervised Learning of an Atlas from Unlabeled Point-Sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:160–172, 2004.
2. B. Combès and S. Prima. An Efficient EM-ICP Algorithm for Symmetric Consistent Non-linear Registration of Point Sets. In *MICCAI'2010*, volume 6362, pages 594–601. LNCS, 2010.
3. D. Doria and A. Gelas. Poisson surface reconstruction for VTK. *The VTK Journal*, 2010.
4. P.B. Fitzgerald, K. Hoy, S. McQueen, J.J. Maller, S. Herring, R. Segrave, M. Bailey, G. Been, J. Kulkarni, and Z.J. Daskalakis. A randomized trial of rTMS targeted with MRI based neuro-navigation in treatment-resistant depression. *Neuropsychopharmacology*, 34(5):1255–62, 2009.
5. S. Granger and X. Pennec. Multi-scale EM-ICP: A Fast and Robust Approach for Surface Registration. In *ECCV'2002*, volume 2353, pages 418–432. LNCS, 2002.
6. <http://www.cgal.org>. Computational Geometry Algorithms Library.
7. M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Eurographics Symposium on Geometry Processing*, 2006.
8. J. Koikkalainen, M. Könönen, J. Karhu, J. Ruohonen, E. Niskanen, and J. Lötjönen. Dynamic probabilistic atlas of functional brain regions for transcranial magnetic stimulation. In *MICCAI'2008*, volume 5241, pages 543–550. LNCS, 2008.
9. F. Lalys, C. Haegelen, J.-C. Ferré, O. El Ganaoui, and P. Jannin. Construction and assessment of a 3-T MRI brain template. *NeuroImage*, 49(1):345 – 354, 2010.
10. C. Nauczyciel, P. Hellier, B. Millet, and X. Morandi. Localization of the dorso-lateral prefrontal cortex in MRI images for neuronavigated TMS. *Submitted to Psychiatry research*.
11. C. Nauczyciel, P. Hellier, X. Morandi, S. Blestel, D. Drapier, J.-C. Ferré, C. Barillot, and B. Millet. Assessment of standard coil positioning in transcranial magnetic stimulation in depression. *Psychiatry research*, 186(2-3):232–238, 2011.

12. M. Petrides. Lateral prefrontal cortex: architectonic and functional organization. *Philos Trans R Soc Lond B Biol Sci*, 360(1456):781–795, 2005.
13. M. Ridding and J. Rothwell. Is there a future for therapeutic use of transcranial magnetic stimulation? *Nature Rev. Neurosci.*, 8:559–567, 2007.
14. H. Wendland. *Scattered Data Approximation*. Cambridge Monographs on Applied and Computational Mathematics, 2005.

Shape based Conditional Random Fields for Segmenting Intracranial Aneurysms

Sajjad Baloch^{*1}, Erkang Cheng¹, Ying Zhu¹, Ashraf Mohamed¹, Haibin Ling²,
and Tong Fang¹

¹ Siemens Corporate Research, Princeton NJ

² Temple University, Philadelphia, PA

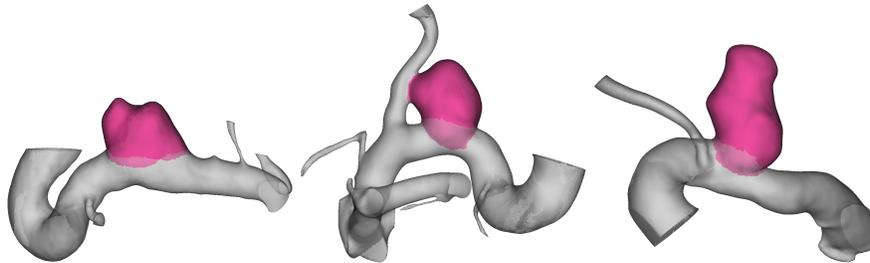


Fig. 1. Examples of intracranial sidewall aneurysms.

Abstract. Studies have found strong correlation between the risk of rupture of intracranial aneurysms and various physical measurements on the aneurysms, such as volume, surface area, neck length, among others. Accuracy of risk prediction relies on the accuracy of these quantities, which in turn, is determined by the precision of the underlying segmentation algorithm. In this paper, we propose an algorithm for the separation of aneurysms in pathological vessels. The approach is based on conditional random fields (CRF), and exploits regional shape properties for unary, and layout constraints for pair-wise potentials to achieve a high degree of accuracy. To this end, we construct very rich rotation invariant shape descriptors, and couple them with randomized decision trees to determine posterior probabilities. These probabilities define *weak priors* in the unary potentials, which are also combined with *strong priors* determined from user interaction. Pairwise potentials are used to impose smoothness as well as spatial ordering constraints. The proposed descriptor is independent of surface orientation, and is richer than existing approaches due to attribute weighting. The conditional probability of CRF is maximized through graph-cuts, and the approach is validated with real dataset w.r.t. the groundtruth, resulting in the area overlap ratio of 88.1%. Most importantly, it successfully solves the “touching vessel leaking” problem.

* Corresponding author: Sajjad Baloch, Siemens Corporate Research, Princeton, NJ, sajjad.baloch@siemens.com

1 Introduction

Intracranial aneurysms is a major vascular disease in the brain, attributed to local weakening of the vessel wall. It manifests in the form of bulging (saccular aneurysm) or dilation (fusiform aneurysm) as shown in Fig. 1. Intracranial aneurysms frequently occur near areas of high arterial curvature or bifurcations, as these regions usually experience more hemodynamic stress [7]. If left untreated, an aneurysm grows in size, thereby further weakening the wall strength and increasing the risk of rupture, which may lead to subarachnoid hemorrhage, neurological deficits, and in up to 56% of cases mortality [14]. In order to prevent its growth and reduce the risk of rupture, surgical intervention is required, where stents, wire coils, and other embolic material or devices are placed not only to enforce the vessel wall, but also to alter the blood flow pattern, thereby reducing the pressure on regions more prone to rupture.

Once diagnosed, aneurysms are carefully monitored and examined, before making a surgical decision. To this end, geometry of aneurysm plays a crucial role. Physicians analyze various measurements of geometric primitives [5] evaluated on an aneurysm, which allows them to carry out surgical planning. Separation of the aneurysm from the healthy vessel, therefore, serves as a critical step, whose accuracy determines the eventual outcome in terms of surgical decisions, device selection, as well as patient recovery. The problem is, however, very challenging due to the complex topology and geometry of the underlying blood vessel, and its large inter-patient variation.

Early approaches on aneurysms separation mainly focused on deformable models. By grouping local shape descriptors, McLaughlin and Nobel [10] employed a region-splitting algorithm to segment the aneurysm from the neighboring vasculature. Their approach, however, fails to yield reliable results for wide-neck saccular aneurysms. Wong and Chung [16] modeled the healthy part of the vessel as a tubular deformable model, and determine the abnormal structure of aneurysms as the complement of the healthy model. Their approach does not provide protection against the leaking of the deformable model inside the aneurysm. Ford et al. [6], on the other hand, presented a method to reconstruct the parent artery by removing the aneurysm. They also utilized deformable model to model vessel of parent. The method needs smooth surface and is limited to the morphology of parent artery surface.

More recently, Mohamed el at. [11] utilized mesh based snakes for computer-aided planning for endovascular treatment of intracranial aneurysms (CAPETA). A major limitation of their approach was its inability to adapt to topological changes in the aneurysm boundary, resulting in erroneous results for meshes with touching vessels. Furthermore, due to the local nature of the snakes, their proposed method fails to guarantee a globally optimal solution. To address these issues, Sgouritsa el at. [13] proposed a curvature based graph-cut strategy for segmenting the 3D vessel mesh, guided by strong priors, which in turn were determined from user input in the form of 3 seed points. Due to its dependence on strong priors, the accuracy of this approach is determined by the accuracy of the prior computation algorithm, and remains highly sensitive to the

manual user input. Also, in practical situations, the gaussian curvature is too coarse to adequately capture underlying shape variations. It is, therefore, required to construct richer shape representations, which adequately capture the vessel-aneurysm differences.

In this paper, we propose a novel method for aneurysm separation based on conditional random fields (CRF) by augmenting both unary and pairwise potentials. Unary potentials are decomposed into strong and weak priors. The former are determined from interactive user input, and the latter are estimated via randomized decision forest as posterior probabilities of novel shape descriptors. More specifically, we propose algorithms for accurate strong prior determination. The shape descriptors, in turn, are constructed from underlying local and regional geometry. Pairwise potentials are encoded to impose spatial ordering as well as smoothness constraints. The conditional probability of CRF is maximized through graph-cuts. This formulation allows multitude of improvements over existing approaches. Like [13], our method is not limited by topological variations. Unlike [13], we construct shape descriptors specifically tailored for the problem under consideration, i.e., for the separation of blob like aneurysms on tubular vessels. Inferencing from examples allows us to specify weak priors, thereby allowing more flexibility. Smoothness and layout constraints penalize the assignment of inconsistent labels. We validate the proposed method with real dataset comprising of 27 3D digital subtraction angiographic (DSA) images in the CAPETA framework [11], and carry out a comparison with [13]. The experimental results demonstrate that (1) in all cases our improved prior seeds were in accordance with the groundtruth, and (2) our aneurysm method consistently outperforms [13].

The rest of the paper is organized as follows: We first formulate aneurysm separation as CRF in Section 2, followed by the construction of shape descriptors in Section 3. Unary and pairwise potentials are developed in Section 4, along with the algorithms for finding strong priors. Experimental results are presented in Section 5, before we concluded in Section 6.

2 Problem Formulation

For robust aneurysm separation algorithm, we formulate the problem as Conditional Random Fields (CRF) that are driven by rich shape descriptors.

2.1 Conditional Random Fields

Given a triangular mesh $\mathcal{T} := (P, E)$ representing a surface \mathcal{M} embedded in \mathbb{R}^3 , with $P := \{p_i = p(v_i)\}$ denoting the set of positions at vertices $V = \{v_i\}$, and $E = \{e_k\}$ denoting the edges connecting the respective vertices. $p : V \rightarrow P, v_i \mapsto p_i$, therefore, forms an isomorphism from the undirected graph $\mathcal{G} := (V, E)$ to \mathcal{T} . The problem under consideration is to find a binary labeling $l : V \rightarrow L = \{l_v, l_a\}, v_i \mapsto l_i := l(v_i)$ that partitions \mathcal{G} into two segments. For aneurysm separation problem, this amounts to segmenting a mesh into the healthy vessel

and the aneurysm regions, and is dictated by some feature properties of the underlying geometry.

Suppose x_i defines a shape or geometric descriptor of v_i in \mathcal{T} possibly with a non-local region of support, then the isomorphism $x : V \rightarrow X, v_i \mapsto x_i := x(v_i)$ captures the underlying geometric description of v_i and the corresponding graph $X := (x_i, e_i)$ may be exploited to find the partitioning. The optimal partition is the one that maximizes the joint distribution of (X, L) :

$$\begin{aligned} l^* &= \arg \max_l P(X, l; \theta) \\ &= \arg \max_l P(X; \theta) P(l|X; \theta) \\ &= \arg \max_l P(l|X; \theta), \end{aligned} \tag{1}$$

where θ represents a distribution parameter. Conditional random field setting allows one to simplify the above expression, by considering a local (1-ring) neighborhood $N(v_i)$ at each vertex, v_i :

$$P(l|X; \theta) = \frac{1}{Z(\theta, X)} \prod_i \phi_i(l_i, X; \theta) \prod_{(i,j)} \psi_{ij}(l_i, l_j, X; \theta), \tag{2}$$

where the *unary potential*, ϕ_i , captures the posterior distribution of labels at v_i , and the *pairwise potential*, ψ_{ij} , models the neighborhood labeling relations allowing one to impose spatial constraints. Depending on the application, one may incorporate various constraints, such as smoothing or spatial ordering, in the form of soft layout consistency or hard layout consistency [18].

The problem is, therefore, reduced to constructing appropriate shape descriptors x , estimating posterior probabilities, and specifying a pairwise potential that is suitable for the application. Maximization of Eq. (2) is identical to the minimization of the following energy functional:

$$E(l|X) = \sum_i \log \phi_i(l_i, X; \theta) + \sum_{(i,j)} \log \psi_{ij}(l_i, l_j, X; \theta). \tag{3}$$

3 Shape Descriptors

We are interested in separating blob like structures, such as aneurysms, from somewhat tubular regions, such as blood vessels. For each point p_i on \mathcal{T} , we extract surface features, \mathbf{F}_i , that are highly discriminating between these kind of regions. They include regional shape as well as local geometry.

3.1 Local Descriptors

For local descriptors, we rely on the curvature information. We exploit Gaussian curvature κ , mean curvature H , and maximum and minimum principal curvatures, κ_1 and κ_2 . In addition, we also consider maximum and minimum principal directions, \mathbf{v}_1 and \mathbf{v}_2 , and the shape index, s :

$$s = \frac{2}{\pi} \arctan \frac{\kappa_2 + \kappa_1}{\kappa_2 - \kappa_1}, \quad \kappa_2 \geq \kappa_1. \quad (4)$$

3.2 Regional Shape Descriptors

The regional shape information is captured through various shape descriptors, namely (1) Wilmore energy, and (2) regional attribute weighted geodesic shape contexts.

Wilmore Energy The Wilmore energy of a vertex $v_i \in \mathcal{G}$ is defined in terms of the isomorphism \mathcal{T} , and its n -ring neighborhood, S_i :

$$W := \int_{S_p} (H^2 - \kappa) dA, \quad (5)$$

where dA is a surface area element of S_p .

Note that $W(p) \geq 0$, with $W(p) = 0$ if and only if p is convex, and v and all of its neighbors, S_p lie on a common sphere[2]. Consequently, big blob like structures, such as aneurysms, are characterized by small Wilmore energy.

Regional Attribute Weighted Geodesic Shape Contexts Shape contexts [1] are defined by creating bins of various spatial parameters, followed by constructing a histogram that counts points falling in each bin. [12] proposed 3D shape contexts for surface matching, which bin the 3D space via spherical coordinates. Such descriptors, however, are not invariant to the orientation of the surface. To overcome this problem, we carry out geodesic binning for each point $p \in \mathcal{M}$, as illustrated in Fig. 2. Geodesic distances are intrinsic to a surface, and, therefore, lead to rotation invariance. If $g(p, \cdot)$ is the geodesic distance from point p , then geodesic binning, within a local neighborhood $\mathcal{G}_r(p) := \{\forall q \in \mathcal{M} : g(p, q) \leq r\}$, is defined as $\{i = 0, \dots, k-1 : g_i \leq g_i(p, q) < g_{i+1}\}$ with $g_k = r$. Histograms are then generated by computing the concentration of various surface attributes within each bin. Although apparently similar to [18], which has also utilized geodesic binning for layout consistent segmentation of ear impressions, we consider regional binning instead of doing it globally. In problem under consideration, global binning causes a confounding effect, as random branching across individuals introduces noise in the sample data. Furthermore, we consider various histograms of diverse surface attributes, such as Gaussian curvature weighted point distribution, area distribution, and the distribution of connected components.

The Gaussian curvature weighted point distribution shape context $\mathbf{f}^g = (f_0^g, \dots, f_{k-1}^g)$ computes the number of points falling within each bin normalized by the total number of points in all bins to create a distribution. Finally each bin is weighted by the Gaussian curvature averaged within it. For meshes with non-uniform triangulation, the discrete number of points does not truly represent the underlying geometry. For this reason, we augment our feature set,

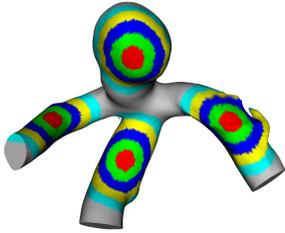


Fig. 2. Geodesic binning at selected vertices.

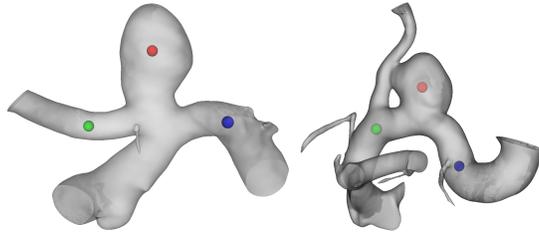


Fig. 3. Reference input provided by a user in the form of dome, proximal, and distal points.

by computing the surface area of each bin to create a second shape context distribution $\mathbf{f}^a = (f_0^a, \dots, f_{k-1}^a)$. It should be noted that this does not make the previous shape context redundant, since for non-uniform triangulation the point density is related to interesting features. We, therefore, retain both descriptors, and will later carry out feature selection strategy to retain the most distinguishing features.

The third and final shape context $\mathbf{f}^c = (f_0^c, \dots, f_{k-1}^c)$ captures the number of connected components in each bin. This shape context is particularly important to differentiate flat or thick regions from narrow tubular areas. The hypothesis is that such a descriptor will help in segmenting a touching vessel from an aneurysm.

Visibility from Reference Point Due to their almost convex shape, most points on an aneurysm are visible from their centroid r . This allows us to consider a very powerful feature in our feature design:

$$\nu_r(p) := \begin{cases} 1 & \text{if } p \text{ is visible from } c \\ 0 & \text{o.w.} \end{cases} \quad (6)$$

Since the centroid is not known a priori, we assume that its rough location is specified by a user as a reference point.

All these features are combined in a feature vector $\mathbf{F}_p := (\kappa, H, \kappa_1, \kappa_2, s, \mathbf{v}_1, \mathbf{v}_2, W, \mathbf{f}^g, \mathbf{f}^a, \mathbf{f}^c, \nu_r)$ as a local and regional descriptor of a point $p \in \mathcal{T}$.

4 Potential Specification

After constructing the shape descriptors, we are in a position to specify unary potentials and pairwise potentials. Unary potentials are determined from the probability of aneurysm and pairwise potentials are used to impose smoothness and layout constraints. In this paper, we decompose the unary potentials into strong and weak priors; the former is determined from the user input and latter incorporates posterior probabilities learned from shape descriptors. The term strong highlights high confidence, and hence, large weights for such priors.

Major differences of our approach from [13] include: (1) rich shape descriptors, as opposed to scalar descriptor (Gaussian curvature) in [13]; (2) we employ CRF framework, which naturally encodes the posterior probabilities of aneurysm; (3) weak prior unary potentials are estimated by randomized decision trees; (4) new robust algorithms for finding strong priors, where [13] finds erroneous priors; (5) smoothness and layout constraints in pairwise potentials.

4.1 Strong Unary Potentials

For strong priors, we adopt an approach similar to recently proposed CAPETA framework [11], where a user specifies so-called *dome*, *proximal* and *distal* points as shown in Fig. 3. The dome point p_D very roughly provides the location of the aneurysm relative to the vessel, and proximal p_p and distal p_d points specify the region of interest for subsequent analysis in CAPETA. The aneurysm, therefore, always falls between the proximal and distal point input. It should be noted that these points do not lie on the surface, but are centered inside the vessel in the viewing direction. The closest point projection \tilde{p}_D of p_D is used as p_{ref} in Eq. (10), and p_D is employed as the reference point r for the visibility feature of Eq. (6).

Vessel Prior We exploit these points for establishing strong priors in our unary potential. [13] used two kind of contours for strong priors. The strong prior for the vessel was determined from the proximal p_p and distal p_d points, by projecting them to the mesh, \tilde{p}_p and \tilde{p}_d , and then using the geodesic h between them as the vessel prior. The limitation of this approach is that frequently touches the boundary between the aneurysm and the healthy vessel, therefore, leading to incorrect specification of the strong vessel priors as illustrated in Fig. 4(a). To overcome this limitation, we enforce the geodesic to stay away from the boundary. The optimal geodesic is the one that simultaneously minimizes the distance between \tilde{p}_p and \tilde{p}_d , as well as the minimum mean curvature along the path:

$$h^*(\tilde{p}_p, \tilde{p}_d) := \arg \min_{g \in \mathcal{M}} \int_{\tilde{p}_p}^{\tilde{p}_d} \varphi(H) \pi(h(\tilde{p}_p, \tilde{p}_d)) dh, \quad (7)$$

where φ is a decreasing functional of mean curvature H , and π is the length of the geodesic. As shown in Fig. 4(b), the modified geodesic leads to more reliable vessel prior A_v .

Aneurysm Prior In [13], aneurysm prior was based on a planar contour $\mathcal{C} := \mathcal{T} \cap P$ found as an intersection between \mathcal{T} and a plane P centered at the dome point p_D . The plane normal \mathbf{n} was defined as the direction of a vector found as follows. First, the closest point p_c to the dome point p_D is determined on the vessel centerline between p_p and p_d . \mathbf{n} is then:

$$\mathbf{n} := \frac{p_D - p_c}{\|p_D - p_c\|_2}. \quad (8)$$

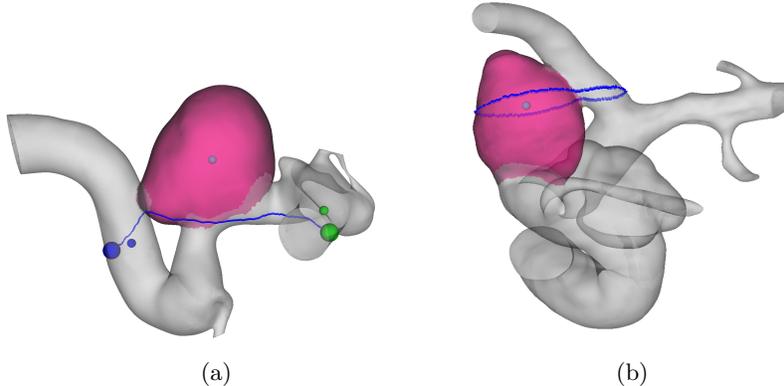


Fig. 4. Problems with the strong priors of [13]: (a) Vessel prior; (b) Aneurysm prior. Small spheres represent the reference points that lie inside the vessel. Big spheres are the mapping of the references points to the vessel surface.

Finally, the intersection contour is employed as the strong prior for the aneurysm region. The utility of this prior is limited to the extent that there is only one intersection contour, which is usually not true. This problem can be resolved by selected the correct single connected component, in the case of touching vessels. In addition, the plane may cut across at the touching point, and the intersection contour may consist of both aneurysm and the touching vessel part. We exploit the minimum principal curvature to decompose the contour into various segments, and eventually use the largest connected segment as the strong aneurysm prior A_a .

4.2 Weak Unary Potential

Weak unary potentials are learned from examples. In general, x_i can reside in some high dimensional space, and the learning process may be quite challenging. Given a dataset of pathological vessels, and the associated expert labeled groundtruth, a randomized decision forest is constructed for the classification similar to [15,18], based on the extraction of the above mentioned shape descriptors. The advantages of randomized forests include: (1) its built-in feature selection mechanism, where maximum information gain is used for node splitting, and (2) its ability to avoid over-fitting without pruning. For test vessels, posterior probabilities are computed and used as weak priors.

The strong priors are incorporated in the unary potential as follows:

$$\log \phi_i(l_i, X; \theta) = \begin{cases} \gamma_3 & \text{if } p_i \in A_a \text{ and } l_i = l_v \\ \gamma_4 & \text{if } p_i \in A_v \text{ and } l_i = l_a \\ \log \phi_i(l_i, X; \theta) & \text{o.w.} \end{cases} \quad (9)$$

where γ_3, γ_4 are the penalization costs for incorrect label assignments.

4.3 Pairwise Potential

In order to derive pairwise potential, we assume that a reference point p_{ref} is given in the region l_a . A spatial layout constraint is then introduced that penalizes the assignment of l_v to a vertex that is closer in geodesic distance sense to p_{ref} than its neighbor with a label l_a . Similarly, we penalize the assignment of different labels to neighboring vertices to ensure a smoothness constraint:

$$-\log \psi_{ij}(l_i, l_j, X; \theta) = \exp(\alpha(\kappa_i + \kappa_j)) + \begin{cases} 0 & \text{if } l_i = l_j \\ \gamma_1 & \text{if } l_i \neq l_j, l_i = l_a, g(p_i, p_{\text{ref}}) \geq g(p_j, p_{\text{ref}}) \\ \gamma_2 & \text{if } l_i \neq l_j \end{cases} \quad (10)$$

where κ_i and κ_j denote the Gaussian curvature at points $p_i, p_j \in \mathcal{T}$. The first term on the right hand side, therefore, penalizes the assignment of identical labels to neighboring vertices with high curvature edges. γ_1 and γ_2 are the costs assigned empirically or inferred from data. α is an exponential shaping coefficient.

Eventually segmentation is carried out by minimizing Eq. (3) through the α -expansion algorithm [4,9,3].

5 Experiments

In this section, we compare the proposed method with [13]. 3D Digital Subtraction Angiographic (DSA) images of 30 patients are acquired, and subsequently thresholded as described in [11], to extract the pathological vessels. 3D triangular meshes were constructed through marching cubes, and were then decimated using quadric decimation [8]. An expert is asked to provide manual labeling on these meshes. Various features were computed to construct regional shape descriptors at each mesh vertex. Curvature estimates were based on adaptive ball approach [17] due to its robustness to noise. 10-ring neighborhood was utilized for Wilmore energy, and 3mm neighborhood was considered for computing geodesic shape contexts.

4 problem cases were identified, where [13] had failed to provide reasonable results. These examples, along with 4 additional randomly selected meshes were considered as the test dataset. The remaining meshes were added to the training dataset, and training of randomized forest (comprising of 60 trees) was carried out, to compute coarse aneurysm separation and posterior estimates. Some examples are given in Fig. 5.

For the test cases, vessel and aneurysm labels for strong priors were determined as described in Section 4.1. For comparison, results are shown in Fig. 6 for the problem cases highlighted in Fig. 4.

Energy functional of Eq. (3) was minimized via two iterations of α -expansion algorithm. A quantitative measure was defined as $QM := (A \cap B) / (A \cup B)$, where A is an automatic separation, whereas B is the groundtruth. $\gamma_1 = 3.5, \gamma_2 = 5, \gamma_3 = \gamma_4 = 10^9$ in Eqs. (9) and (10) were determined by maximizing QM

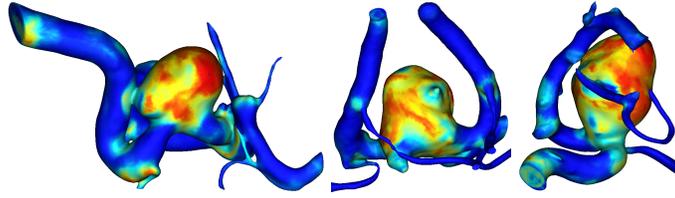


Fig. 5. Posterior label probabilities for example cases.

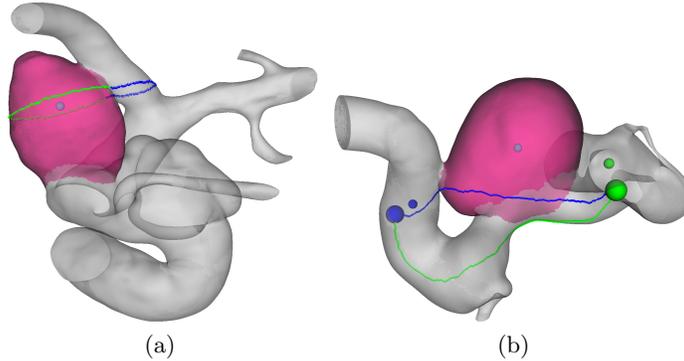


Fig. 6. Strong prior labels determined by the method proposed in Section 4.1 for the problem cases of Fig. 4: (a) Aneurysm strong prior labels; (b) Vessel strong prior labels. Note that the modified aneurysm prior (green contour) does not go inside the touching vessel, and the vessel prior (green contour) stays away from aneurysm region.

via alternating variables. Results are given in Fig. 7, providing a qualitative comparison with [13]. Overall, the average QM for our method was computed to be 88.1%, compared with 73.0% for [13]. Most importantly, the leaking problem with [13] into the touching vessels is completely resolved by our approach.

6 Conclusions

We have proposed a CRF based method for the separation of aneurysms from healthy vessel regions. A unique strength of the method is that it effectively combines two sources of information, patient specific strong priors in the form of user input, and data-driven shape priors learned from a large number of aneurysm examples. Along with strong prior obtained from user’s interactive input, posterior probabilities learned by randomized decision trees with rich shape descriptors are considered as unary potentials. Smoothness of the segmentation is ensured through pairwise potentials, which are also enriched with spatial ordering constrains. Final segmentation is achieved by minimizing the resulting energy functional with graph-cuts. The proposed method is validated with a

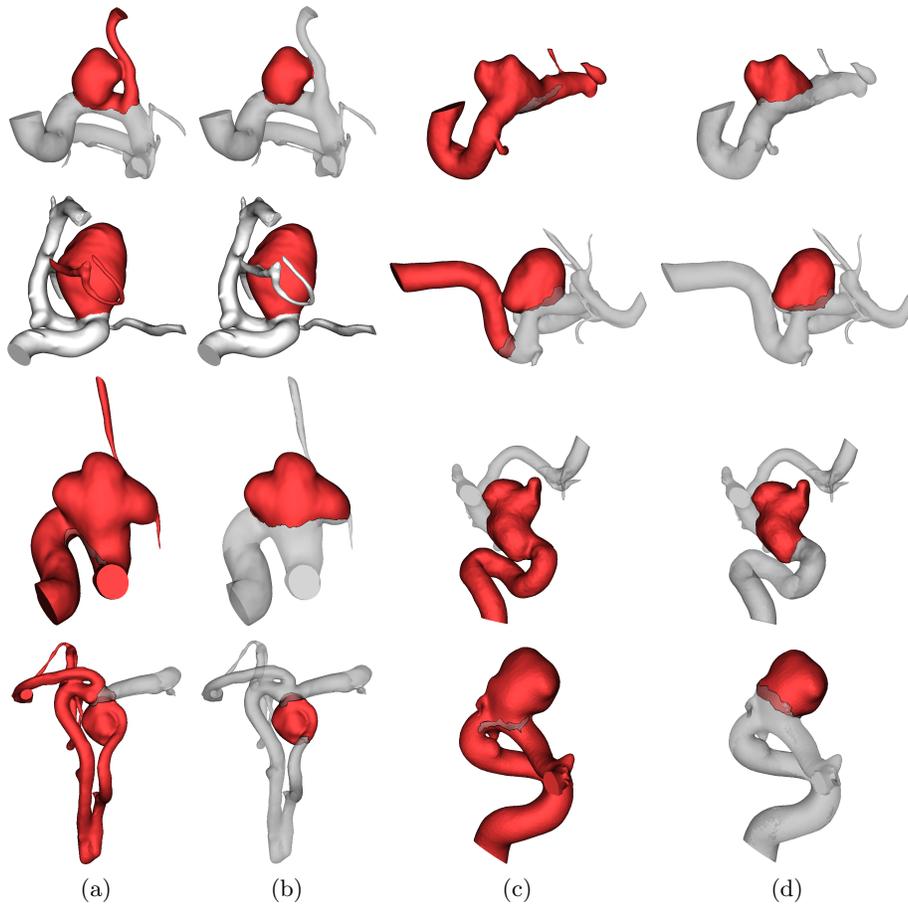


Fig. 7. Comparison with [13]. (a) and (c) [13]; (b) and (d) The proposed method. Latter outperforms [13] especially for the touching vessels.

real dataset with outstanding results (an accuracy measure of 88.1%). It also perfectly resolves the touching vessel leaking problem.

7 Acknowledgements

We gratefully acknowledge Dr. Michael E. Mawad, St. Luke’s Episcopal Hospital, for providing with the valuable data used in the experiments in the paper.

References

1. Belongie, S.J., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 509–522 (2002) 5

2. Bobenko, A.I.: A conformal energy for simplicial surfaces. *Combinatorial and computational geometry* 52, 133–143 (2005) [5](#)
3. Boykov, Y., Kolmogorov, V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *PAMI* 26, 1124–1137 (2004) [9](#)
4. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. *PAMI* 23, 1222–1239 (2001) [9](#)
5. Debrun, G., Aletich, V., Kehrl, P., Misra, M., Ausman, J., Charbel, F., Showkneen, H.: Aneurysm geometry: an important criterion in selecting patients for Guglielmi detachable coiling. *Neurologia Medico-Chirurgica* 38, S1–S20 (1998) [2](#)
6. Ford, M., Hoi, Y., Piccinelli, M., Antiga, L., Steinman, D.: An objective approach to digital removal of saccular aneurysms: technique and applications. *British Journal of Radiology* 82(Special Issue 1), S55–S61 (2009) [2](#)
7. Foutarakis, G., Yonas, H., Scialabassi, R.: Saccular aneurysm formation in curved and bifurcating arteries. *American Journal of Neuroradiology* 20(7), 1309 (1999) [2](#)
8. Garland, M., Heckbert, P.: Surface simplification using quadric error metrics. In: *Annual Conference on Computer Graphics*. pp. 209–216 (1997) [9](#)
9. Kolmogorov, V., Zabih, R.: What energy functions can be minimized via graph cuts? In: *ECCV*. pp. 65–81 (2002) [9](#)
10. McLaughlin, R.A., Noble, J.A.: Demarcation of aneurysms using the seed and cull algorithm. In: *MICCAI*. pp. 419–426 (2002) [2](#)
11. Mohamed, A., Sgouritsa, E., Morsi, H., Shaltoni, H., Mawad, M.E., Kakadiaris, I.A.: Computer-aided planning for endovascular treatment of intracranial aneurysms (CAPETA). In: *SPIE Conf. Med. Imaging*. vol. 7625 (2010) [2](#), [3](#), [7](#), [9](#)
12. Park, G.J., Novotni, M., Klein, R.: 3D shape matching with 3D shape contexts. In: *Int. Conf. in Central Europe on Comput. Graph. and Vis.* (2003) [5](#)
13. Sgouritsa, E., Mohamed, A., Morsi, H., Shaltoni, H., Mawad, M.E., Kakadiaris, I.A.: Neck localization and geometry quantification of intracranial aneurysms. In: *ISBI*. pp. 1057–1060 (2010) [2](#), [3](#), [7](#), [8](#), [9](#), [10](#), [11](#)
14. Weir, B.: Unruptured intracranial aneurysms: a review. *Journal of Neurosurgery* 96(1), 3–42 (2002) [2](#)
15. Winn, J.M., Shotton, J.: The layout consistent random field for recognizing and segmenting partially occluded objects. In: *CVPR*. pp. 37–44 (2006) [8](#)
16. Wong, W.C.K., Chung, A.C.S.: Augmented vessels for quantitative analysis of vascular abnormalities and endovascular treatment planning. *IEEE TMI* 25(6), 665–684 (2006) [2](#)
17. Yang, Y.L., Lai, Y.K., Hu, S.M., Pottmann, H.: Robust principal curvatures on multiple scales. In: *Eurographics Symposium on Geometry processing*. pp. 223–226. *SGP '06* (2006) [9](#)
18. Zouhar, A., Baloch, S., Tsin, Y., Fang, T., Fuchs, S.: Layout consistent segmentation of 3-D meshes via conditional random fields and spatial ordering constraints. In: *MICCAI*. pp. 113–120 (2010) [4](#), [5](#), [8](#)

3D Surface Realignment Tracking for Medical Imaging: A Phantom Study with PET Motion Correction

Oline V. Olesen^{1,2,3}, Rasmus R. Paulsen¹, Rasmus R. Jensen¹, Sune H. Keller², Merence Sibomana², Liselotte Højgaard², Bjarne Roed³, and Rasmus Larsen¹

¹ Informatics and Mathematical Modelling, Technical University of Denmark
Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark
ovol@imm.dtu.dk, <http://www.imm.dtu.dk/>

² Department of Clinical Physiology, Nuclear Medicine & PET, Rigshospitalet,
Copenhagen University Hospital, University of Copenhagen

³ Siemens Healthcare, Siemens A/S, Denmark

Abstract. We present a complete system for motion correction in high resolution brain positron emission tomography (PET) imaging. It is based on a compact structured light scanner mounted above the patient tunnel of the Siemens High Resolution Research Tomograph PET brain scanner. The structured light system is equipped with a near infrared diode and uses phase-shift interferometry to compute 3D representations of the forehead of the patient. These 3D point clouds are progressively aligned to a reference surface and thereby giving the head pose changes. The estimated pose changes are used to reposition a sequence of reconstructed PET frames. To align the structured light system with the PET coordinate system a novel registration algorithm based on the PET transmission scan and an initial surface has been developed. The performance of the complete setup has been evaluated using a custom made phantom based on a plastic mannequin head equipped with two positron emitting line sources. Two experiments were performed. The first simulates rapid and short head movements, while the second simulates slow and continuous movements. In both cases, the system was able to produce PET scans with focus the PET reconstructions. The system is near ready for clinical testing.

1 Introduction

Patient head movement during high resolution brain positron emission tomography (PET) scanning will cause blurring and ghosting [3]. The low count rate and resulting low contrast makes it near impossible to perform motion correction on raw PET data and therefore most methods rely on external tracking of the head movement [10, 11, 12, 16]. The Polaris Vicra (Northern Digital Inc.) tracking system has been used as the reference on many PET installations [6]. While the Polaris system is well tested and accurate it suffers from problems related to attaching optical markers to the patients head. Experience shows that

in a clinical setting, the markers are difficult to attach such that they will stay in position during the entire scan. A markerless system that fits into the narrow PET tunnel will improve the clinical acceptance and the diagnostic value of PET brain scans.

We have previously described a structured light (SL) based system that is based on a small projector and two small cameras [7] for tracking patient head pose. It has been designed to fit to the patient tunnel of the Siemens High Resolution Research Tomograph (HRRT) PET brain scanner. The HRRT PET scanner has a spatial resolution down to 1.4 mm [8] and is therefore well suited for testing new motion correction methods. The SL system is based on the progressive reconstruction of 3D surfaces of the upper face region of the patient placed in the scanner. The pose changes are found by computing the rigid transformation between the current scan and the initial surface scan. The system described in [7] was based on visible light and did not operate in real time. Visible light scanners are not suited for repeated human facial scans. Furthermore, to be functional the system needs to acquire frames sufficiently fast to faithfully capture patient head movements. In this paper, we describe a system using invisible light with a camera acquisition rate of 30 frames per second.

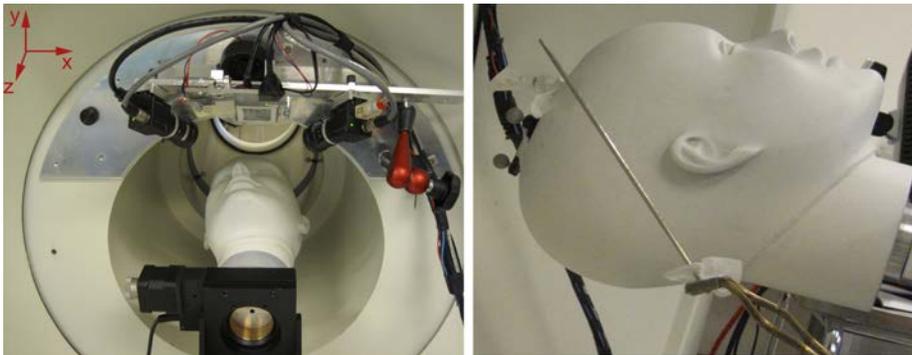


Fig. 1. Left: The patient tunnel of the HRRT PET scanner with the invisible structured light system mounted. The phantom is mounted on a rotating stage rotated to the right (-10 degrees). Right: The phantom including one of the two radioactive line sources. It is placed in the head in the same angle as shown here simulating the brain.

While the previous paper has focused on measuring the accuracy of the SL tracking system using a rotation stage as ground truth [7], the real interest is the improvement of the PET scans. To be able to evaluate the quality improvement of the PET scan a scan using a radioactive tracer must be performed. A common approach is to use a phantom and compare the resulting PET scan with the known geometry of the phantom [8]. We have therefore designed a custom phantom with a radioactive source and used that in the evaluation of the SL system based motion correction. Our system setup with the SL system and the

custom phantom can be seen in Fig. 1. Compared to other external tracking systems, where the geometric alignment between the tracking system and the PET scanner can be problematic, we investigate a novel alignment approach based on aligning the SL system scan directly to the PET transmission scan.

2 Experiments and Methods

The SL system consists of two Point Grey Flea2 cameras (1288 x 964 pixels) each running at 30 frames per second. The Pico Digital Light Processing (DLP) projector from Texas Instruments is used to project phase modulated patterns. One of the light diodes of the Pico projector has been replaced with a near infrared (NIR) diode resulting in a non-visible projected pattern. The projector is controlled by a GFM Pico developer kit board that also sends trigger pulses to the cameras, in order to synchronize the projected patterns and the shutter of the cameras. A multithreaded C++ program running on a standard portable computer acquires the real time camera data and stores them as image files. The 3D point cloud generation, surface reconstruction, and alignment are done in a post-processing step. The SL system and HRRT PET acquisition computer are synchronized through an internal network time protocol (NTP) server.

The 3D point cloud generation is based on phase-shifting interferometry (PSI) [5] where a set of 2D interferograms are projected and projector-camera correspondence can be found using phase unwrapping. This is explained in detail in [5, 7]. While three patterns are used in [7], the system is now extended to use six patterns with varying wavelengths to make the phase unwrapping more robust to discontinuities in the surface. Since each point cloud computation requires six frames, the effective tracking frequency is 5 Hz. A surface is reconstructed using a modern algorithm [9] based on the point cloud acquired in the initial position. The facial pose changes are then found by rigidly aligning the following surface scans to this reference surface using an optimized iterative closest point (ICP) algorithm [15].

To correct for motion, we need to know the transformation between the HRRT PET scanner coordinate system and the SL system. To estimate this transformation, we use the transmission scan of the HRRT PET scanner, which is also used for the attenuation and scatter correction within normal PET reconstruction. The transmission scan is a voxel volume similar to a computed tomography (CT) scan. The initial reference surface scan is captured by the SL system during the transmission scan thus creating correspondence. The transformation is computed using a pseudo-ICP algorithm [13]. The surface scan is scaled to fit the volume, and manually rotated and translated into an initial position. To find correspondence between surface sample points and the volume, the volume is sampled in the normal direction (both positive and negative) of the surface scan to find the point with maximum gradient. Knowing the general orientation of the patient in the PET scanner, we use the absolute gradient in the x (left/right) and z (axial) direction and the negative gradient for the y

(anterior/posterior) direction:

$$\Delta f = \left| \frac{\partial f}{\partial x} \right| - \frac{\partial f}{\partial y} + \left| \frac{\partial f}{\partial z} \right|$$

With a point correspondence, a rigid transformation is found using the closed-form loop to estimate the absolute transformation [4]. With an initialization, transformation this process is iterated until the transformation of the SL scan converges to the volume data.

In order to apply the motion correction to the PET data we apply the multiple acquisition frames (MAF) method [10]. In [8] the MAF method was demonstrated on the Siemens HRRT PET scanner using the tracking input from the Polaris Vicra system. We divide the PET emission list mode data into equal time length intervals and for each interval, a PET frame is reconstructed using the 3-D ordered subset expectation maximization (3D-OSEM) algorithm with resolution modeling and incorporating a spatially invariant point spread function [14]. These frames are then repositioned to a reference position using the SL tracking system.

A custom phantom with known geometry was designed. It consists of a hollow plastic mannequin head with a very low attenuation coefficient. Two radioactive line sources are inserted into the head to provide activity for the HRRT PET scanner. The activity of the line sources are 2×7.8 MBq each created by a positron-emitting germanium-68/gallium-68 generator. As can be seen in Fig. 1, the line sources goes through the head from the back of the skull to the forehead. The phantom was mounted onto a rotation stage by Thorlabs and placed in the patient tunnel of the HRRT PET scanner.

Two experiments where performed using the stage to rotate the head. In experiment one, the head was rotated from -20 to 20 degrees in steps of 5 degrees. At each position a 30 s frame was PET reconstructed and repositioned. Data with motion was excluded from the reconstruction. In experiment two, the head was rotated from -10 to 10 degrees in a continuous motion with a maximum speed of one degree per second. The PET data was reconstructed using one second frames in experiment two. Experiment one simulates the clinical situation where the patient is performing a rapid head motion followed by a stationary period. State of the art practise is to discard PET data during such rapid motions. The second experiment simulates e.g. a patient falling asleep where the head is slowly drifting from side to side.

We evaluate the effect of the motion correction on the reconstructed PET images by calculating Dice's coefficient (percent volume overlap) [1] between a reference image recorded without phantom motion and the motion distorted and the SL system based motion corrected image. The number of voxels, N , included in the calculation was set to a value corresponding to the number of voxels inside the tubes 2.5 times the diameter of the PET sources used (outer diameter 3.2 mm and active length 168 mm) [8]. The extended volume is used in order to compensate for partial volume effects. In each image to be studied the set of the N most intense voxels are extracted and used for the Dice's coefficient computation presented as the percentage overlapping voxels. In addition we compute the

normalized cross correlation between the reference image and each imageframe either motion corrected or uncorrected [2].

3 Results and Discussion

The rigid transformation between the coordinate system of the SL system and the PET image frame is obtained from using the described surface to volume alignment. Figure 2(a) shows the reference surface aligned to maximum gradient points in the transmission scan.

Figure 2(b) shows the results of the first experiment with stepwise rotation of the phantom. The top figure shows the percentage overlap between PET frames of the line sources in the reference position and a scan position as a function of the performed rotation of the head. The bottom plot of Fig. 2(b) shows similarly the correlation coefficient between the reference image and a motion correction/uncorrected image for the different scan positions. Results based on tracking information from the left and right camera of the SL system are shown in green and blue colors respectively while the red curve represents the uncorrected image results. The overlap and the correlation measures are in agreement. The results of the uncorrected frames decreases with the size of performed rotation from an overlap of 100% down to 2% at ± 20 degrees. The overlap of the motion corrected reconstruction is improved significantly for all positions with percentage overlap of 71-93%. The overlap is not 100% which is mainly due to the internal calibration of the SL system, the ICP alignment, and the geometrical alignment with the HRRT PET scanner. In addition the interpolation error, which combined with the straight and narrow line sources of the phantom (with a diameter similar to the voxel size of 1.2 mm) induces partial volume effects and thereby decreasing the overlap and correlation measure. The differences between the left camera and the right camera could be explained by the construction of the reference surface scan where left camera was chosen as basis. The result is similar to [8] where the overlap was 65-85% for a 10 degrees corrected rotation. However, the two studies cannot directly be compared since the phantom designs are different.

A visual evaluation of the motion correction is shown in Fig. 3 for the maximum rotation of 20 degrees. The PET images are summed along one dimension and visualized on top of each other pairwise in the red and green color channels. The overlapping pixels of the two PET images appear yellow. The top row of the figure shows the reference image and the uncorrected image as two sets of rods rotated approximately 2 cm at the end points. These correspond to the relevant brain regions: the frontal lobe and cerebellum. The bottom row shows the reference image and the corrected image seen as two yellow rods demonstrating a near perfect motion correction. This position has an overlap of 89% in contrast to the rotation of -20 degrees with an overlap of 71%.

The results of the second experiment with a continuous rotation of 20 degrees demonstrating the realtime pose registration of the SL system is presented in Fig. 4. The one second PET frames are summed and fused with the transmission

image of the phantom. The top of the figure shows a row of uncorrected images, where the motion of the line sources is seen as blurred circle parts. The bottom row shows the motion corrected image, where the previously blurred parts appear in focus and with high intensity. The cross section of the line sources shows dots with a diameter of only a few pixels. Long drift motion is a very complex problem to overcome using image registration methods for motion estimation, why an external tracking system is of great value. Our latest results show that continuous motion can be tracked in real time and PET frames successfully corrected.

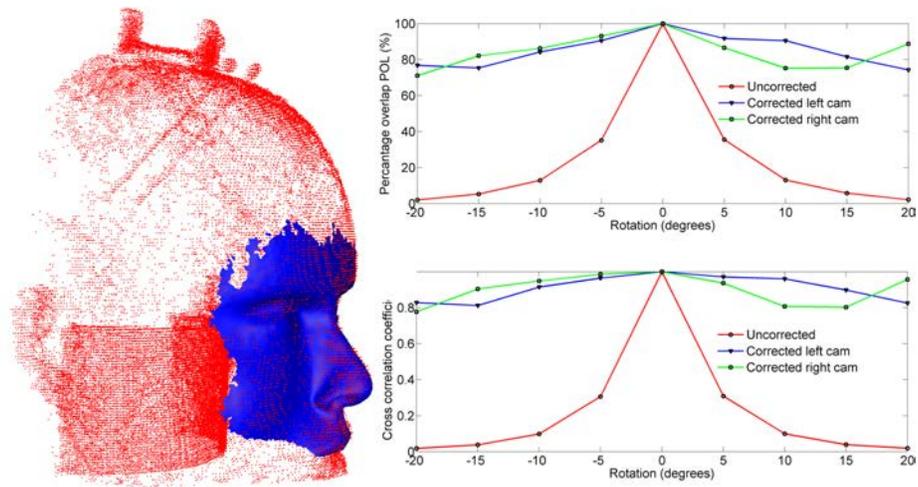


Fig. 2. Figure 2(a) shows the alignment between maximum gradient points in the transmission scan shown as red dots and the SL face scan shown as a blue surface. Figure 2(b) shows quantitative results of the stepwise experiment. Top: the percentage number of overlapping points in the reference PET image compared to the unaligned/aligned PET images based on either the right or the left camera. Bottom: the cross correlation between the reference PET image and the unaligned/aligned PET images.

4 Summary and Conclusions

This paper describes a complete system for motion correction in high resolution PET brain imaging. It is based on a small and flexible structured light scanner mounted above the patient tunnel of the PET scanner. The scanner is equipped with a near infrared light source making it suitable for future patient examinations. Furthermore, the system is tracking the head pose changes with a frequency of 5 Hz, which is suitable for the head movement experienced during real clinical PET scanning. In order to align the two systems a novel algorithm

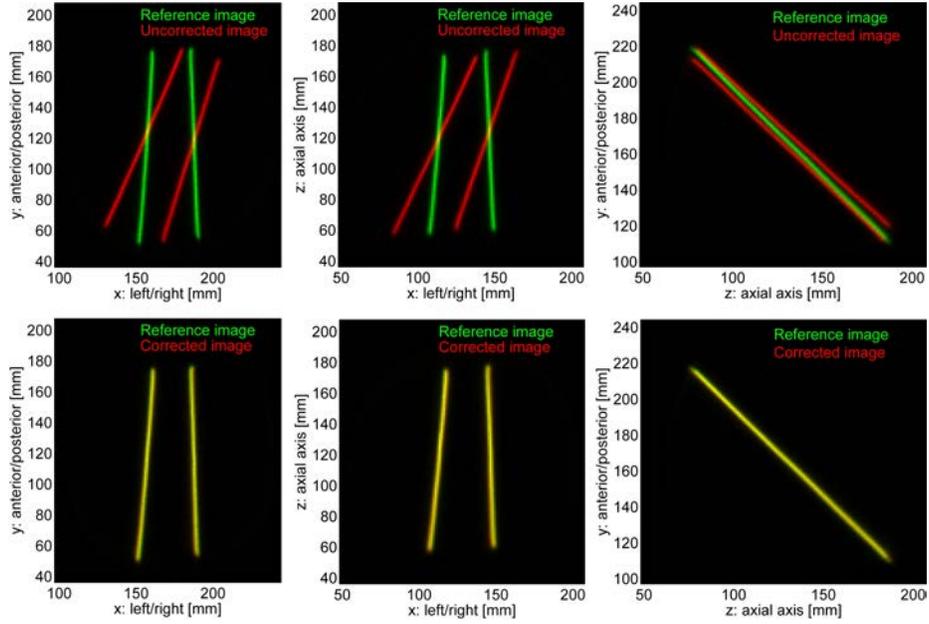


Fig. 3. The figure shows the summation of the PET images along 3 different axes for a reference image (shown in green) and a target image (shown in red) with a 20 degree rotation (overlap = 89%, shown in yellow). The uncorrected image is shown as captured in the first row, while the second row shows the image after motion correction.

using the HRRT PET transmission scan and the initial surface scan was presented. The performance of the system was evaluated using a custom designed phantom with two radioactive line sources mounted on a programmable rotation stage. The results of the two experiments are very promising. The first experiment simulates rapid but short head movements and the second experiments slow but longer head movements. Quantitative analysis show that the combined system is able to robustly reduce motion artifacts and greatly improve PET scans for scenarios involving both slow and rapid movements. The system is near ready for actual clinical testing.

Acknowledgements The John and Birthe Meyer foundation is thanked for the generous donation of the HRRT PET scanner and the cyclotrons for isotope production.

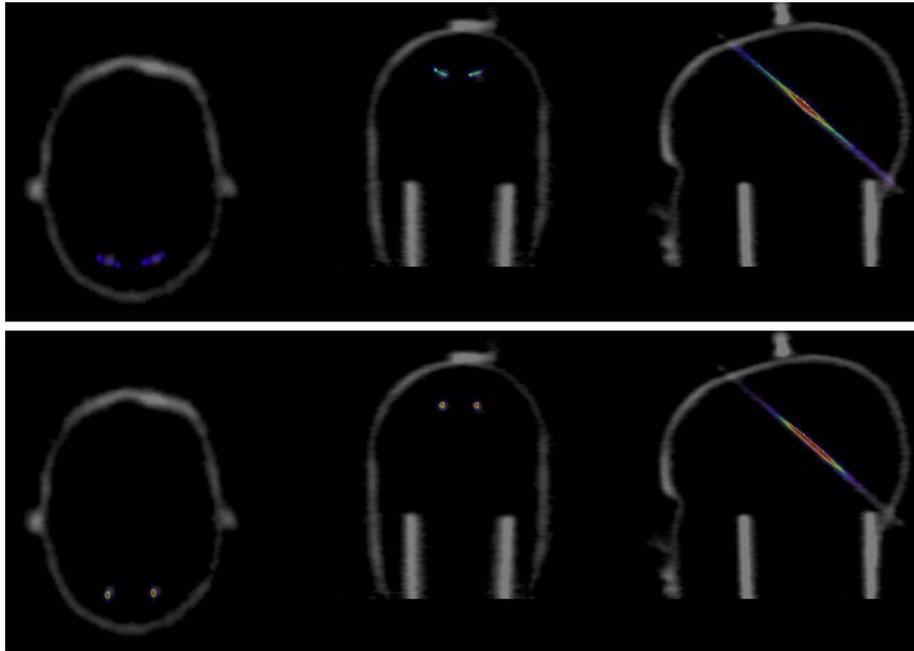


Fig. 4. Results of the dynamic PET scan. One hundred one second frames uncorrected (top) and MAF motion corrected (bottom) are summed and fused with a transmission scan. The frame repositioning is based on the left camera alignment.

References

1. Dice, L.: Measures of the amount of ecologic association between species. *Ecology* 26(3), 297–302 (1945)
2. Gonzalez, R., Woods, R.: *Digital image processing*, a. dourkin, ed. upper saddle river, nj. Prentice Hall pp. 205–208, 701–704 (2002)
3. Green, M.V., Seidel, J., Steina, S.D., Tedder, T.E., Kempner, K.M., Kertzman, C.: Head movement in normal subjects during simulated PET brain imaging with and without head restraint. *J. Nuclear Medicine* 35(9), 1538–1546 (1994)
4. Horn, B.K.P.: Closed form solution of absolute orientation using unit quaternions. *Journal of the Optical Society A* 4(4), 629–642 (1987)
5. Huang, P., Hu, Q., Jin, F., Chiang, F.: Color-encoded digital fringe projection technique for high-speed three-dimensional surface contouring. *Optical Engineering* 38, 1065 (1999)
6. Lopresti, B.J., Russo, A., Jones, W.F., Fisher, T., Crouch, D.G., Altenburger, D.E.: Implementation and performance of an optical motion tracking system for high resolution brain PET imaging. *IEEE Trans. on Nuclear Science* 46(6), 2059–2067 (1999)
7. Olesen, O., Paulsen, R., Højgaard, L., Roed, B., Larsen, R.: Motion Tracking in Narrow Spaces: A Structured Light Approach. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010* pp. 253–260 (2010)

8. Olesen, O., Svarer, C., Sibomana, M., Keller, S., Holm, S., Jensen, J., Andersen, F., Højgaard, L.: A Movable Phantom Design for Quantitative Evaluation of Motion Correction Studies on High Resolution PET Scanners. *Nuclear Science, IEEE Transactions on* 57(3), 1116–1124 (2010)
9. Paulsen, R.R., Bærentzen, J.A., Larsen, R.: Markov Random Field Surface Reconstruction. *IEEE Transactions on Visualization and Computer Graphics* 16(4), 636–646 (2010)
10. Picard, Y., Thompson, C.J.: Motion correction of PET images using multiple acquisition frames. *IEEE Trans. on Medical Imaging* 16(2), 137–144 (1997)
11. Raghunath, N., Faber, T.L., Suryanarayanan, S., Votaw, J.R.: Motion correction of PET brain images through deconvolution: II. Practical implementation and algorithm optimization. *Physics in Medicine and Biology* 54(3), 813 (2009)
12. Rahmim, A., Dinelle, K., Cheng, J.C., Shilov, M.A., Segars, W.P., Lidstone, S.C., Blinder, S., Rousset, O.G., Vajihollahi, H., Tsui, B., Wong, D.F., Sossi, V.: Accurate event-driven motion compensation in high-resolution PET incorporating scattered and random events. *IEEE Trans. on Medical Imaging* 27(8), 1018–1033 (2008)
13. Rusinkiewicz, S., Levoy, M.: Efficient variants of the ICP algorithm. In: 3dim. p. 145. Published by the IEEE Computer Society (2001)
14. Sureau, F., Reader, A., Comtat, C., Leroy, C., Ribeiro, M., Buvat, I., Trebossen, R.: Impact of image-space resolution modeling for studies with the high-resolution research tomograph. *Journal of Nuclear Medicine* 49(6), 1000 (2008)
15. Wilm, J., Olesen, O.V., Paulsen, R.R., Højgaard, L., Roed, B., Larsen, R.: Real time surface registration for pet motion tracking. *Proc. SCIA* (2011)
16. Woo, S.K., Watabe, H., Choi, Y., Kim, K.M., Park, C.C., Bloomfield, P.M.: Sinogram-based motion correction of PET images using optical motion tracking system and list-mode data acquisition. *IEEE Trans. on Nuclear Science* 51(3), 782–788 (2004)

Example based style classification

Katarzyna Welnicka, Jakob Andreas Bærentzen, Henrik Aanæs, and Rasmus Larsen

DTU Informatics, Technical University of Denmark

Abstract. We address the problem of analysis of families of shapes which can be classified according to two categories: the main one corresponding usually to the coarse shape which we call the function and the more subtle one which we call the style. The style and the function both contribute to the overall shape which makes the general analysis and retrieval of such shapes more challenging. Also there is no single way of defining the style as this depends much on the context of the family of shapes used for the analysis. That is why the definition needs to be given through the examples.

The straight forward way of finding the shape descriptors 'responsible' for a given category would be to use well known statistical methods and find through them such descriptors with which we are able to classify shapes according to a given category. When a function is dominating this approach might not suffice - we might be unable to find a set descriptors which are independent of a given function. We show how to decouple the effect of the style from that of the function by considering the shapes of the same function but different styles. We also propose a metric coanalysis approach: if two styles are similar this similarity should be reflected across different functions.

We show the usability of our methods first on the example of a number of chess sets which our method helps sort. Next, we investigate the problem of finding a replacement for a missing tooth given a database of teeth.

1 Introduction

While digital shapes are starting to have a number of medical applications, for instance related to hearing aid production and dental work, the use of digital shapes does not necessarily lead to complete automation. Typically, certain procedures are still left to human operators. However, it is an important goal to be able to help the human operator as much as possible. The particular scenario which we address in this paper is the selection of tooth shapes which can serve as the starting point for digital models of crowns.

There is a lot of work in the shape analysis and especially shape retrieval community with a task of finding the most similar shape to a query one. However, many shapes might be classified not only according to a single category, e.g. as being a table or a chair, which we will call the function, but also according to the style: A table and a chair of the same style share common geometric properties which are different from the overall shape. The style and the function both

interact and contribute to the overall shape of the object. It is not always easy to separate them and point out geometric elements responsible for a function or a style.

The general distinction between the specific shape properties, which tells which ones are responsible for the style and which for the function, is not possible as this depends on a context. That is why we define the style and function through examples.

1.1 Existing Work related to style function recognition

Style and function separation in the context of man made three dimensional shapes was recently mentioned by Xu et al. [14], where the style of an object is defined by the proportions (anisotropic scaling) of its parts. It seems to be very intuitive and reasonable approach but this does not exhaust the subject. The style might be hidden in details, repetitions of some patterns or some other types of deformation as well. Very often it is hard to define it mathematically although the human brains usually do not have problems in recognizing it.

In many shape processing articles, even if the problem of style is not addressed in an explicit way there are situations where the space of given shapes is broken into two different independent classification systems. In the deformation transfer [10] different kinds of animals can take similar poses in which case it is quite easy to localize them, as the type of animal is described by an intrinsic metric of the shape surface, and the pose is its embedding in three dimensional space. The idea of geometric texture [1] fits within this framework as it aims to separate overall shape from its geometric details. Application of example based priors for surface reconstruction [9, 3] can also be seen as imposing style of the object.

In the image processing field Hertzmann et al. [6] presented a method that given three images, an image with style A and function 1, an image with style B and function 1, an image with style A and function 2, created an image with style B and function 2. The same concept was also explored by this group in the field of curve styles [7]. Other related problems can be present when dealing with images of fonts, separating lighting conditions from the scene and distinguishing between the spoken language and the accent - all of those three cases were examined through bilinear models by Tenenbaum et al. [12].

Tenenbaum's framework requires establishing one to one correspondences of the parts both for the style and with the function - for example fonts are compared through pixels of a bitmap: in general for different types of shapes obtaining such correspondences is usually hard to achieve. Similar correspondences need to be established across the styles for Hertzman's work. Our approach does not require any correspondence finding, which usually is a costly task and sometimes it is not possible as for example in the problem of registering a table to a chair. Instead we do shape comparisons through the shape descriptors. There is a lot of current work on content based shape retrieval and different descriptors might capture different properties of the shape and produce different notion of their similarity. So a good approach is to extract many different shape descriptors and combine them in a proper way.

1.2 Metric learning

If the feature space is available, many well established statistical methods can be used such as Linear Discriminant Analysis [8] which modifies the feature space so that, for a given training set containing objects from different classes, it maximizes intra class variance and minimizes within class variance. Similar approach was also used by [13] which gives the possibility of defining the similarity and dissimilarity relationships between selected pairs of objects.

As mentioned by Giorgi et al. [5] for the case of shapes there are many useful shape descriptors like skeletons, trees, weighted point sets, which do not provide multidimensional feature space. Still with such descriptors there is usually a way of establishing a notion of similarities between different shapes which results in some kind of pseudodistance.

Giorgi et al. [5] customize a way of combining a set of distances between shapes so that user defined similarity is captured. In this work the metric is modified in order to reflect the user defined constraints of nearby or far away shapes. The final metric is taken as a maximum distance from distances given by all of the metrics, however the particular metrics are scaled according to a similarity feedback provided by the user.

The approach of combining different metrics relies on the fact that at least there exists a set of shape descriptors which can capture the similarity imposed by virtue of shared stylistic or functional properties. For function, which usually is easier to distinguish such an approach would be very suitable. However when a style needs to be extracted it might not be enough and not even single descriptor might exist which is purely responsible just for the style.

One of our main observations concerning this problem is that knowing what is the function of an object enhances the possibilities for style recognition. For many descriptors information on style is coupled with information on function. In general, when the distance between two shapes is small, it might be both due to similarity in the style and similarity in the function. The retrieval of style related information can be achieved when providing a set of shapes sharing the function and having different styles.

The requirement of recognizing the object of the same function or the same style as being close is not enough in such case. We also want our dissimilarity measures between shapes to be consistent across different functions. This requirement stems from the fact that we want to be able to find the most similar styles and most similar functions. However, for our style-function task case we do not have a direct input which indicates which styles are similar and which are not. Instead we have some notions of similarities which are induced by different shape descriptors and there is a need to choose the ones which are relevant. This relevance is not defined directly by indicating the shapes which should be treated as similar but indirectly as a consistency requirement: dissimilarity or similarity between the styles should be reflected in a similar way for different functions.

Similar indirect consistency approach methodology can be found in [15] which removes incorrect mappings of sets of different views. The assessment of the quality view mappings is done through analyzing them in broader context of

the consistent mapping loops. If the loop is inconsistent it means that one of the mappings that belongs to it is wrong and the consistent loop means that mappings are likely to be correct. Having evaluated the correctness of many loops the bad mappings are spotted through a loopy belief propagation.

1.3 Contribution

This paper focuses on an issue, which we think has many application areas, but was not very much explored yet: the analysis and classification of shapes according to more than one category, when categories may be coupled together which in our case is the style and the function.

We propose here a general methodology which can be applied in order to deal with the style-function determination problem. Because the style and the function strongly depend on the context, defining it by providing example shapes seems to be the most general approach.

We show the method for decoupling the effect of the style from that of the function. By having as a training dataset the shapes of the same function and different styles, we can factor out the function and determine the most likely style of an unknown shape as the closest shape from the set. In an analogous way by using the shapes of the same style but different functions the unknown function may be retrieved.

We realize that the key to success is to find a good metric between the shapes: metric which can capture both stylistic and functional features. Using the example of chess pieces we show what are the desired properties of such a metric (section 2) and how to decouple the style from the function when only one metric is available.

We also show how to find an appropriate metric by combining the metrics obtained through different shape descriptors (section 3). Novel in our case is that we do not only use standard similarity notions but also explore the metric consistency approach. The problem is illustrated with the example for a tooth dataset.

After the example of chess pieces, we focus on teeth as an example medical application. Note that our framework is fairly generic. It could be applied to any type of biological surface which exhibits variation due to both style and function.

2 Decoupling metric

In this section we will show how an information about style and function hidden in the same metric can be decoupled. This is illustrated by the example of style - function classification based on the chess pieces. Since the chess pieces are rotationally symmetric, their three dimensional representation can be reduced to the space of plane curves by taking the outline curve obtained through rotating the chess piece by the rotational symmetry axis. The Translation Invariant Dynamic Time Warping [2] is used in order to establish a similarity metric $d(.,.)$ between the objects.

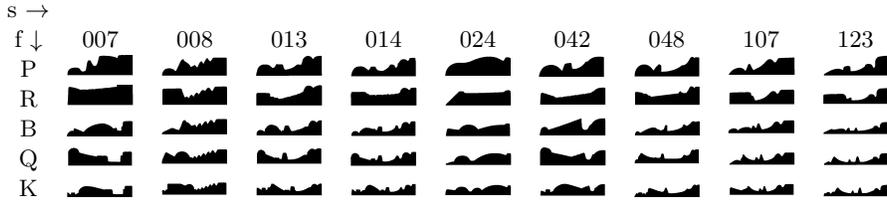


Table 1. The outline curves of the chess pieces. Our dataset has 45 chess pieces, which are the scans taken from 9 existing chess sets. The function is the type of the chess piece (pawn, rook, bishop, king, queen) and the style is the set the chess piece belongs to.

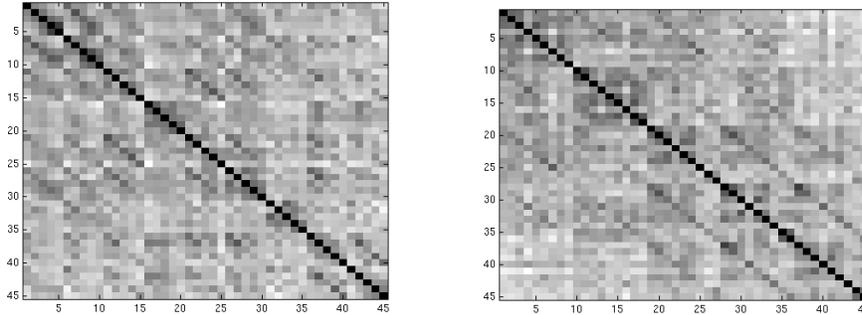


Fig. 1. Similarities between the curves. Each block has the same function or style, the diagonals of blocks are darker which reflects the smaller distance when the function or the style is the same.

2.1 Likelihoods computation

In our setup the proximity of two shapes can be affected by two factors: the similarity of the style and the similarity of the function. Also dissimilarity with respect to one factor, which usually is a style might be more subtle than the other one. However if we have in the training set pieces which share the same style (or function) but have different function (style), then it is possible that we may factor the style (function) out. Instead of taking absolute distances one may use the relative distance information: the difference of the distances. For example if the distance to a king is smaller than a distance to a bishop of the same style we may say that the unknown shape is more likely to be the king than to be something else and that will affect the sign of the distance difference.

The partial likelihood of unknown shape x to be a function K , when we have two example shapes of the same style S_i of which one (denoted as KS_i) is of a function K and other one NS_i is of a function other than K , is equal to:

$$l_f^{S_i, N}(x, K) = d(NS_i, x) - d(KS_i, x).$$

In our training dataset TS_i , for a given style S_i , we may have more than just one shape not being of a function K so we take the mean plus the minimum of all of the partial likelihoods:

$$l_f^{S_i}(x, K) = \text{mean}_{K \neq NS_i \in TS_i} l_f^{S_i, N}(x, K) + \min_{K \neq NS_i \in TS_i} l_f^{S_i, N}(x, K).$$

Note that minimum is equal to the distance to the closest of the known shapes from style S_i other than KS_i , minus the distance to KS_i . If the function K is the closest of the shapes from that style, then the minimum will be positive otherwise it will be negative. The mean value stabilizes the results by taking into account distance measures of all of the shapes of this style.

In order to gather the information from all of the training styles we take the mean value plus the maximum of all the styles, for which in a training set there is a function K and some shapes not being of function K .

$$l_f(x, K) = \max_{\substack{S_i \in S, KS_i \in TS_i, \\ K \neq NS_i \in TS_i}} l_f^{S_i}(x, K) + \text{mean}_{\substack{S_i \in S, KS_i \in TS_i, \\ K \neq NS_i \in TS_i}} l_f^{S_i}(x, K).$$

Here by taking the maximum we are favoring the style for which the K function is most likely. The mean is again added to get the distance information from all known styles.

There might be cases when we do not have enough information in the training set for establishing likelihoods. This happens when there is no set which has a training representative for the function K and for some shape which is not of a function K . In such a case we set the likelihood to zero.

The likelihood computation of "x being the style i" is done in an analogous manner. Then for a given x the cost of assigning to it style j and function i is equal to: $l(x, F_i, S_j) = l_f(x, F_i) + l_s(x, S_j)$.

2.2 Chess classification example with the assignment problem

We use the likelihoods as negative costs and solve the minimum linear assignment problem for the unknown labels and loose chess pieces.

Table 2 contains the results of the assignment problem if the training dataset is one set and one function, and we are searching for other chess pieces. The results depend a lot on the type of the set and function imposed as an example shape. Some of the sets contain a lot style and function information but some other do not. The sets 024 and 008 are performing the worst also the results for the rooks is always worse than for other functions.

2.3 Multiple step assignment

An assignment problem with the costs defined above does not make use of the information about all of the distances between the shapes. If we are able to locate the chess pieces of which we can expect that the initial matching went correctly we can add those into a training dataset with the labels obtained by the

	P	R	B	Q	K	mean
123	9	24	12	17	18	18
f	13	11	5	9	10	9.6
s	13	21	9	10	10	12.6
107	17	20	16	12	17	16.4
f	13	12	10	5	9	9.8
s	12	15	11	7	10	11
014	20	20	10	9	14	14.6
f	6	11	2	2	6	5.4
s	17	13	8	7	10	11
024	28	27	19	25	28	25.4
f	20	22	15	20	22	19.8
s	16	12	8	12	14	12.4
013	14	24	6	6	19	13.8
f	7	11	2	2	8	6
s	8	15	4	4	13	8.8
048	16	17	11	10	17	14.2
f	7	6	4	2	8	5.4
s	14	12	8	8	12	10.8
008	22	24	19	16	24	21
f	18	21	15	10	17	16.2
s	14	18	9	7	11	11.8
042	20	24	21	14	18	19.4
f	14	14	13	9	10	12
s	11	16	9	7	12	11
007	15	23	8	10	21	15.4
f	8	12	6	7	12	9
s	14	15	4	6	16	11
mean	19	22.56	13.56	13.22	19.56	17.58
f	11.78	13.33	8	7.33	11.33	10.36
s	13.22	15.22	7.78	7.56	12	11.16

Table 2. Mismatches of the single assignment problem with one style and one function given. The table contains the general number of pieces with mismatched total label, mismatched function and the mismatched style.

initial assignment. In order to estimate the labeling reliability, we calculate the diagonal cost of the assignment which we define as the average sum of similarities between all the pieces having the same style or function labels. For a hypothetical unknown chess piece we might add it for a moment to the training set and calculate what is the diagonal cost when assignment is solved with the use of this piece. We discovered that instead of calculating diagonal costs directly it is better do the inverse assignment, which is performed by swapping the unknown data with the known and then calculating the diagonal cost. Then the smaller the inverse diagonal cost is the more reliable is the hypothetical assignment of the unknown chess piece to its label.

In order to minimize bad choices we always take the piece having minimum inverse diagonal cost and is reliable according to the additional reliability criteria. We add it to the initial dataset and repeat the assignment and the most reliable pieces addition until there is no reliable piece to be added. Then we use the assignment from the last step as the final assignment.

In the results (table 3) we observe an average improvement of the assignment tasks by approximately 3 chess pieces. Usually if initial guess is quite good but not perfect then correctness of the matching may be improved quite well. If there are too many mismatches the improvement does not occur: as then we also take as reliable the matchings which are not correct. Usually it does not make the solution worse but keeps it at a similar level as it was with the initial problem.

3 Finding the good metric

The case of chess pieces was special problem as we were able to reduce the shape information to the space of the curves and had a way of establishing

	P	R	B	Q	K	mean
123	19	23	9	7	15	14.6
f	11	11	4	2	9	7.4
s	10	18	5	5	10	9.6
107	15	14	10	5	11	11
f	6	7	6	0	5	4.8
s	13	13	5	5	8	8.8
014	14	12	5	4	17	10.4
f	8	8	0	0	9	5
s	9	9	5	4	11	7.6
024	24	25	18	26	27	24
f	20	18	14	20	20	18.4
s	14	18	7	12	11	12.4
013	15	21	0	0	19	11
f	8	11	0	0	11	6
s	9	14	0	0	12	7
048	19	19	2	4	12	11.2
f	11	9	2	0	5	5.4
s	14	14	0	4	9	8.2
008	22	28	23	9	8	18
f	18	22	14	4	7	13
s	13	25	14	6	5	12.6
042	22	21	15	17	23	19.6
f	13	14	10	15	13	13
s	14	15	7	2	16	10.8
007	17	25	6	8	11	13.4
f	8	9	2	2	6	5.4
s	13	19	4	6	5	9.4
mean	18.56	20.89	9.78	8.89	15.89	14.8
f	11.44	12.11	5.78	4.78	9.44	8.71
s	12.11	16.11	5.22	4.89	9.67	9.6

Table 3. Mismatches of the multiple assignment problem with one style and one function given. The table contains the general number of pieces with the mismatched total label, mismatched function and mismatched style.

similarity between those curves by using Translation Invariant Dynamic Time Warping. In general for three dimensional shapes we do not know a good metric in advance, instead we have many propositions of metrics $d_i(\cdot)$ which can be obtained through different kinds of shape descriptors D_i .

The task is to choose such a metric d_i or some combinations of metrics with which we can distinguish between different styles. As mentioned in the introduction it is not enough to be able for the objects of the same style to be close but also the dissimilarity measures should be consistent across different functions and we don't know which one will work best for specific problem.

This requirement can be illustrated with the problem of tooth shapes. Suppose a patient has one tooth destroyed. In order to be able to reproduce its shape, we want to find from a database a tooth which is mostly similar to the existing tooth he has. We have a molar missing but because a premolar is still in the patient's mouth, we wish to search in our database for a mouth which has the most similar premolar to the patient's. From that mouth we take a molar as a template for our new tooth. This approach assumes that similarity for premolars induces a similarity between molars.

This case shows that the metric consistency requirement is necessary as it aids in many concrete tasks - like searching for the closest to missing data. Here we do not know directly what 'close' means, as we have many metric but don't know which one is a correct. Usually a correct metric combination in such a case can be found by giving example pairs of shapes which are similar and which are dissimilar[5]. In our case we do not have such information. Instead we can impose the metric consistency requirement: the distances between shapes having different styles and function A should be close to the distances of the shapes of the same styles and function B .

3.1 Metric consistency

Assume we have a set of training shapes $F_{i=1..n_i}S_{j=1..n_j}$, where i indicates the function and j style. We also have a k_n potential distances $d_k(,)$

Let us take all distances $d_k(F_{i_1}S_{j_a=1..n_j}, F_{i_1}S_{j_b=1..n_j, j_b \neq j_a})$ between different shapes of the function i_1 . In order to be comparable those distances need to be normalized which we do by dividing them by the median from obtained distances. This results in a $\binom{n_j}{2}$ dimensional vector of k-distances between shapes with function i_1 which we will denote $v(d_k, f_{i_1})$.

For each pair $i_1 \neq i_2$ of two different functions we can establish the **consistency score** $cs_{d_k}^{i_1, i_2}$ with respect to a distance k and function i_1 and i_2 as a norm of difference of distance vectors:

$$cs_{d_k}^{i_1, i_2} = \sqrt{\sum_{l=1..\binom{n_j}{2}} (v(d_k, f_{i_1})_l - v(d_k, f_{i_2})_l)^2}$$

In order to calculate total consistency factor (TCF_k) for a distance measure k sum of the differences for all function pairs is taken. Note that smaller TCF_k is the more consistent is d_k with respect to style.

We construct the final metrics by summing the metric obtained through different shape descriptors with weights that promote consistency.

$$D_f(,) = \sum_k e^{(-2 \frac{TCF_k}{mean(TCF)})} \frac{d_k(,)}{\sigma_{d_k}}$$

where σ_{d_k} median distance from distances $d_k(,)$ between all training shapes.

From the final metrics we can also compute consistency scores $cs_{D_f}^{i_1, i_2}$. This consistency measures can be used in order to asses what kind of tooth types are better correlated. For example two neighbor upper molars can be more correlated than molar and incisor. So if a molar is missing and we have the neighbor molar and incisor, we should give higher weight for query of closest mouth with respect to a molar than with respect to incisor. We can also compute mean distances between styles by summing $v(D_f, f_j)$ for all function types j .

3.2 The tooth problem



Fig. 2. Front view of molar, premolar and incisor from 3 different mouths

In the teeth analysis task we take a type of a mouth as style and a tooth type as function. An example dataset we use for this problems contains teeth shapes

(figure 2) from 6 different mouths. In order to make number of styles larger, we assume that the left side of a mouth will be treated separately from the right part. Thus we have 12 styles which we will label as A,B,C,D,E,F,a,b,c,d,e,f, where big letter means one left part of a mouth and small the other one. We have taken 10 tooth types 2 upper molars, lower molar, 2 upper premolars, lower premolar, upper canine, upper incisor, 2 down incisors. They are labeled and placed in the following order: 7M,6M,6m,5P,4P,4p,3C,1I,1i,2i, where upper case means respectively upper tooth.

In order to get independence of meshing we uniformly sampled the surface of teeth and computed descriptors out of those samples. We used local shape descriptors which rely on neighborhood at some distance from a given position. As neighborhood size we have taken 0.01 0.04 0.16 and 0.64 of the radius of a bounding sphere of a tooth. For slippage we used 0.01 0.04 and 0.16. In Total we had: 2x4 descriptors for main curvatures obtained by fitting primitives [11], 3x4 eigenvalues of covariance matrix of points sampled from the neighborhood area and 12x3 slippage coefficients [4] which are 6 eigenvalues of slippage covariance matrix and 6 is a translational contribution to its eigenvectors. We took 2 samples for 1000 points, for which soft histograms were computed. Histograms from two independent sampling were compared. The mean across all training shapes, of their difference was taken in order to estimate the measure error coming from different samplings. Then the mean of the 2 sample histogram is taken. However in order to compare two histograms for shapes S_i and S_j the distance between two bins is reduced by the previously computed measure error. Then the sum of those values is taken across all bins as our distance $d_k(S_i, S_j)$.

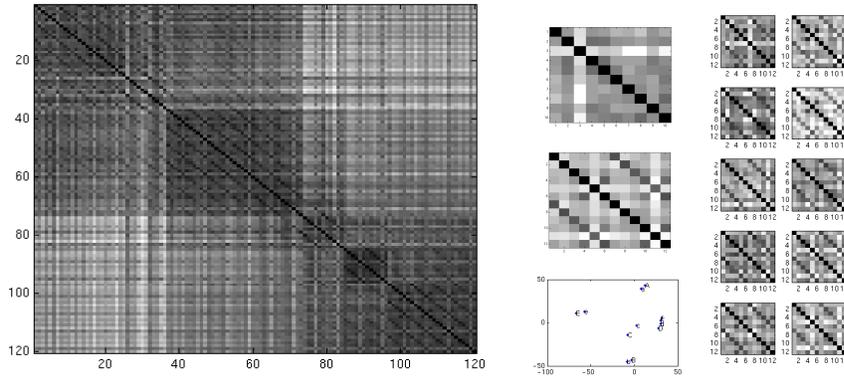


Fig. 3. Left: final metric D_f obtained with teeth database (indices grouped with respect to teeth type). Center: $cs^{i,j}$ for different teeth types, the average distance between the styles, and the multidimensional scaling plot for the average distances. Right: metrics between different styles when a function is fixed, obtained form final metric with the training styles $AFade$.

Then the total consistency factors are computed as mentioned in the previous section and the final metric is computed. Figure 3 contains the resulting metric, where all of the available teeth were used. It is worth mentioning that the consistency score for a resulting metric is smaller than the scores from any particular metrics. Note that the styles that come from the same mouth (left or right part) are being found as close. Also note that neighbor teeth tend to have more consistent scores. This information might be used when searching for a missing tooth. Let us consider the case when *AFade* styles were taken as training styles and a metric T was created. Then a patient comes with mouth of style C and with missing $4P$. We have scans of his $4p$, $5P$ and $1i$ and we have $cs_T^{4P,4p} = 1.6391$, $cs_T^{4P,5P} = 1.5819$ and $cs_T^{4P,1i} = 1.7543$, so we use tooth $5P$ as it has the best consistency. We use $v(T, f_{5P})$ instead of unknown $v(T, f_{4P})$ in order to evaluate proximity between teeth (5th and 6th plot on the Right of figure 3). We evaluate distances between C and *AFade* among teeth of type $5P$ and the mouths sorted with respect to distance will be *FdaAE* if we checked the ground truth we have *dFaEA*. Despite the swaps which was a result of a very close similarity values of dF and EA we can see that in general dissimilar teeth remain dissimilar.

We also tested on how the consistency properties of metric change when different subset of styles was used as training dataset. We generated metric from this information and evaluated the results on all of the data.

Usually removing only small number of mouths did not increase or even slightly decrease the consistency scores. Only when using 3 or 4 mouths, the results seemed be different. This might come from the fact that there was always some symmetric tooth left in the set which was able to set the consistency scores in a correct way. The increase was mostly noticeable when styles which are close to each other are used as training set (table 4).

training	TCS	training	TCS	training	TCS	training	TCS	training	TCS
<i>DdFf</i>	111.007	<i>EFb</i>	102.088	<i>cEF</i>	98.17	all	97.94	none	117.22
<i>Eef</i>	117.107	<i>ABd</i>	99.138	<i>ADEFc</i>	97.35	<i>AFade</i>	98.494	<i>ABCDFaf</i>	95.929

Table 4. Total consistency factors when using different mouth subsets as training data.

4 Conclusion

In this article we presented methods of working with shapes that can be classified into having two categories: style and function. One of them decouples style and function when they are incorporated into the same metric. The second finds a metric as a combination from existing ones when a consistency between different function types is needed. Those methods were illustrated by the chess and tooth datasets. We are aware that for a further analysis and development of our methods more data will be needed but we think the results obtained so far are promising.

Acknowledgments: We would like to thank 3Shape for the tooth dataset and for making their scanner available when scanning the chess pieces.

References

1. Vedrana Andersen, Mathieu Desbrun, J. Andreas Baerentzen, and Henrik Aanaes. Height and tilt geometric texture. In *Proceedings of the 5th International Symposium on Advances in Visual Computing: Part I*, ISVC '09, pages 656–667, Berlin, Heidelberg, 2009. Springer-Verlag.
2. Alon Efrat, Quanfu Fan, and Suresh Venkatasubramanian. Curve matching, time warping, and light fields: New algorithms for computing similarity between curves. *J. Math. Imaging Vis.*, 27(3):203–216, 2007.
3. Ran Gal, Ariel Shamir, Tal Hassner, Mark Pauly, and Daniel Cohen-Or. Surface reconstruction using local shape priors. In *SGP '07: Proceedings of the fifth Eurographics symposium on Geometry processing*, pages 253–262, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.
4. N. Gelfand and L. Guibas. Shape segmentation using local slippage analysis. In *Proc. Symp. Geom. Processing*, pages 219–228, 2004.
5. D. Giorgi, P. Frosini, M. Spagnuolo, and B. Falcidieno. 3d relevance feedback via multilevel relevance judgements. *Vis. Comput.*, 26:1321–1338, October 2010.
6. Aaron Hertzmann, Charles E. Jacobs, Nuria Oliver, Brian Curless, and David H. Salesin. Image analogies. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 327–340, New York, NY, USA, 2001. ACM.
7. Aaron Hertzmann, Nuria Oliver, Brian Curless, and Steven M. Seitz. Curve analogies. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering*, pages 233–246, Aire-la-Ville, Switzerland, Switzerland, 2002. Eurographics Association.
8. G.J. McLachlan. *Discriminant Analysis and Statistical Pattern Recognition*. Wiley, 1992.
9. Mark Pauly, Niloy J. Mitra, Joachim Giesen, Markus Gross, and Leonidas J. Guibas. Example-based 3d scan completion. In *SGP '05: Proceedings of the third Eurographics symposium on Geometry processing*, page 23, Aire-la-Ville, Switzerland, Switzerland, 2005. Eurographics Association.
10. Robert W. Sumner and Jovan Popović. Deformation transfer for triangle meshes. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*, pages 399–405, New York, NY, USA, 2004. ACM.
11. Gabriel Taubin. Estimating the tensor of curvature of a surface from a polyhedral approximation. pages 902–902. IEEE Computer Society, IEEE Computer Society, 1995.
12. Joshua B. Tenenbaum and William T. Freeman. Separating style and content with bilinear models. *Neural Comput.*, 12:1247–1283, June 2000.
13. Eric P. Xing, Andrew Y. Ng, Michael I. Jordan, and Stuart Russell. Distance metric learning, with application to clustering with side-information. In *Advances in Neural Information Processing Systems 15*, volume 15, pages 505–512, 2002.
14. Kai Xu, Honghua Li, Hao Zhang, Daniel Cohen-Or, Yueshan Xiong, and Zhiqian Cheng. Style-content separation by anisotropic part scales. *ACM Transactions on Graphics, (Proc. of SIGGRAPH Asia 2010)*, 29(5):to appear, 2010.
15. Christopher Zach, Manfred Klopschitz, and Manfred Pollefeys. Disambiguating visual relations using loop constraints. In *CVPR'10*, pages 1426–1433, 2010.

Template Deformation with User Constraints for Live 3D Interactive Surface Extraction

Benoit Mory, Oudom Somphone, Raphael Prevost, and Roberto Ardon

Medisys Research Lab, Philips Healthcare France

Abstract. We describe an algorithm for 3D interactive surface extraction by non-rigid implicit template deformation, with two main original features. First, our formulation incorporates user input as inside/outside labeled points to drive the deformation and improve both robustness and accuracy. This yields inequality constraints, solved using an Augmented Lagrangian approach. Secondly, a fast implementation of non-rigid deformation enables interactions with a live visual feedback of the deformation. We validate this technique in Contrast Enhanced Ultrasound images of kidneys in preoperative radio-frequency ablation planning.

Keywords: Interactive, Segmentation, Surface Extraction, Warping, Deformation

1 Introduction

In medical applications, surface extraction of anatomical structures in difficult conditions such as tissue inhomogeneities, noise, loss-of-contrast, can be significantly facilitated by the incorporation of prior knowledge. This approach has been extensively studied in terms of shape prior by constraining the solution to remain close to a pre-defined shape. For instance, statistical methods have been proposed to model shapes, such as the *active shape models* [1]. In the level-set framework, shape priors have also been used, penalizing the dissimilarity between the implicit object representation and the one embedding the prior shape, *via* an additive shape constraint [2–5]. These two approaches have been combined by embedding training shapes in distance functions and defining a statistical model for the shape term [6–9].

Template-to-image registration is a possible alternative, recently applied to medical applications such as liver segmentation in CT [10], in which segmentation is performed by geometrically deforming a binary template towards the image [10–12]. The prior is the template itself and the shape constraint consists in a regularization of the deformation.

In pathological cases presenting extreme variabilities of image features and organ shapes, shape priors may be helpful but insufficient; expert input is then essential to guide the segmentation. Designing intuitive and reliable interactive tools remain a key challenge, particularly in 3D. Few attempts to combine shape priors and interactivity have been made [13]. Along the same line, our proposed formulation incorporates user input constraints in the form of inside/outside labels. Our fast implementation results in an application that provides live visual feedback of the current non-rigid deformation.

2 Surface Extraction by Implicit Template Deformation

In this work, we formulate surface extraction as the partitioning of an image I to provide the best trade-off between regions homogeneity and boundary regularity. Formally, the solution surface (\mathcal{S}) minimizes an objective criterion, defined as the sum of a regularization term $\mathcal{R}(\mathcal{S})$ and image-dependent homogeneity measures r_1 and r_2 integrated over the inside and outside of \mathcal{S} :

$$\min_{\mathcal{S}} \left\{ \mathcal{R}(\mathcal{S}) + \int_{\text{inside } \mathcal{S}} r_1(\mathbf{x}) d\mathbf{x} + \int_{\text{outside } \mathcal{S}} r_2(\mathbf{x}) d\mathbf{x} \right\} \quad (1)$$

For instance, with known intensity distributions p_i for each region, it is common to set $r_i(\mathbf{x}) = -\log p_i(I(\mathbf{x}))$ [14].

An equivalent formulation can be written using an implicit representation of \mathcal{S} , using a function $\Phi : \Omega \rightarrow \mathbb{R}$, positive inside \mathcal{S} and satisfying $\Phi^{-1}(0) = \mathcal{S}$. Denoting H the Heaviside function ($H(a) = 1$ if $a > 0$, 0 otherwise), $H(\Phi)$ is the characteristic function of the interior of \mathcal{S} , and (1) is equivalent to:

$$\min_{\Phi} \left\{ \mathcal{R}(\Phi) + \int_{\Omega} H(\Phi(\mathbf{x})) r(\mathbf{x}) d\mathbf{x} \right\} \quad \text{with} \quad r(\mathbf{x}) = r_1(\mathbf{x}) - r_2(\mathbf{x}) \quad (2)$$

The regularization $\mathcal{R}(\Phi)$ can be complemented with an additional shape prior term that forces the solution to remain close to a predefined implicit representation (the authors of [2] constrain Φ to be a distance function in order to penalize its L_2 -distance to a globally transformed template). Nevertheless, this technique does not guarantee that the zero level-set of the solution preserves the topology of the shape prior. Moreover, $\mathcal{R}(\Phi)$ is classically chosen to penalize the perimeter of the region boundary, which is counter-productive since it inevitably smooths out possible important details of the prior shape.

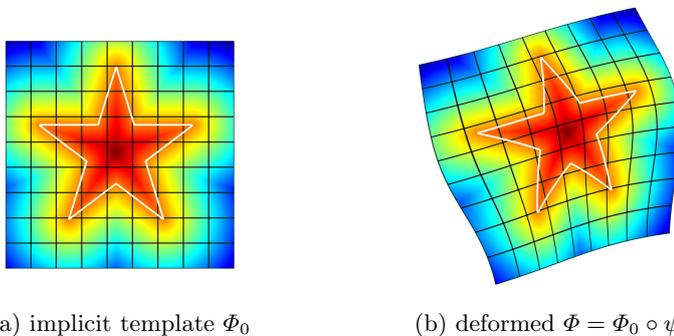


Fig. 1. Implicit template deformation

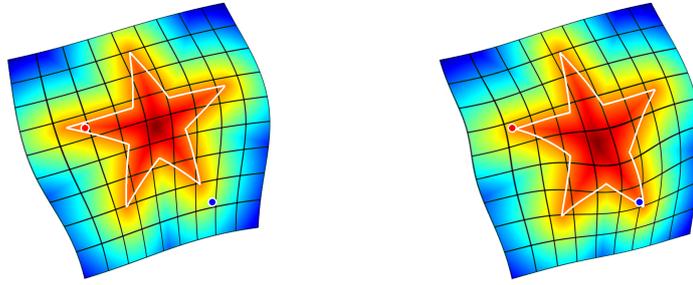
To cope with these problems, alternative approaches have been proposed [10–12] to rather deform a template surface \mathcal{S}_0 defined in a referential Ω_0 with a geometric transformation $\psi : \Omega_0 \rightarrow \Omega$. This template can be implicitly represented by a function Φ_0 such that $\mathcal{S}_0 = \Phi_0^{-1}(0)$ (illustrated in Fig. 1.a). The unknown becomes the transformation ψ and the penalization $\mathcal{R}(\Phi)$ in (2) is substituted with a shape term $\mathcal{R}(\psi)$, consisting in a regularization constraint acting on ψ . The extracted surface is

then obtained as $\mathcal{S} = \Phi^{-1}(0)$ with $\Phi = \Phi_0 \circ \psi$ (see, Fig. 1.b). A general formulation of surface extraction by implicit template deformation then reads:

$$\min_{\psi} \left\{ E(\psi) = \mathcal{R}(\psi) + \int_{\Omega} H(\Phi_0 \circ \psi(\mathbf{x}))r(\mathbf{x}) d\mathbf{x} \right\} \quad (3)$$

Compliance with the shape prior is determined by both the deformation model ψ and its associated constraint $\mathcal{R}(\psi)$. In the non-rigid case, Saddi *et al.* [10] represented the deformation with a diffeomorphic fluid model. In [11], Somphone *et al.* proposed deformations based on finite elements with partition of unity. In [12], Huang and Metaxas adopted Free Form Deformations in their *Metamorphs*.

3 User Interactions, Deformation Model and Shape Term



(a) $\Phi = \Phi_0 \circ \psi$ violating constraints (b) $\Phi = \Phi_0 \circ \psi$ satisfying constraints

Fig. 2. User constraints as inside/outside labeled points

User Interactions. Additional control and robustness can be obtained by allowing the user to indicate whether some specific points shall lie inside or outside the surface to extract. Denoting $\{\mathbf{x}_k \in \Omega\}_{k \in \{0, \dots, N-1\}}$ these labeled points, the user input can be translated into N constraints on the sign of the transformed template $\Phi = \Phi_0 \circ \psi$, at $\{\mathbf{x}_k\}$:

$$\forall k \in \{0, \dots, N-1\} \quad \gamma_k \Phi_0 \circ \psi(\mathbf{x}_k) \geq 0 \quad (4)$$

where $\gamma_k = 1$ (resp. -1) for inside (resp. outside) points. Fig. 2 illustrates the constraint induced by an inside point (in blue) and an outside point (in red) on the deformation of the star-shaped object of Fig. 1.

Putting together criterion (3) and the constraints (4) yields a general formulation of implicit template deformation with user interactions, as the minimization problem

$$\min_{\psi} \left\{ \begin{array}{l} E(\psi) = \mathcal{R}(\psi) + \int_{\Omega} H(\Phi_0 \circ \psi(\mathbf{x}))r(\mathbf{x}) d\mathbf{x} \\ \text{s.t.} \quad \gamma_k \Phi_0 \circ \psi(\mathbf{x}_k) \geq 0, \quad \forall k \in \{0, \dots, N-1\} \end{array} \right\} \quad (5)$$

Note that forcing the resulting surface $\mathcal{S} = \Phi^{-1}(0)$ to pass through a specific point \mathbf{x} can be obtained by adding at this same point both inside/outside constraints.

Transformation The choice of the space of possible solutions ψ to problem (5) is, in our case, intrinsically linked to the notion of *shape*. A *shape* can be considered as a set of objects sharing the same visual aspect. It should be invariant to geometric transforms such as translation, rotation, scaling or shearing. We will refer to such a global transformation as the *pose*. To set up a clear distinction between the pose and the subsequent shape *deformation*, similarly to [15], we design our template transformation model ψ as a *functional composition* of a global transformation (\mathcal{G}) and a non-rigid local transformation (\mathcal{L}):

$$\psi = \mathcal{L} \circ \mathcal{G} \quad (6)$$

Fig. 3 illustrates this composition. The template surface, a sphere (in green in Fig. 3.a), is deformed to reach a target shape, undergoing a global transformation (scaling and translation, in yellow in Fig. 3.b), followed by a non-rigid deformation (in red in Fig. 3.b).

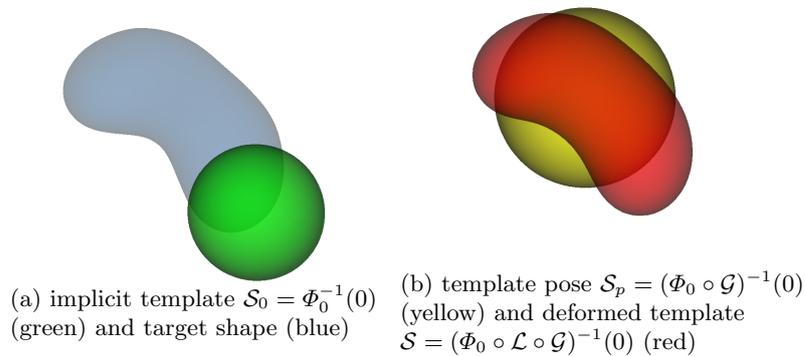


Fig. 3. Transformation decomposition in pose and deformation

Pose. $\mathcal{G} : \Omega \rightarrow \Omega_0$ is chosen as a parametric transform that coarsely aligns the template with the target surface in the image. For anatomical structures in 3D medical images, similarities (which preserve the aspect ratio) are particularly adapted. \mathcal{G} is thus represented by a matrix in homogeneous coordinates defined by 7 parameters $\mathbf{p} = \{p_i\}_{i=1\dots 7}$ and noted $\mathcal{G}_{\mathbf{p}}$.

Deformation. $\mathcal{L} : \Omega_0 \rightarrow \Omega_0$ is expressed using a displacement field \mathbf{u} in the template referential $\mathcal{L} = \mathbf{u} + \mathbf{Id}$. Similarly to problems in image registration and optical flow algorithms [16], \mathbf{u} should be smoothly-varying in space. While adding penalizations on differential terms on \mathbf{u} to $\mathcal{R}(\psi)$ is a valid approach, efficient implementations are difficult to derive. Taking advantage of efficient linear filtering, smoothness of the displacement \mathbf{u} is set as a built-in property defining it as a filtered version of an integrable unknown displacement field \mathbf{v}

$$\mathbf{u}(\mathbf{x}) = [K_\sigma * \mathbf{v}](\mathbf{x}) = \int_{\Omega_0} K_\sigma(\mathbf{x} - \mathbf{y})\mathbf{v}(\mathbf{y})d\mathbf{y} \quad (7)$$

where K_σ is a Gaussian kernel of scale σ .

Shape Term. Decomposing $\psi = \mathcal{L} \circ \mathcal{G}$ allows to define the shape prior term independently from the pose: $\mathcal{R}(\psi) = \mathcal{R}(\mathcal{L})$. \mathcal{R} thus quantifies how much the extracted

surface \mathcal{S} deviates from the prior *shape* \mathcal{S}_0 . Using the L_2 norm we choose to constraint \mathcal{L} towards the identity \mathbf{Id} :

$$\mathcal{R}(\mathcal{L}) = \frac{\lambda}{2} \|\mathcal{L} - \mathbf{Id}\|_2^2 = \frac{\lambda}{2} \int_{\Omega_0} \|\mathbf{u}(\mathbf{x})\|^2 d\mathbf{x} \quad (8)$$

λ is a positive scalar parameter that controls the strength of the shape prior.

The constrained optimization problem to solve finally reads:

$\min_{\mathbf{p}, \mathbf{v}} \left\{ E(\psi_{\mathbf{p}, \mathbf{v}}) = E(\mathbf{p}, \mathbf{v}) = \frac{\lambda}{2} \int_{\Omega_0} \ K_\sigma * \mathbf{v}(\mathbf{x})\ ^2 d\mathbf{x} + \int_{\Omega} H(\Phi_0 \circ \psi_{\mathbf{p}, \mathbf{v}}(\mathbf{x})) r(\mathbf{x}) d\mathbf{x} \right\}$
$\text{s.t} \quad \gamma_k \Phi_0 \circ \psi_{\mathbf{p}, \mathbf{v}}(\mathbf{x}_k) \geq 0, \quad \forall k \in 0..N-1$
$\text{with} \quad \psi_{\mathbf{p}, \mathbf{v}} = \mathcal{G}_{\mathbf{p}} + (K_\sigma * \mathbf{v}) \circ \mathcal{G}_{\mathbf{p}}$

(9)

4 Implementation for Live 3D Surface Extraction

Constraints Management Since $E(\psi_{\mathbf{p}, \mathbf{v}})$ is a non-convex functional to be minimized under a set of non-linear constraints, no specifically tailored algorithms are available. For this matter, we follow a general augmented Lagrangian methodology and define an equivalent unconstrained problem [17] that can be locally minimized by gradient descent. Although we adopted $\psi = \psi_{\mathbf{p}, \mathbf{v}}$ as specified in (9), the following applies to a large majority of transformation models.

The constrained problem (9) can equivalently be written as an unconstrained minimization problem of the form

$$\min_{\psi_{\mathbf{p}, \mathbf{v}}} \left\{ \tilde{E}(\psi_{\mathbf{p}, \mathbf{v}}) = \max_{\alpha \geq 0} \left\{ E(\psi_{\mathbf{p}, \mathbf{v}}) - \sum_{k=0}^{N-1} \alpha_k c_k(\psi_{\mathbf{p}, \mathbf{v}}) \right\} \right\} \quad \text{with } c_k(\psi_{\mathbf{p}, \mathbf{v}}) = \gamma_k \Phi_0 \circ \psi_{\mathbf{p}, \mathbf{v}}(\mathbf{x}_k) \quad (10)$$

where α_k is the Lagrange multiplier associated to the k^{th} constraint. (10) has the same set of solutions as the original problem (9): if $\psi_{\mathbf{p}, \mathbf{v}}$ satisfies all the constraints c_k ($\psi_{\mathbf{p}, \mathbf{v}}$ is feasible), then $\tilde{E}(\psi_{\mathbf{p}, \mathbf{v}}) = E(\psi_{\mathbf{p}, \mathbf{v}})$, otherwise $\tilde{E}(\psi_{\mathbf{p}, \mathbf{v}})$ is infinite. Since \tilde{E} jumps from finite to infinite values at the boundary of the feasible set, a more practical minimization requires to introduce a smooth approximation \hat{E} . Within an iterative minimization process, in order to constrain the maximizers $\alpha = \{\alpha_k\}_{k=0, \dots, N-1}$ to finite values, one has to explicitly introduce at each iteration a quadratic penalty parameter μ and a set of Lagrange multipliers α^j (at the j^{th} iteration) to define

$$\hat{E}_\mu(\psi_{\mathbf{p}, \mathbf{v}}, \alpha^j) = \max_{\alpha \geq 0} \left\{ E(\psi_{\mathbf{p}, \mathbf{v}}) - \sum_{k=0}^{N-1} \alpha_k c_k(\psi_{\mathbf{p}, \mathbf{v}}) - \frac{1}{2\mu} \sum_{k=0}^{N-1} (\alpha_k - \alpha_k^j)^2 \right\} \quad (11)$$

The maximizing Lagrange multipliers associated to each constraint $c_k(\psi_{\mathbf{p}, \mathbf{v}})$ can then be found as functions of previously estimated values:

$$\alpha_k^{j+1} = \begin{cases} 0 & \text{if } \alpha_k^j - \mu c_k(\psi_{\mathbf{p}, \mathbf{v}}) \leq 0 \\ \alpha_k^j - \mu c_k(\psi_{\mathbf{p}, \mathbf{v}}) & \text{otherwise.} \end{cases} \quad (12)$$

Substituting (12) in (11) yields the expression of the smooth approximation \hat{E} :

$$\hat{E}_\mu(\psi_{\mathbf{p},\mathbf{v}}, \boldsymbol{\alpha}^j) = E(\psi_{\mathbf{p},\mathbf{v}}) + \sum_{k=0}^{N-1} \Psi_\mu \left(c_k(\psi_{\mathbf{p},\mathbf{v}}), \alpha_k^j \right) \quad (13)$$

$$\text{with } \Psi_\mu(a, b) = \begin{cases} -ab + \frac{\mu}{2}a^2 & \text{if } \mu a \leq b \\ -\frac{1}{2\mu}b^2 & \text{otherwise.} \end{cases} \quad (14)$$

The alternate scheme described in Algorithm 1, in which the penalty parameter μ is gradually increased, will provide a local minimizer of (9) that eventually satisfies the user constraints.

Algorithm 1: Augmented Lagrangian Scheme For Inequality Constraints

given starting penalty parameter μ^0 , and $\boldsymbol{\alpha}^0 = 0$,

repeat

 choose $\mu^t > \mu^{t-1}$,

repeat

 (1) $\psi_{\mathbf{p},\mathbf{v}}$ being fixed, update the Lagrange multipliers $\boldsymbol{\alpha}^{j+1}$ using (12)

 (2) $\boldsymbol{\alpha}^j$ being fixed, update $\psi_{\mathbf{p},\mathbf{v}}$ by minimizing $\hat{E}_{\mu^t}(\psi_{\mathbf{p},\mathbf{v}}, \boldsymbol{\alpha}^j)$ using (13)

until convergence;

until a local minimum of $E(\psi_{\mathbf{p},\mathbf{v}})$ satisfying $\forall k, c_k(\psi_{\mathbf{p},\mathbf{v}}) \geq 0$ is found;

Gradient Descent An important feature of our 3D template deformation application is its interactivity. Therefore, iterations of the gradient descent chosen to minimize the smooth approximation $\hat{E}_\mu(\psi_{\mathbf{p},\mathbf{v}}) = \hat{E}(\mathbf{p}, \mathbf{v})$ in (13) should be fast enough to provide a real-time display of the surface evolution.

The gradient descent evolution equations are obtained by applying standard calculus of variations; recall that

$$\hat{E}(\mathbf{p}, \mathbf{v}) = E(\mathbf{p}, \mathbf{v}) + \sum_{k=0}^{N-1} \Psi_\mu \left(c_k(p_i, \mathbf{v}), \alpha_k^j \right) \Rightarrow \begin{cases} \frac{\partial p_i}{\partial t} = -\frac{\partial \hat{E}}{\partial p_i} \\ \frac{\partial \mathbf{v}}{\partial t} = -\frac{\partial \hat{E}}{\partial \mathbf{v}} \end{cases} \quad (15)$$

By setting

$$\mathcal{A}_i(\mathbf{x}) = \left\langle \nabla \Phi_0 \circ \mathcal{L}(\mathbf{x}), (\mathbf{I} + \mathbf{J}_{\mathbf{u}} \circ \mathcal{G}) \frac{\partial \mathcal{G}}{\partial p_i} \circ \mathcal{G}^{-1}(\mathbf{x}) \right\rangle \quad (16)$$

and $b_k = \gamma_k \frac{\partial \Psi_\mu}{\partial a} \left(c_k, \alpha_k^j \right)$

where \mathbf{I} is the Identity matrix, $\mathbf{J}_{\mathbf{u}}$ is the Jacobian matrix of \mathbf{u} , the final expression of the evolution equations for p_i and \mathbf{v} reads

$$\begin{cases} \frac{\partial p_i}{\partial t} = & - \int_{\Omega_0} \delta(\Phi_0 \circ \mathcal{L})(r \circ \mathcal{G}^{-1} \mathcal{A}_i) - \sum_{k=0}^{N-1} b_k \mathcal{A}_i \circ \mathcal{G}(\mathbf{x}_k) \\ \frac{\partial \mathbf{v}}{\partial t} = -K_\sigma * \left[\begin{array}{l} \lambda \mathbf{u} \\ \text{shape} \end{array} + \left(\underbrace{\delta(\Phi_0 \circ \mathcal{L}) r \circ \mathcal{G}^{-1}}_{\text{image force}} + \underbrace{\sum_{k=0}^{N-1} b_k \delta_{\mathcal{G}(\mathbf{x}_k)}}_{\text{constraints}} \right) \nabla \Phi_0 \circ \mathcal{L} \right] \end{cases} \quad (17)$$

A quick analysis of (17) reveals several key aspects for an efficient implementation. Interpolating $\Phi_0 \circ \mathcal{L}$ and $\nabla \Phi_0 \circ \mathcal{L}$ over the whole domain Ω_0 would be extremely time-consuming. Nevertheless, since it is multiplied either by $\delta(\Phi_0 \circ \mathcal{L})$ or $\delta_{\mathcal{G}(\mathbf{x}_k)}$, the warped gradient field $\nabla \Phi_0 \circ \mathcal{L}$ is only needed on the set $\{\Phi_0 \circ \mathcal{L} = 0\}$ and at a limited number of points $\{\mathbf{x}_k\}$ (Fig. 4.a) which highly reduces computational burden. Moreover, precise knowledge of the warped template $\Phi_0 \circ \mathcal{L}$ is only necessary near its zero level set. Setting Φ_0 to a distance function to the template surface \mathcal{S}_0 allows a coarse-to-fine approach using octrees. At each level a decision is made to further refine the cell depending on the distance measure (Fig. 4.b) drastically dropping complexity. Finally, stemming from the displacement model (7), extrapolating image and pointwise forces to the whole space boils down to a convolution with K_σ (Fig. 4.c). In practice, our current 3D implementation supports up to 100 time steps per second when discretizing Ω_0 with a lattice containing 48^3 points. This execution speed allows a visual feedback of the deforming surface with a live response to constraints.

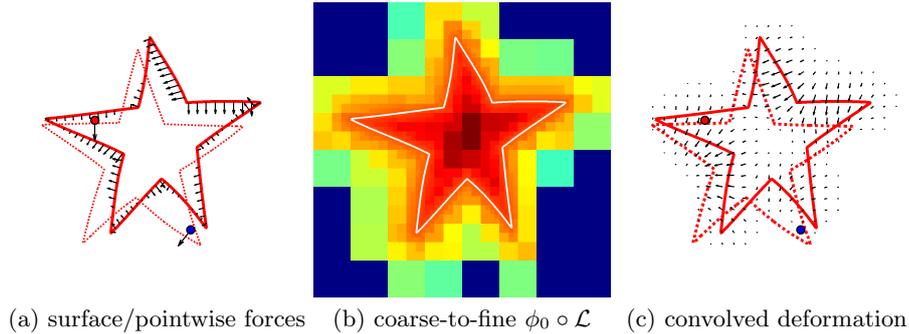


Fig. 4. Fast template deformation with coarse-to-fine distance warp and convolutions

5 Workflow and Validation

The generality of Eq (2) permits a wide field of applications. In medical imaging, going from one application to another is done by redesigning homogeneity measures (r_1 and r_2) and the surface model Φ_0 .

Workflow description An important part of many preoperative planning in liver cancer treatment is the measurement of liver volume in MRI imaging. Large variability in liver shape as well as the presence of important lesions contribute to the difficulty

of the segmentation problem. Our workflow is illustrated through an example in this modality. The simultaneous processing of user interactions and solving of problem (9) enables an intuitive workflow. As illustrated in Figure 5, the user first places the model Φ_0 , which initializes pose \mathcal{G} (Fig 5.a). From this point and further on the optimization algorithm is launched as a background task. With real-time feedback the user observes the surface evolution based on image descriptors (Fig 5.b). At any time constraints can be added by simple clicks (for instance left/right clicks for inside/outside points) and directly integrated into the optimization algorithm, thus guiding the surface extraction (Fig 5.c) with live feedback on the effect of each interaction.

In this application, the surface model Φ_0 is generated from an average mesh built from manual segmentations, while the image-based terms r_1 and r_2 are set to log probability distributions (see Eq (2)). Note that even though image quality is fair, user interactions are necessary for the correct surface extraction given the observed lesion.

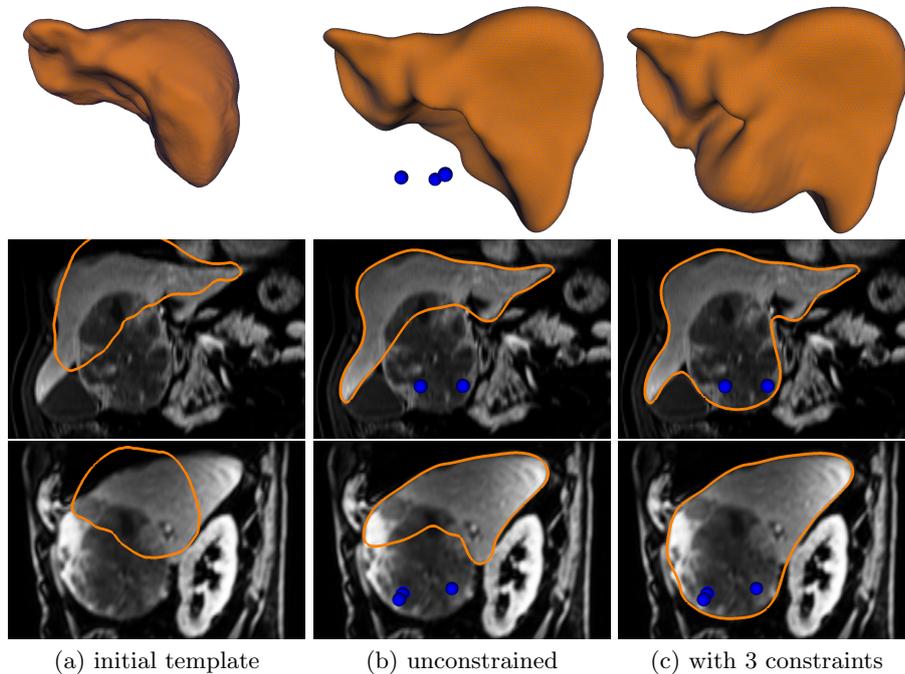


Fig. 5. MRI of the liver for preoperative resection planning. A 3D template of the liver (a) is deformed to extract most of the liver tissue (b). In this case, 3 inside user constraints are required to include a severe and unpredictable protuberant lesion (c).

Validation on Contrast-Enhanced Ultrasounds (CEUS) of kidneys

We now validate our method more deeply on an imaging technique that has proven essential for both its relatively low toxicity and cost, in particular during radio-frequency (RF) ablation planning of kidney tumors. Comparing kidney and tumor volumes (before the operation) with kidney and resection volumes (after the operation) is necessary to evaluate the success of the intervention. However, segmenting the kidney in CEUS images is a particularly difficult task: the presence of contrast agents generates noisy

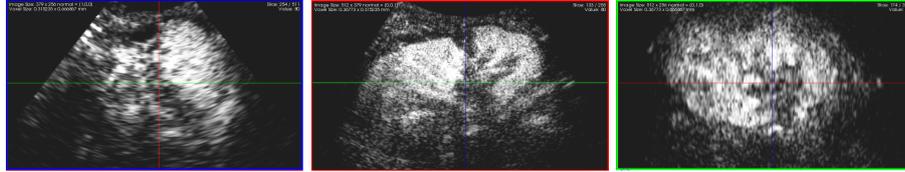


Fig. 6. 3D CEUS image of a kidney, here shown three orthogonal cross-sections (courtesy of Prof. J.M. Correas, Hospital Necker - Paris, France)

data, the limited field of view of probes often prevents the acquisition of the whole kidney and the presence of lesions induces variations in the usual shape (see Figure 6). In collaboration with the radiology department of Hospital Necker (Paris, France) - the main medical center for renal RF ablation in Europe - we have quantified the benefits of our method. For thirteen patients, a specialist was asked to produce a ground truth segmentation to measure the quality of our algorithm as a function of the number of interactions. The model surface Φ_0 was set to an ellipsoid and r_i to $-\log(p_i(I))$, $i = 1, 2$. \mathcal{G} was initialized by the user and deformation \mathcal{L} to identity ($\mathbf{u} = \mathbf{0}$). The graphs in Figures 7.a and 7.b show that without interactions, the minimization of E (see Eq (3)) already considerably reduces the segmentation error. As the user interacts the closer the segmentation gets to the ground truth, often less than three clicks are needed for a satisfactory result (in an average of 15 seconds). The ground truth is never exactly matched due to high intra-operator variability. Figures 7.c and 7.d show two examples where image features, shape prior and user interactions are all essential to the correctness of the final segmentation.

6 Conclusion

Although proven a solid approach for medical image segmentation, template deformation is still unable to cope with the wide spectrum of shapes that pathologies can generate. Providing users with reliable interactions is essential for the general acceptance of these segmentation tools. In this context, we introduced 3D interactions in a template deformation framework as points inside or outside the target anatomical structure. Since the control is even more intuitive with a real-time visualization of the evolution, a special care has been devoted to algorithmic efficiency.

Acknowledgments

The authors wish to thank Prof. Jean-Michel Correas from Hospital Necker (Paris, France) for providing images and clinical advice.

References

1. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models: Their training and application. *CVIU* **61** (1995) 38–59

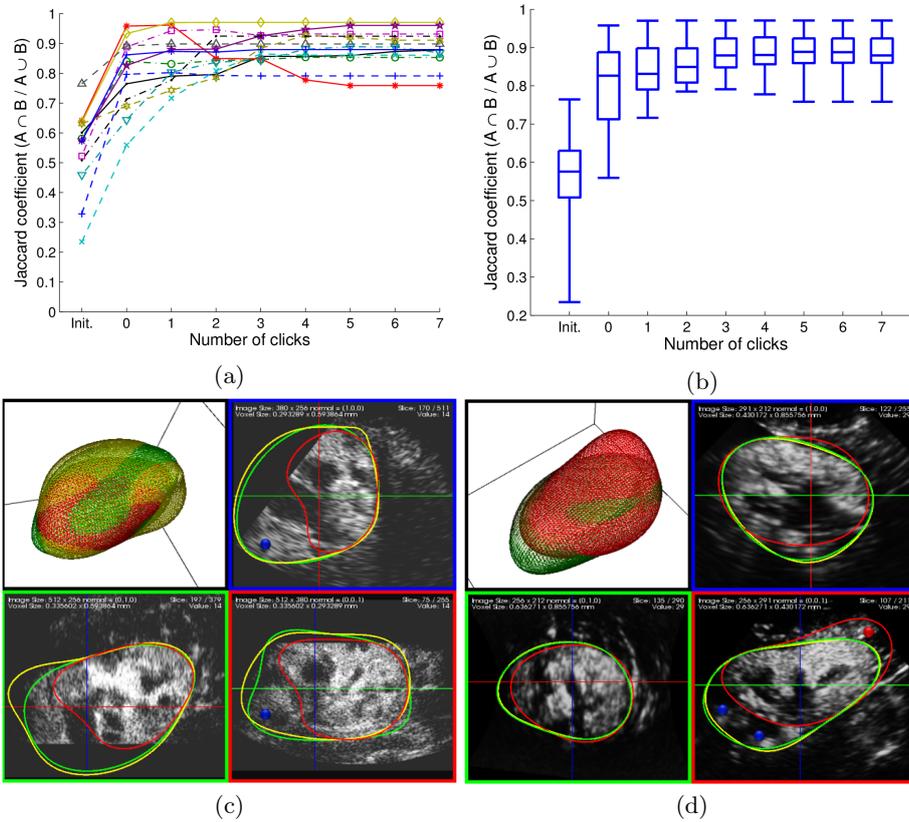


Fig. 7. (a) and (b): Similarity between the different segmentations and the ground truth defined by $J(A, B) = \frac{|A \cap B|}{|A \cup B|}$, as a function of the number of interactions. Left: each curve correspond to a single image. Right: box plot summarizing curves. (c) and (d): Example of two different images, with different segmentations superimposed: ground truth (yellow), deformation result with no constraint (red) and final result (green) with interactions. Blue/red dots indicate points marked as inside/outside.

2. Paragios, N., Rousson, M., Ramesh, V.: Matching distance functions: A shape-to-area variational approach for global-to-local registration. In: ECCV. (2002) II: 775–789
3. Cremers, D., Sochen, N.A., Schnorr, C.: Towards recognition-based variational segmentation using shape priors and dynamic labeling. In: Scale Space. (2003) 388–400
4. Chan, T.F., Zhu, W.: Level set based shape prior segmentation. In: IEEE CVPR. (2005) II: 1164–1170
5. Foulonneau, A., Charbonnier, P., Heitz, F.: Multi-reference shape priors for active contours. IJCV **81** (2009) 68–81
6. Leventon, M.E., Grimson, W.E.L., Faugeras, O.D.: Statistical shape influence in geodesic active contours. In: IEEE CVPR. (2000) 316–323

7. Tsai, A., Yezzi, A.J., Wells, III, W.M., Tempany, C., Tucker, D., Fan, A., Grimson, W.E.L., Willsky, A.S.: A shape-based approach to the segmentation of medical imagery using level sets. *IEEE TMI* **22** (2003) 137–154
8. Bresson, X., Vandergheynst, P., Thiran, J.P.: A variational model for object segmentation using boundary information and shape prior driven by the Mumford-Shah functional. *IJCV* **68** (2006) 145–162
9. Cremers, D., Kohlberger, T., Schnorr, C.: Shape statistics in kernel space for variational image segmentation. *Pattern Recognition* **36** (2003) 1929–1943
10. Saddi, K.A., Rousson, M., Chefd’hotel, C., Cheriet, F.: Global-to-local shape matching for liver segmentation in ct imaging. In: *MICCAI*. (2007)
11. Somphone, O., Mory, B., Makram-Ebeid, S., Cohen, L.D.: Prior-based piecewise-smooth segmentation by template competitive deformation using partitions of unity. In: *ECCV*. (2008)
12. Huang, X., Metaxas, D.: Metamorphs: Deformable shape and appearance models. *IEEE Trans. PAMI* **30** (2008) 1444–1459
13. Freedman, D., Zhang, T.: Interactive graph cut based segmentation with shape priors. In: *CVPR*. Volume 1. (2005) 755–762 vol. 1
14. Zhu, S.C., Yuille, A.: Region competition: Unifying snakes, region growing, and bayes/mdl for multiband image segmentation. *PAMI* **18** (1996) 884–900
15. Yezzi, A., Soatto, S.: Deformotion: Deforming motion, shape average and the joint registration and approximation of structures in images. *IJCV* **53** (2003) 153–167
16. Schunck, B., Horn, B.: Determining optical flow. In: *Image Understanding Workshop*. (1981) 144–156
17. Nocedal, J., Wright, S.J.: *Numerical Optimization*. Springer (1999)

Automatic feature identification in dental meshes^{*}

Yokesh Kumar¹, Ravi Janardan¹, Brent Larson², and Joe Moon²

¹ Dept. of Computer Science & Engg., Univ. of Minnesota,
{kumaryo,janardan}@cs.umn.edu

² Division of Orthodontics, Univ. of Minnesota,
{larso121,moonx162}@umn.edu

Abstract. The goal of virtual orthodontic treatment planning is to re-position the teeth in a digital dental model so that the desired alignment of the teeth on each dental arch and occlusion (i.e., matching) of the upper and lower arches is achieved. The input to the planning process is a collection of individual tooth objects obtained by segmenting a noisy 3D surface mesh that is generated by laser-scanning a plaster model of the dental arch built from patient-specific dental impressions. A key step in the planning is the identification of features on the surface of the teeth such as cusps, grooves, incisal edges, marginal ridges, and occlusal surface boundary, that are important both for carrying out the alignment and evaluating its quality. This paper presents a collection of techniques to identify such features automatically, with minimal user intervention. Experimental results show the effectiveness of the approach.

Keywords: Alignment, clustering, curvature, watershed method

1 Introduction

Orthodontic treatment planning seeks to re-position the teeth in the dental arches of a patient in order to achieve an outcome that is both aesthetically pleasing and functionally optimal. Technological advances now make it possible for the treatment plan to be simulated virtually (in 3D) and allow the clinician to choose between multiple alternatives in order to achieve the best possible outcome. Virtual planning begins with the acquisition of a 3D surface mesh of each arch. This mesh is generated by building a plaster model of each arch from patient-specific dental impressions and then laser-scanning the model. The mesh is then segmented into individual tooth objects (i.e., submeshes), using, for instance, the algorithms in [3,4]. These tooth objects are the input to the alignment process, which involves proper repositioning of teeth relative to one another in each arch and proper matching of the contact regions between teeth in different arches.

^{*} This research was supported, in part, by the Orthodontics Education and Research Fund at the University of Minnesota.

Dental features and their importance: Proper tooth alignment and functionality (e.g., chewing) depend on a number of intrinsic features on the tooth surface, including cusps, grooves, incisal edges, ridges, occlusal contact region, etc. (These terms are defined in Section 2.) These features are important for many reasons: They provide a set of landmarks that can be used to define alignment requirements quantitatively (e.g., height difference between adjacent marginal ridges) [7]. Also, the features themselves do not change over time (as they are intrinsic to the tooth surface), so they can be used to monitor the progress of the alignment. Moreover, since certain (derived) features need to remain invariant throughout treatment (e.g., the distance between canine tips), the intrinsic features dictate the best possible alignment that can be achieved and provide a means for evaluating a computed alignment. Finally, features such as number of cusps are useful in classifying teeth automatically (e.g., molars, premolars, etc.)

Goals and contributions: The identification of a relevant feature set is crucial for orthodontic treatment planning. However, identifying features manually, by having a clinician “eyeball” each tooth is labor-intensive, time-consuming, and prone to error. Our goal is to identify intrinsic tooth surface features (cusps, ridges, grooves, etc.) automatically, with intervention by the clinician to verify or correct identified features only in difficult or unusual cases. In this paper, we discuss the computational issues associated with automatic feature identification and present a collection of algorithms to do this effectively. Our algorithms are based on curvature analysis, clustering on 2D cross-sections of tooth surfaces, and an adaptation of the watershed method for segmentation. We have incorporated our algorithms in a software tool and we present experimental results that show that our methods are effective at automatic feature identification on noisy and incomplete real-world datasets.

Challenges: The input to the feature identification task is a collection meshes representing individual tooth objects. The feature identification task is complicated by the fact that the meshes are almost always noisy and incomplete. This is due to the limitations (resolution) of the laser scanning process itself and due to the segmentation process subsequently “cutting through” areas of the mesh that are shared between adjacent teeth that very close to each other. A practical requirement of the feature identification algorithms that we seek is that they should be robust to noise and missing information.

Organization of the paper: Section 2 defines relevant anatomical terms and the features of interest. Section 3 describes an algorithm for cusp identification based on the familiar watershed method. Section 4 introduces a general approach based on curvature analysis and clustering on 2D cross-sections of tooth surfaces for identifying other features such as incisal edges, grooves, marginal ridges, and occlusal surface boundary. Section 5 discusses how these features are identified and presents experimental results. Section 6 offers concluding remarks.

2 Dental anatomy and dental features

Dental anatomy: All definitions below are illustrated in Figure 1.

Teeth are classified as *incisors*, *canines*, *premolars* and *molars*. Each *dental arch*, i.e., row of teeth, can be divided into a *left* and a *right* side. Each side has two incisors (*central* and *lateral*), one canine, two premolars (*first* and *second*) and three molars (*first*, *second* and *third*). The incisors and canines are collectively called *anterior*s and are used in cutting action. The premolars and molars are called *posterior*s and are involved in chewing action.

The inner (resp. outer) part of the tooth on the tongue (resp. face) side is called the *lingual* (resp. *facial*) side. The face side on posteriors (resp. anterior) is called the *buccal* (resp. *labial*) side. A tooth's surface towards the front (resp. back) of the arch, i.e., towards (resp. away from) the central incisors, is called the *mesial* (resp. *distal*) side.

A suitably chosen line through the mesial and distal side of each tooth defines the *mesiodistal* line of the tooth. Similarly, the lingual and buccal (or labial) sides define the *buccolingual* line. These lines are important in feature identification and in understanding tooth functionality.

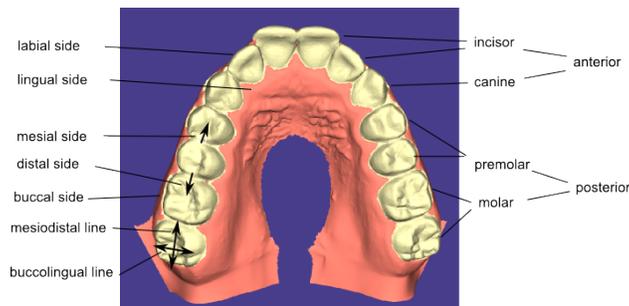


Fig. 1. Dental anatomy. (*All figures in the paper are best viewed in color.*)

Dental features: All definitions below are illustrated in Figure 2. More information can be found in a standard dental text such as [10].

Incisal edges: These are the sharp ridges on the anterior (incisors and canines) running along the mesiodistal line.

Cusps: On premolars and molars these are the mountain peak-like structures on the surface, at the corners of the tooth. Each cusp has *cusp ridges* radiating from its tip, similar to ridges that connect mountain peaks to valleys on a terrain; these can be used to define other features such as the occlusal surface boundary (explained below).

Canines have a single cusp which plays an important role in determining the overall quality of the alignment. The premolars (resp. molars) have 2 or 3 (resp. 4 or 5) cusps depending on the arch (upper or lower) and the individual. These cusps are named according to their position on the tooth surface. For example, a molar would have a mesiolingual cusp situated on the mesial and lingual side

of the tooth. Similarly, for mesiobuccal, distolingual, distobuccal, lingual and buccal cusps.

Occlusal surface and marginal ridges: The *occlusal surface* of a posterior tooth is the area of the tooth surface where chewing occurs. This is also the contact area between corresponding posterior teeth from opposing arches. Thus, the occlusal surface complements the functionality of incisal edges.

The *marginal ridges* are located at the mesial and distal ends of the occlusal surface. These are the regions where the mesial or distal walls of a tooth make contact with the occlusal surface. Thus, each tooth has a mesial and a distal marginal ridge. For incisors, the marginal ridges are on the vertical sides of the teeth since the occlusal surface of incisors is the lingual (resp. labial) surface on the upper (resp. lower) arch.

The occlusal surface of a posterior tooth is bounded on the buccal and lingual sides by the cusp ridges. On the mesial and distal sides, the occlusal surface is bounded by the two marginal ridges. This provides a boundary around the occlusal surface area called *occlusal surface boundary*. Thus, the occlusal surface boundary is a curve on the tooth surface that connects the marginal ridges, cusp ridges, and the cusp peaks.

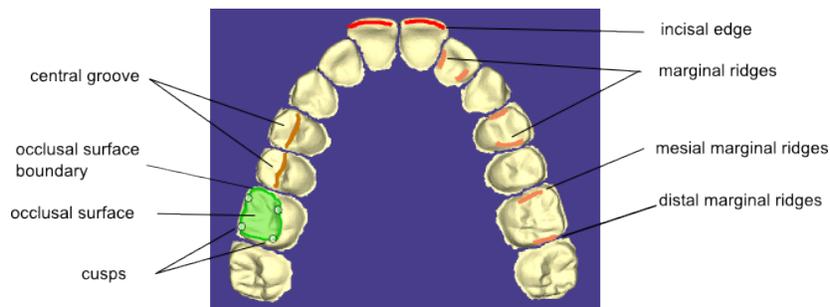


Fig. 2. Cusps, incisal edges, grooves, marginal ridges, and occlusal surface boundary in an upper arch, as found by our feature identification algorithms.

Grooves: These are the depressions and fissures on the occlusal surface of a posterior tooth that resemble riverbeds and valleys on a terrain. There are various types of grooves and corresponding classification and naming conventions. For our study, we are interested in the long grooves running along the mesiodistal line of the tooth, called *central developmental grooves*, or just *central grooves*.

3 Cusp identification

Our approach is based on the *watershed* approach to mesh segmentation [6]. Although this method is not generally suitable for segmenting very noisy meshes (such as dental meshes), it is well-suited for identifying cusps on the tooth surface. Cusps are defined as mesh vertices that are local minima with respect to

an appropriately defined height function. (For teeth on the upper arch, cusps are local minima while for those on the lower arch they are local maxima; we treat both cases uniformly, as local minima, by negating vertex coordinates in the latter case.)

The key to watershed-based cusp extraction is designing a good height function \mathcal{H} . Most generic mesh segmentation schemes define \mathcal{H} solely on the basis of surface curvature, so as to be able to segment meshes from a variety of sources. Here we take advantage of the fact that we are dealing with tooth surfaces and we define \mathcal{H} on the basis of both surface curvature and the elevation of the vertices (i.e., their z -coordinates, since the scanned meshes are provided with a base that is parallel to the xy -plane). More precisely, for a vertex v on the surface of a tooth in the upper arch, we define its height as $\mathcal{H}(v) = (1 - \alpha) \cdot (-K_v) + \alpha \cdot v_z$, where v_z is v 's z -coordinate, K_v is the surface curvature at v , and α is a parameter that controls the relative influence of K_v and v_z on $\mathcal{H}(v)$. In our implementation, we found that choosing α from $[0.4, 0.6]$ worked well for all the models tested.

Figure 3 shows the result of applying the cusp identification algorithm with the above height function on three different models. Figure 3(a) shows a lower arch in which teeth are not well-aligned. Figures 3(b) and 3(c) show two upper arches where the molars are rotated away from the vertical z -axis by different amounts. Note that the number of cusps for teeth on the lower arch can be four or five, and we may sometimes have to deal with partially erupted molars.

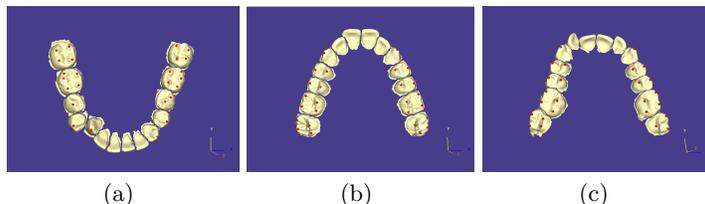


Fig. 3. Cusps (shown red) identified using watershed method.

After the cusps have been identified automatically, our software tool allows the user to fine-tune the results manually (for very difficult cases) by clicking on the tooth surface to add/delete (hence move) cusps. For instance, a click to add a cusp triggers a search for the local minimum in a small neighborhood of the surface around the location of the click and places a cusp at this minimum. Thus the user does not have to “eyeball” the mesh to find the exact location of the desired cusp.

4 A general approach to feature identification

In this section we discuss how to identify the remaining features of interest (beyond cusps), i.e., incisal edges, central grooves, marginal ridges, and occlusal surface boundary. It turns out that these features can all be identified using a general approach that is based on computing certain 2D cross-sections of the

tooth surface, analyzing the curvatures of the (piecewise-linear) curves that define these cross-sections to determine points of sufficiently high (positive or negative) curvature, and identifying clusters of such points that form the features of interest.

Our approach starts by creating a sequence of parallel planar cross-sections of the tooth surface, one set oriented along the mesiodistal line and one along the buccolingual line. For each cross-section in each set, we perform 2D curvature analysis at the vertices of the piecewise-linear curve that defines the cross-section and identify points of high curvature and cluster these to identify regions of interest on each curve. Then appropriate regions are selected from each curve and stitched together with those from curves of neighboring parallel cross-sections to obtain a connected and coherent feature (e.g., a ridge or groove). Thus, this approach provides a means to extract the desired 3D features in a “guided” fashion from a set of 2D regions of interest on individual cross-sections. This has several advantages over an approach that attempts to extract features directly, via curvature analysis in 3D, as we will discuss in Section 4.4.

We will now explain the main steps of the approach in detail.

4.1 Computing a medial curve

Consider the left side of the arch in Figure 1. We can define a curve that passes suitably through the mesial and distal sides of each tooth on the left side starting from the last molar and ending at the central incisor. A similar curve can be defined for the right side of the arch. These two curves meet at the mesial sides of the two central incisors. We define the *medial curve* of the arch as the curve obtained by joining the curves from the left and right sides of the arch. The medial curve will guide the feature identification process outlined above. The medial curve represents an approximate layout of all the teeth on the arch and also provides a mesiodistal line for each tooth that will be needed to compute the mesiodistal cross-sections, as we will see in Section 4.2.

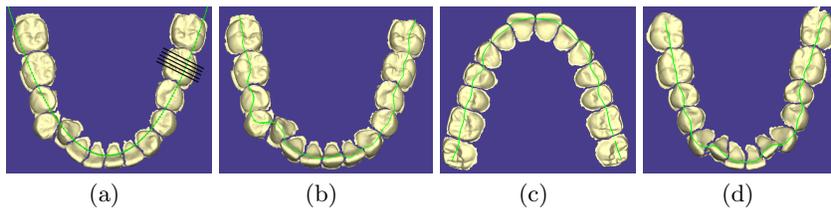


Fig. 4. Computation of a medial curve (shown green): (a) Initial fit of the arch with a cubic curve. (b) Computed medial curve for arch in (a). Additional examples of medial curves on an upper arch (c) and a lower arch (d).

The medial curve is not unique since many different curves can satisfy the given definition. For our purposes, it suffices to have an approximate medial curve that reliably identifies the mesial, distal, buccal (labial) and lingual sides of a tooth. Accordingly, we compute the medial curve as follows.

1. Project the triangular faces on each tooth surface to a set of triangles, F_{xy} , on the xy -plane.
2. Find a cubic curve C_{fit} that best fits the 2D projections of the vertices in F_{xy} (Figure 4(a)).
3. Sample the length of C_{fit} at suitably small intervals to get an ordered sequence, $P = p_1, p_2, \dots, p_n$, of points on C_{fit} .
4. At each point $p_i \in P$, define a line l_i orthogonal to C_{fit} (Figure 4(a)). Find the median point, m_i , of the intersection points of l_i with F_{xy} . Output the medial curve as the piecewise-linear curve that connects the ordered sequence, $C_{\text{med}} = m_1, m_2, \dots, m_n$, of medians (Figure 4(b)).

Figure 4 shows the results of this algorithm.

4.2 Computing planar cross-sections

Using the medial curve described above, we compute an approximation, L_{md} , of the mesiodistal line, of a given tooth by finding the best-fitting line segment for the section of the medial curve segment that goes through the tooth. We can also compute an approximation, L_{bl} , of the buccolingual line of the tooth by taking the line orthogonal to the mesiodistal line and passing through the centroid of the xy -projection of the tooth. We will refer to the approximation of the mesiodistal (resp. buccolingual) line as simply the *mesiodistal* (resp. *buccolingual*) *line*.

We define the *cross-section* of (the mesh of) a tooth object with respect to a plane as the piecewise-linear curve obtained by intersecting the plane with the triangular faces of the mesh. A plane that is parallel to both L_{md} and the z -axis and intersects the tooth object generates a *mesiodistal cross-section*. Let P_{md} be a set of such uniformly-spaced planes and let the set of corresponding cross-sections generated be \mathcal{M} . Similarly, a plane that is parallel to both L_{bl} and the z -axis and intersects the tooth object generates a *buccolingual cross-section*. Let P_{bl} be a set of such uniformly-spaced planes and let the set of corresponding cross-sections generated be \mathcal{B} . Figure 5(c) is an example of a cross-section.

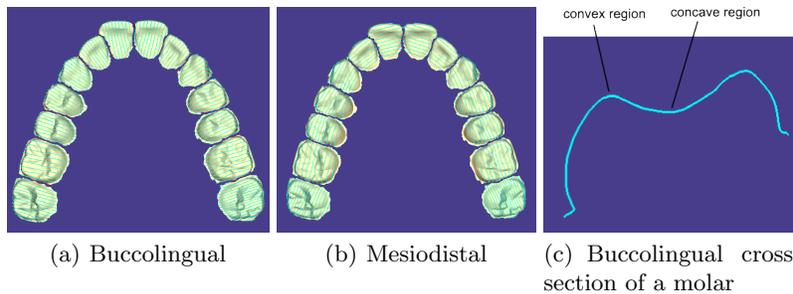


Fig. 5. Examples of planar cross-sections on the tooth surface (shown in cyan).

4.3 Clustering of high-curvature regions

The planar cross-sections generated are piecewise-linear and composed of a sequence of points (vertices) in 3D. The curvature of a point p_i on such a curve C (relative to the underlying smooth curve) can be approximated as $\kappa(p_i) = \frac{\Delta\theta}{|p_{i-1}p_i| + |p_i p_{i+1}|}$, where p_{i-1} (resp. p_{i+1}) is the previous (resp. next) point with respect to p_i on the curve C , $|p_i p_j|$ is the length of the edge from p_i to p_j , and $\Delta\theta$ is the (signed) angle between the lines passing through line segments $p_{i-1}p_i$ and $p_i p_{i+1}$ [2]. The signed angle $\Delta\theta$ allows us to define positive and negative curvatures at points. The points corresponding to the positive (resp. negative) curvatures form the convex (resp. concave) regions on the curve with respect to the plane containing the curve (Figure 5(c)).

Given the buccolingual and mesiodistal cross-sections, we find the vertices on the curves for which the magnitude of curvature is greater than a specified threshold, T_{curv} . The sign of the curvature threshold may be positive or negative depending on the targeted feature (e.g., ridges (resp. grooves) have high positive (resp. negative) curvatures). The set, S , of high-curvature vertices resulting from the thresholding serves as the input to the clustering algorithm.

Our goal behind clustering is to aggregate the vertices in S into connected components of high-curvature vertices that correspond to features of interest. The extracted clusters reveal these features and also cause noise regions to be grouped into small isolated components. The latter can be identified and deleted easily using a threshold, T_{noise} , on the number of noise vertices.

We compute the connected components as follows. Initially all vertices in S are taken to be singleton components. Starting with this, we repeatedly merge a pair of components if the closest distance between them is less than a specified threshold, T_{ccd} . The closest distance between two components is defined as the minimum of the pairwise distances between their corresponding vertices.

The clustering method requires three user-specified threshold parameters: T_{curv} , T_{noise} , and T_{ccd} . We have observed experimentally that, for each feature type of interest, values of these parameters can be found that work across a broad range of patient datasets. Thus, the parameter values can be predetermined and our algorithm does not, in general, require user input during routine operation. Section 5 below, which describes how the methods of this section can be used to identify various features, provides the specific parameter values that our implementation uses.

4.4 Discussion

The features of interest to us could be defined naturally in terms of the curvatures at the mesh vertices (relative to the underlying 3D surface) [1,5,9]. For example, cusps, incisal edges, and ridges consist of clusters of vertices of high positive curvature (convex regions), whereas, grooves consist of clusters of vertices of high negative curvature (concave). Hence it is natural to wonder if 3D curvature analysis can be used for feature recognition.

We implemented this approach and found that, unfortunately, it did not perform well on most datasets. There are several reasons for this: (a) Tooth surfaces tend to exhibit sharp changes in curvature over small neighborhoods. This makes it difficult to compute 3D curvatures reliably and use them to identify partitioning ridges and valleys. (b) For each feature type, curvature threshold settings vary quite considerably from one model to the next (in contrast to the situation in the 2D approach); this requires users to set the thresholds on a case-by-case basis and imposes an undue burden. (c) The presence of spurious ridge- and groove-like structures on the occlusal surface and very close to the target features makes the identification and cleanup of actual features difficult. We tried using skeletonization methods for 3D meshes that are based on morphological operators [8], but these did not perform well. Moreover, as Figure 6 shows, the sizes of these noise regions are comparable to those of actual features, which makes identification of true features extremely difficult.

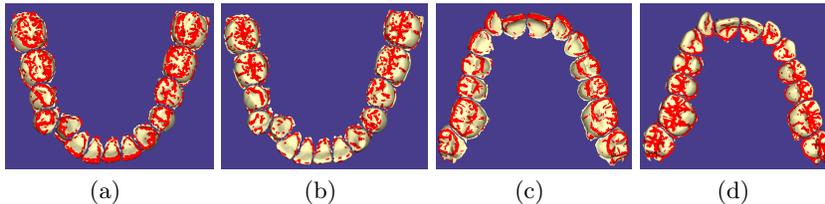


Fig. 6. Examples of poor feature recognition using 3D surface curvatures on a lower and an upper arch.

It is for these reasons that we chose to develop the 2D method, based on cross-sections, outlined earlier. It is easy to see that using a large number of cross-sections effectively reproduces the entire tooth surface to high accuracy. Since the cross-sections are generated by a set of parallel planes, this process can be thought of as a “guided” reconstruction of the 3D surface. An advantage of this is that we now have the flexibility to select only those parts of the cross-sections that yield interesting information. For example, on a buccolingual cross-section of a molar, the outermost convex regions corresponds to the mesial and distal cusp ridges and the innermost concave regions correspond to the grooves (Figure 5(c)). The notions of “outermost” and “innermost” are easy to define relative to the natural directionality provided by a 2D curve, but it is not clear how to do this on a surface. Furthermore, the parallel orientation of the cross-sections makes it easy to stitch together the portions of features found on each cross-section into a 3D feature.

5 Feature identification

This section describes how the techniques of Section 4 can be used to automatically identify features other than cusps, such as incisal edges, grooves, marginal

ridges, and occlusal surface boundary. (Cusp identification, via a different approach, was discussed in Section 3.)

5.1 Incisal edges

These can be extracted as a cluster of convex regions of high-curvature vertices on the buccolingual cross-sections of the anterior teeth. The results of our incisal edge detection algorithm are shown in Figure 7 for two pairs of upper and lower arches. Note that the algorithm performs correctly even when the tooth is in a rotated position from the global archform (Figures 7(b) and 7(d)).

As mentioned in Section 4.3, there are three parameters that control the clustering process. Their specific values for incisal edge detection are as follows: $T_{curv} = 0.6$ (the curvatures are normalized to the range $\pm[0, 1]$ here and later), $T_{noise} = 1/10$ th of the number of buccolingual cross-sections and $T_{ccd} = 0.5$ mm. (We reiterate that these parameters work well on all models tested here and later, and do not require adjustment by the user.)

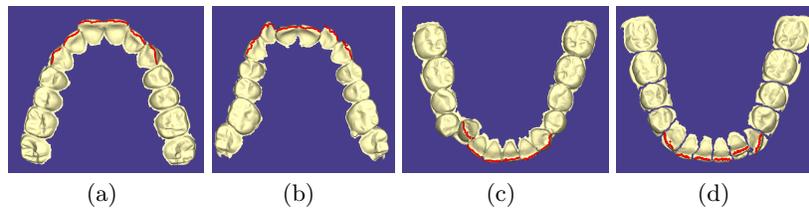


Fig. 7. Incisal edges (shown red) on anteriors identified using buccolingual cross-sections. Upper arches shown in (a) and (b); lower arches in (c) and (d).

5.2 Grooves

Grooves are found on the posterior teeth and are extracted as the concave corners in the middle-section of the buccolingual cross-sections. (Figure 5(c).) As mentioned in Section 2, there are many types of grooves on the posteriors, especially the molars. We are interested in the central grooves that run along the mesiodistal line of the tooth. Figure 8 shows the results of our groove identification algorithm on two pairs of upper and lower arches. As expected, the grooves on the upper arches have more variation than the ones on the lower arches and consist of multiple “branches”. For the purpose of treatment planning, we may choose to approximate these grooves by a straight line segment or a curve.

The specific parameter values used for groove extraction are as follows: $T_{curv} = -0.4$, $T_{noise} = 1/10$ th of the number of buccolingual cross-sections and $T_{ccd} = 0.5$ mm.

5.3 Marginal ridges

These are extracted using the mesiodistal cross-sections of the tooth surface. Figure 9 shows the results of our algorithm on two pairs of upper and lower

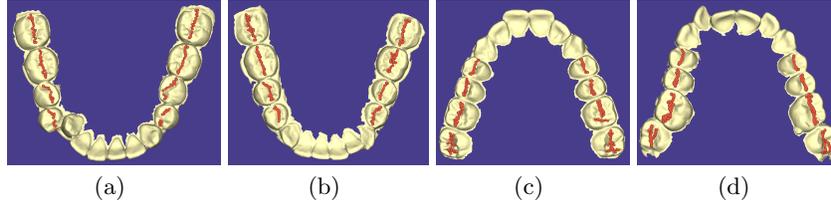


Fig. 8. Central grooves (shown brown) on posteriors identified using buccolingual cross-sections. Note the difference in the groove patterns between the lower arches ((a), (b)) and upper arches ((c), (d)).

arches. The marginal ridges on the lower arch premolars are difficult to identify because they are not easily separable from the cusp ridges and also may not be properly aligned with the mesiodistal line (see Figure 9(a) and (b)). Also, the distal marginal ridge on the last available molars often cannot be accurately identified because of the inherent ambiguities in the structure of these teeth.

The specific parameter values used for marginal ridge detection on an upper arch is $T_{curv} = 0.3$, $T_{noise} = 1/10$ th of the number of mesiodistal cross-sections and $T_{ccd} = 0.5$ mm. The mesial and distal walls on the lower arch posteriors are higher than those on the upper arch and, thus, we set $T_{curv} = 0.6$. The threshold for T_{noise} and T_{ccd} are the same as for the upper arch.

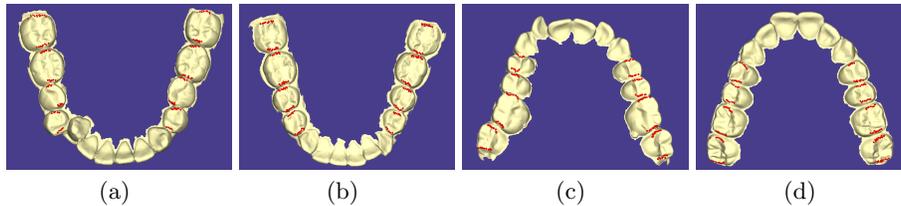


Fig. 9. Marginal ridges (shown red) on posterior teeth from mesiodistal cross-sections.

5.4 Occlusal surface boundary

The occlusal surface boundary of a posterior tooth is defined by the marginal ridges on the mesial and distal sides, and the cusp ridges on the buccal and lingual sides. (Section 2.) Section 5.3 describes how to identify the marginal ridges. Cusp ridges on the buccal and lingual sides are also identified via buccolingual cross-sections, via an approach similar to that for the incisal edges, but using both the buccal and the lingual sides of the cross-sections. Figure 10 shows the result of our cusp ridge identification algorithm.

The specific parameter values used for cusp ridge extraction are $T_{curv} = 0.5$, $T_{noise} = 1/10$ th of the number of buccolingual cross-sections, and $T_{ccd} = 0.5$ mm.

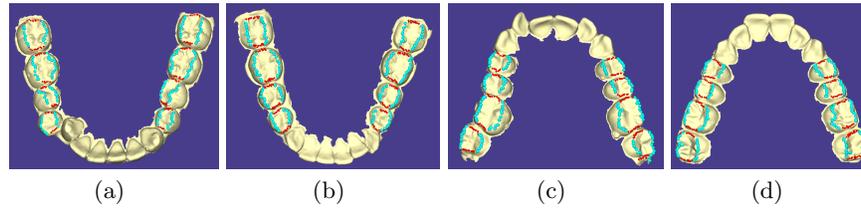


Fig. 10. Occlusal surface boundary identification on posterior teeth using buccolingual cross-sections. The occlusal surface is bounded by the marginal ridges found earlier (shown red) and the cusp ridges on the buccal and lingual sides (shown cyan).

6 Conclusion

We have presented techniques to automatically identify tooth surface features in noisy and incomplete dental mesh models. These features are important in virtual orthodontic treatment planning. Our methods use buccolingual and mesiodistal cross-sections of the tooth surface to facilitate a “guided” extraction of incisal edges, central grooves, marginal ridges, cusp ridges, and occlusal surface boundary. We also described a watershed-based cusp identification algorithm. The algorithms have been implemented and experiments on real-world datasets show that the methods are effective in automatic feature identification.

References

1. T. Gatzke and C. M. Grimm. Estimating curvature on triangular meshes. *Intl. J. of Shape Modeling*, 12(1):1–28, 2006.
2. R. Klette and A. Rosenfeld. *Digital geometry*. Morgan Kaufman, 2004.
3. T. Kronfeld, D. Brunner, and G. Brunnett. Snake-based segmentation of teeth from virtual dental casts. *Comp.-Aided Design. and Appl.*, 7(2):221–233, 2010.
4. Y. Kumar, R. Janardan, B. Larson, and J. Moon. Improved segmentation of teeth in dental models. *Comp.-Aided Design. and Appl.*, 8(2):211–224, 2011.
5. Y. Lee, S. Lee, A. Shamir, D. Cohen-Or, and H. Seidel. Intelligent mesh scissoring using 3D snakes. In *Proc. Comp. Graphics and Appl.*, pages 279–287, 2004.
6. A. P. Mangan and R. T. Whitaker. Partitioning 3D surface meshes using watershed segmentation. *IEEE Trans. on Vis. and Comp. Graphics*, 5(4):308–321, 1999.
7. American Board of Orthodontics. Grading system for dental casts and panoramic radiographs, 2010. <http://tinyurl.com/ygstyfu>.
8. C. Rössl, L. Kobbelt, H. Seidel, and I. Stadtwald. Extraction of feature lines on triangulated surfaces using morphological operators. *AAAI Symp. on Smart Graphics*, pages 71–75, 2000.
9. T. Weinkauff and D. Günther. Separatrix persistence: Extraction of salient edges on surfaces using topological methods. *Computer Graphics Forum (Proc. SGP '09)*, 28(5):1519–1528, July 2009.
10. J. B. Woelfel and R. C. Scheid. *Dental Anatomy: Its Relevance to Dentistry*. Lippincott Williams & Wilkins, 2011.

3D Shape Analysis of the Knee Extensor and Flexor Muscles in Patients with COPD using Mesh Projection-based Features

Hengameh Mirzaalian¹, Ghassan Hamarneh¹,
Bahareh HajGhanbari², W. Darlene Reid²

¹Medical Image Analysis Lab, Simon Fraser University, BC, Canada

²Muscle Biophysics Laboratory, Vancouver Coastal Health Research Institute, Department of Physical Therapy, University of British Columbia, BC, Canada

Abstract. Loss of muscle mass in severe chronic obstructive pulmonary disease (COPD) increases the risk of mortality more than ten-fold compared to those with mild disease. Magnetic resonance imaging (MRI) has been used as a valid, reliable and non-invasive tool to investigate the changes in muscle mass in people with COPD. Using MR images to perform 3D shape analysis of thigh muscles, we extended and applied several state-of-the-art 3D shape descriptors to examine the classification accuracy of a trained support vector machine classifier to distinguish 4 thigh muscles in 20 people with COPD versus 20 healthy controls. The results of the study show high classification accuracy (with an average of 94%). Such a classification system may identify patients at risk of COPD more readily so that early interventions to reverse muscle atrophy can be provided.

1 Introduction

Chronic obstructive pulmonary disease (COPD) is defined by progressive air-flow limitation that is not fully reversible, which causes loss of thigh muscle mass (muscle atrophy) in people with COPD. The statistics reported by the World Health Organization show that COPD shares 4th and 5th places with HIV/AIDS as a single cause of death (after cardiovascular disease, cerebrovascular disease and acute respiratory infection) [4]. Magnetic resonance imaging (MRI) has been widely used as an appropriate tool for monitoring and evaluating the differences in muscle disease distribution and severity [12]. Complimentary techniques to MRI that can better detect atrophy-related changes in individual thigh muscles might facilitate targeting of interventions such as strength training and gene therapy in people with COPD [12]. Thus, more effective therapeutic approaches and preventive strategies can expedite the improvement of muscle function, exercise tolerance, and physical activity in people with COPD [12]. Such a diagnostic technique could also have more widespread application to other chronic diseases that demonstrate muscle atrophy. Recently, in [10], we investigated 3D shape and size measurements to examine the classification

accuracy of a trained support vector machine (SVM) classifier in distinguishing individual thigh muscles in a group of COPD patients compared to healthy people, whom were range-matched for age, gender, and body mass index. The aim of this study is to further explore, extend and employ state-of-the-art 3D shape descriptors (e.g. wavelet transform-based methods [14]) to achieve higher classification accuracy.

Depending on the spatial extent of the region of interest of anatomical structures, shape descriptors can be classified into two main categories: *local* or *global* features. Shape classification using local features requires establishing point correspondence between the given shapes. For neurological (non-musculoskeletal) applications, several approaches have been proposed to compare the extracted corresponding local shape descriptors for classification [21]. However, in the absence of identifiable anatomical landmarks to establish correspondence between the shapes, which is the case for thigh muscles, global shape descriptors could be more useful. Several studies have been performed to investigate the relationship between the shapes of anatomical structures and their pathological groups using global shape descriptors. For example, Durrleman et al. [6] characterized 3D shapes of brain structures via vector fields and constructed global shape descriptors by taking the integral of the vector fields. Gutman et al. [9] proposed to extract a rotation invariant shape descriptor of an inverse conformal map for the hippocampus surface using a spherical harmonics representation. Wang et al. [19] and Chen et al. [3] applied Fourier descriptors to the MR images to diagnose autism from MR images and to the description of shape changes in the human mandible, respectively.

To the best of our knowledge, there is a conspicuous scarcity of studies that have examined the relationship between (non-cardiac) muscle pathology and morphology outside our own group’s work. Kaick et al. [17] used the statistics of the Fourier coefficients extracted from 2D contours of the muscle to classify supraspinatus muscle and differentiate a normal supraspinatus muscle from several pathologies. However, these features were not used for thigh muscle classification. Ward et al. [20] computed a number of 3D shape descriptors (e.g. volume, surface area, 3D moments described in Table 1) for supraspinatus muscle classification. HajGhanbari et al. [10] used the same shape descriptors applied by Ward et al. [20] for thigh muscle classification and made them more localized by dividing each muscle into four equal quarters (regions) along its longest axis and calculated the measures for each region. In our aforementioned works, we have not yet utilized state-of-the-art 3D shape descriptors for thigh muscle classification. In Section 5, we compare the classification accuracy of our proposed method to these earlier techniques [10,20].

In this work, to study thigh muscle anatomy, we utilize a method that encodes localized morphological properties without the prerequisite of explicitly calculating a point-to-point correspondence between shapes. In particular, we first apply the wavelet transform (WT)-based shape descriptors proposed by Papadakis et al. [14] (Section 3); second, we propose a methodological extension of Papadakis et al.’s work [14] in a way that the descriptor follows the natural geometry of the

	Measurement	Description
1	3D moment J_1	Capture characteristics of the spatial distribution of the voxels that make up the muscle shape.
2	3D moment J_2	
3	3D moment J_3	
4	Mean of distances to centroid	Measure surface non-sphericity.
5	Std. dev. of distances to centroid	
6	Eigenvalue ratio λ_1/λ_2	Characterize the tubular, planar, and spherical geometry of the shape.
7	Eigenvalue ratio λ_1/λ_3	
8	Eigenvalue ratio λ_2/λ_3	
9	Surface area	Calculated as the sum of the areas of the triangular mesh of the muscle.
10	Volume	Calculated as the number of voxels inside the mesh of the muscle multiplied by the size of each voxel in mm^3 .
11	Surface area/volume	The ratio between the surface area and volume.

Table 1: Measurements taken by Ward et al. [20] for supraspinatus muscle classification.

muscle more accurately (Section 3.2). We validate the classification accuracy of the extracted features on MR images taken from the individual thigh muscles of 40 subjects (20 healthy versus 20 people with COPD) (Section 2). The results of our study show that SVM-classifier can differentiate individual thigh muscles in COPD group from those of the healthy group with an average accuracy of 94%. A comparison between the extracted feature vectors reveals that the WT-based shape descriptors outperform Ward et al.’s [20] and HajGhanbari et al.’s [10] shape descriptors for all of the four knee flexor and hip adductor muscles. In addition, the classification accuracy of our proposed extension of the WT-based descriptors is greater than that of the original method in [14] (Section 5).

2 Material

Forty subjects, twenty people with COPD and twenty healthy adults, matched for age, gender, and body mass index participated in this study. T1-weighted MR images (field of view 40 cm^2 , matrix 256×256 , 5 mm thick) were taken from each subject’s thigh from the anterior superior iliac spine to the tibial plateau, yielding 80 to 100 axial slices for each subject. Once the images were collected using the DICOM image file format, they were loaded into ITK-SNAP software [10]. Slice-by-slice segmentation of 4 individual thigh muscles was performed by an expert clinician (Figure 1). Finally, triangulated meshes were extracted from the ITK-SNAP 3D segmentation and were used for the shape representation and classification steps. Note that, in this study, in order to prepare scale and translation invariant shape descriptors, the muscles are normalized to femur length for each subject and are centered around the origin. Also, the imaging protocol ensured a consistent orientation of the knee extensor and flexor muscles



Fig. 1: Example segmentation of the knee extensor and flexor muscles in ITK-SNAP: (a) Horizontal, (b) sagittal, (c) coronal, and (d) 3D mesh view. The different knee extensors and flexors are represented by different colors.

across all subjects. More details about the data acquisition and surface mesh preparation can be found in [10].

3 2D Wavelet Transform Shape Descriptors

Papadakis et al. [14] computed 3D shape descriptors using the mean and variance of the wavelet transform (WT) coefficients of panoramic views of a 3D object. In the following sections, we illustrate how to extract the panoramic views of the 3D mesh of the muscle (Section 3.1) and then describe our proposed extensions to the cylindrical projection (Section 3.2). Finally, we discuss the extracted WT-based shape descriptors (Section 3.3).

3.1 Extraction of Panoramic Views by Cylindrical Projection

A panoramic view of a 3D object is obtained by projecting the 3D object onto the lateral surface of a *projection cylinder*; a cylinder that is parallel to one of the three principal axes (X, Y, or Z). The panoramic views are used to capture the position and orientation of the object’s surface in 3D space. Figure 2(b) depicts the discretized lateral surface of the projection cylinder parallel to the Z-axis, whereas Figure 2(e) shows a cross section of the 3D object (blue curve) and the cylinder (green circle). In the cylindrical projection step of Papadakis et al.’s method [14], each point of the cylinder (e.g. point p in Figure 2(e)) would be assigned the radius r of the cylindrical coordinate of the furthest point inside the sector containing that point (e.g. point q in the gray sector in Figure 2(e)). Then, a 2D gray-scale image is created by unfolding the cylinder, such that the image pixels and intensities correspond to the cylinder’s vertices and the radii assigned to them, respectively (Figure 2(f)). Performing similar projections onto the other cylinders parallel to the X and Y axes, results in two additional gray-scale images.

Construction of the WT-based shape descriptors from the extracted 2D images resulting from X, Y and Z cylindrical projections is described in Section 3.3. Let WT-X, WT-Y and WT-Z denote the extracted WT-based shape descriptors of the projection cylinders oriented along the X, Y, and Z-axis, respectively. Papadakis et al. [14] constructed the shape descriptors by concatenating WT-X, WT-Y and WT-Z into WT-XYZ, i.e. $WT-XYZ = [WT-X, WT-Y, WT-Z]$. Since

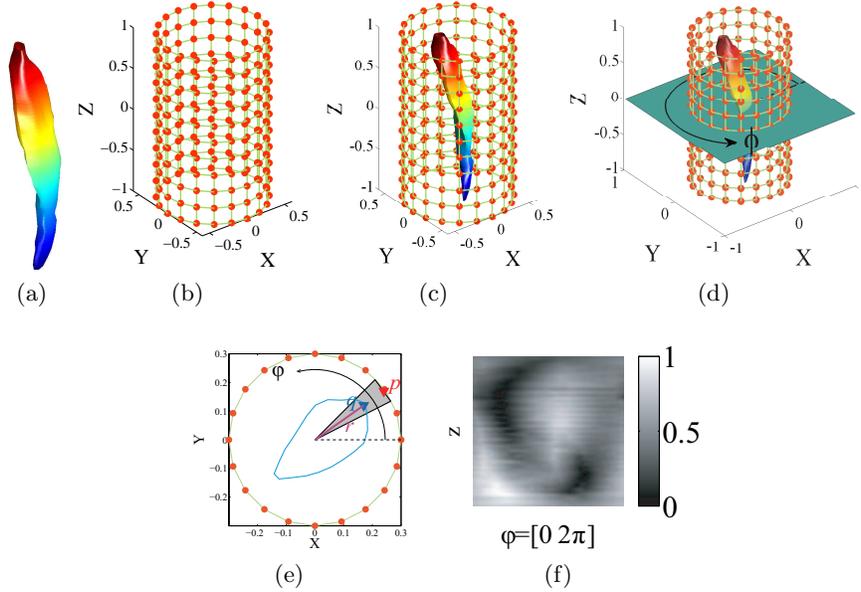


Fig. 2: Panoramic view of a 3D shape. (a) 3D mesh of a thigh muscle. The change in color from blue to red reflect the increasing Z coordinates of the mesh vertices. (b) Discretized (or sampled) cylinder, on which the muscle in (a) is projected. (c) 3D representation of both the muscle and the cylinder prior to projection. (d) 3D representation of a cross section (green plane) of the mesh and the projection cylinder. Note how the angle ϕ is defined on the plane. (e) The 2D cross section of (d). Each point on the cylinder surface (e.g. point p in (e)) is assigned the radius (r) of the cylindrical coordinate of the furthest point of the muscle (blue curve) inside the sector containing that point (e.g. point q in the gray sector). (f) 2D image obtained from unfolding the cylinder in (b). The intensity values at each pixel of the 2D image in (f) correspond to the radii r assigned to the cylinder’s vertices.

each thigh muscle in our study has a tube-like prolate shape that is elongated and aligned along the Z -axis (i.e. the Z -axis is the principle axis of each muscle as shown in Figure 2(c)), the projections along the cylinder parallel to the Z -axis are the most informative. Our results in Section 5 confirm that using only WT- Z as the shape descriptors achieves a classification accuracy similar to that of WT-XYZ but requires only about one third of the computation.

3.2 Extraction of Panoramic Views by Mesh Projections

Generally, the linear axes of the muscles are aligned with the Z -axis. However, to be more precise, each muscle has a *non-linear* (curved) axis that extends along the muscle’s central axis (or medial axes). For this reason, cylindrical projections that follow the natural curved axis of the muscle geometry are more accurate and are expected to be even more descriptive than those resulting from a linear axis. In this subsection, we describe the details of developing this curved axis extension.

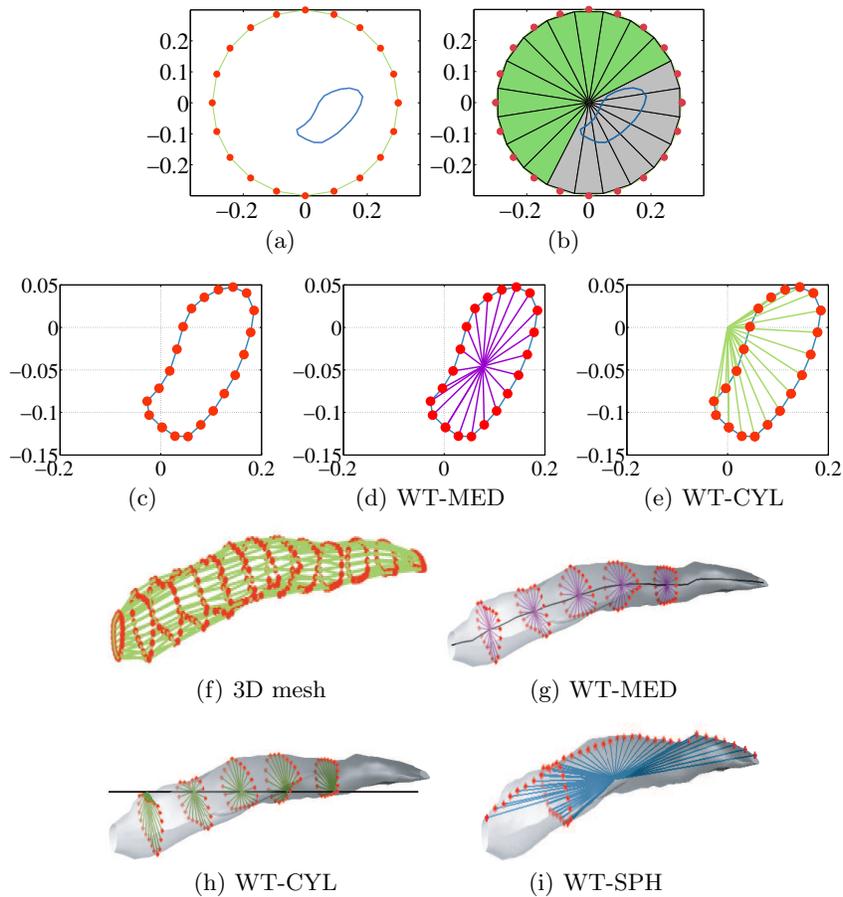


Fig. 3: Extraction of panoramic views by mesh projection. (a) Cross section of the muscle (blue curve) and the cylinder (green circle). (b) There is no intersection of the muscle (blue curve) with any of the green sectors. (c) Vertices of the contour of the mesh in the cross section. (d) Distances (purple lines) from the centroid of the curve to the mesh vertices in the cross section. (e) Radius of the cylindrical coordinate of the points (green lines), i.e. the distance from the origin $(x, y) = (0, 0)$ to the mesh vertices in the cross section. (f) 3D mesh of the muscle. (g-i) 3D representation of the distances assigned to the vertices based on: (g) the medial curve of the muscle; (h) the radius of the cylindrical; and (i) the radius of the spherical coordinates.

Our methodological extension is illustrated with the example in Figure 3. Given a cross section of the muscle (blue curve) and the Z-axis-aligned cylinder (green circle) shown in Figure 3(a), in the cylindrical projection step of Papadakis et al.’s method [14], the values assigned to the vertices belonging to the green sectors in Figure 3(b), would be zero. The reason is that none of the points of the muscle’s cross-section contour (blue curve) lie inside any of those sectors.

As the primary goal is to characterize local geometrical properties, assigning zero values to the sectors that do not intersect with the muscle would cause

a problem because this results in having sectors devoid of mesh points in the projection images; and due to that, the unfolded image would contain numerous missing pixel values. Clearly, this approach would lose information that captures the muscle mesh geometry. To address this problem, we perform the projections on a curved cylinder whose axis follows the geometry (the central or medial axis) of the muscle, rather than performing the projections on the Z-axis-aligned cylinder. In other words, we effectively perform the projections onto the muscle mesh itself.

Given the 3D mesh of the muscle (Figure 3(f)), we compute the 2D projection image by first assigning values to the mesh vertices and then unfolding the mesh into the 2D image with scalar-valued pixels. The values assigned to the vertices are chosen to be the radii of the polar coordinates of the muscle cross section, with the caveat that the center of the coordinate system has to be translated from its position along the Z-axis (Figure 3(d)) into its new position at the center of the muscle cross section curve (Figure 3(e)). Collecting these new cross section centers from cross sections at different Z values form the curved axis of the muscle (black curve in Figure 3(g)); an approximation of the muscles medial axis [15,16]. Finally, we collect the WT-based shape descriptors of the resulting image into the vector WT-MED (short for wavelet-medial).

We also evaluate two additional variant approaches:

- I . In contrast to WT-MED, the first variant does not translate the center of the polar coordinate system but rather maintains its position along the Z-axis (similar to the original WT-Z). However, different from WT-Z, which assigns a value to each vertex of the projection cylinder parallel to the Z-axis, this method assigns a distance to each vertex of the muscle’s mesh (Figure 3(h)). The assigned values are effectively the radii of a cylindrical coordinate system representation. Therefore, we refer to the extracted WT-based shape descriptors from the resulting image by WT-CYL.
- II . The second variant assigns the distance from the center of the muscle in 3D (not the center of the 2D cross section of the muscle) to the mesh (Figure 3(i)). This essentially encodes the radius of the spherical coordinates of the mesh vertices and, hence, we refer to this approach as WT-SPH.

3.3 Statistics of the WT Coefficients

All of the projection approaches presented earlier result in 2D scalar images, which we denote by $f(u, v)$. The wavelet transform is then performed on these 2D images and wavelet coefficients are collected to construct the shape descriptors. Similar to the Fourier transform (or Fourier analysis or decomposition), which captures the magnitudes of different sinusoidal harmonics or frequencies existing in the whole 1D function (or signal) or 2D image, the wavelet transform achieves the same goal with the following extensions [2]. Instead of a global view of the sinusoidal frequencies that make up the whole signal, the wavelet transform focuses on analyzing the similarities between different *localized* regions of the signal and scaled and translated versions of certain function, known as the mother wavelet (instead of the sinusoidal functions in the Fourier analysis). In

particular, the discrete wavelet transform dwt coefficients of the 2D image $f(u, v)$ are extracted by:

$$dwt(s, a, b) = \frac{1}{\sqrt{s}} \sum_{u=1}^M \sum_{v=1}^N f(u, v) \psi\left(\frac{u-a}{s}, \frac{v-b}{s}\right) \quad (1)$$

where ψ is the mother wavelet, and s and (a, b) are scale and translation parameters, respectively. In the WT framework, the WT coefficients are extracted from the different subband images of f resulting from filtering and sub-sampling of f at different scales. At each scale s , there are three detail images, denoted by f_s^{LH} , f_s^{HL} and f_s^{HH} , each of which respectively contains the horizontal, vertical and diagonal high frequency information of the image. There is also an approximation image, denoted by f_s^{LL} , which contains the low frequency information (f_s^{LL} is recursively decomposed as shown in Figure 4). For N different decomposition levels, the total number of subband images f_i is $3 \times N + 1$: $f_i \in \{f_s^{LH}, f_s^{HL}, f_s^{HH}, f_s^{LL}\}$ for $s \in \{1, 2, \dots, N\}$.

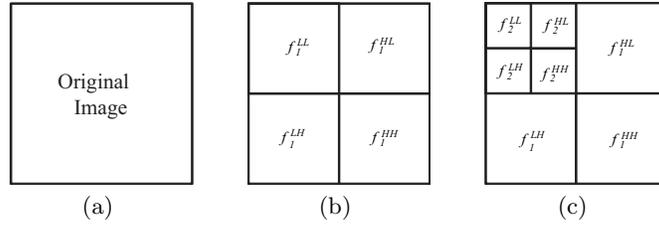


Fig. 4: Wavelet decomposition of a 2D image. A schematic diagram showing how an original image (a) is decomposed into 4 subbands (b) (i.e. one-level, $N = 1$), and then how f_1^{LL} is further decomposed (c) resulting in 7 subbands (i.e. two-level, $N = 2$).

The WT-based feature vector FV of f is computed as the concatenation of the means and variances, μ_i and σ_i , of the WT coefficients of every subband image f_i , and is given by:

$$FV = \{FV_1, FV_2, \dots, FV_{3 \times N + 1}\}, \quad FV_i = \{\mu_i, \sigma_i\}. \quad (2)$$

where μ_i and σ_i of the computed coefficients dwt_i (1) of the i^{th} subband image f_i of size $m_i \times n_i$ are given by:

$$\mu_i = \frac{\sum_{u=1}^{m_i} \sum_{v=1}^{n_i} |dwt_i(s, u, v)|}{m_i \times n_i}, \quad \sigma_i = \sqrt{\frac{\sum_{u=1}^{m_i} \sum_{v=1}^{n_i} (dwt_i(s, u, v) - \mu_i)^2}{m_i \times n_i}}. \quad (3)$$

The WT-based features in [14] (WT-XYZ) are extracted using (2) for the computed X, Y and Z panoramic views of the object (Section 3.1).

4 Muscle Shape Classification

Given a mesh representing a segmented muscle from a novel 3D medical image, our goal is to distinguish abnormal muscles from healthy muscles. In particular, we would like to perform 4 independent classification tasks for the 4 muscles: rectus femoris (RF), vastus lateralis (VL), biceps femoris-short (BS), and sartorius (SS), which belong to the muscle groups: knee extensors, knee flexors, and hip adductors (Table 2). In order to evaluate the accuracy of the classifier, we perform a leave-one-out (LOO) cross-validation [18]. We use a non-linear SVM for thigh muscle classification, which requires the setting of two parameters: C , which assigns a penalty to errors, and γ , which defines the width of a radial basis function [18]. We compute the false positive (FP) and true positive (TP) rates of the classifier for different values of C and γ in a logarithmic grid search (from 2^{-8} to 2^8) to create a receiver operating characteristic (ROC) curve. Therefore, each pair of the parameters (C_i, γ_j) would generate a point (FP_{ij}, TP_{ij}) in the graph. The ROC curve is constructed by selecting the set of optimal operating points. Point (FP_{ij}, TP_{ij}) is optimal if there is no other point (FP_{mn}, TP_{mn}) such that $FP_{mn} \leq FP_{ij}$ and $TP_{mn} \geq TP_{ij}$. We use the area under the generated ROC curves (AUC) obtained from classification involving different shape descriptors to compare their discriminatory power.

5 Results

Figure 5 and Table 2 show the ROC curves and the areas under them computed for the global shape descriptors used in [20] and [10], which we denote by GLOB1 and GLOB2, respectively (Section 1), and WT-based shape descriptors (WT-XYZ, WT-Z, WT-CYL, WT-SPH and WT-MED described in Section 3) to classify 4 individual thigh muscles into normal vs. COPD cases.

Comparing the classification accuracies reported in Table 2, we make the following observations:

- i. Averaged over all the 4 muscles (column 1 in Table 2), the highest SVM classification accuracy is obtained using WT-MED, where individual muscles in the COPD group are differentiated from those in the healthy group with an average classification accuracy of 93.69%.
- ii. WT-based shape descriptors (WT-XYZ, WT-Z, WT-CYL, WT-SPH and WT-MED) outperform GLOB1 [20] and GLOB2 [10].
- iii. The classification accuracy of the 1D WT-based descriptors (WT-Z, WT-CYL, WT-SPH and WT-MED) is greater than that of the 3D descriptor (WT-XYZ) [14]. It can be seen that, for any muscle, we can find at least one 1D shape descriptor with accuracy better than 3D.
- iv. WT-based descriptors using the proposed *mesh projection* outperform the WT-Z using the *cylindrical projection*.
- v. The last column shows our recommended shape descriptor for each muscle.

Furthermore, Tables 3-4 show the recall $R = TP/(TP + FN)$ and precision $P = TP/(TP + FP)$ resulting from the optimum setting of the parameters (C, γ)

in the SVM-classifier. The results indicate that, on average over all the 4 muscles, the maximum precision and recall are achieved by our proposed WT-MED.

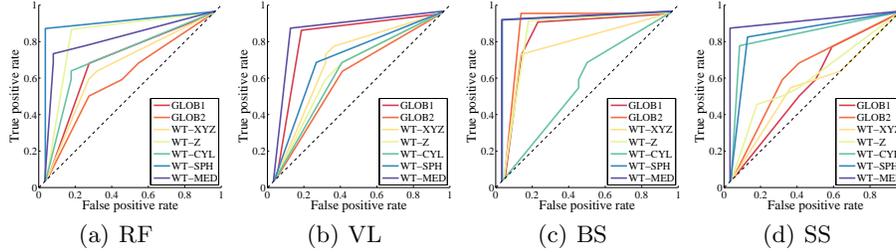


Fig. 5: ROC curves of the classifiers resulting from using the different shape descriptors for the 4 muscles. Areas under the ROC curves are reported in Table 2.

Muscle	Area under the ROC curves							Selected Descriptor(s)
	GLOB1	GLOB2	WT-XYZ	WT-Z	WT-CYL	WT-SPH	WT-MED	
RF	0.7250	0.6200	0.6837	0.8750	0.7462	0.9500	0.8475	WT-SPH
VL	0.6188	0.5938	0.7500	0.5413	0.7900	0.7250	0.9750	WT-MED
BS	0.9025	0.9500	0.8250	0.9000	0.5813	0.9750	0.9750	WT-SPH, WT-MED
SS	0.5850	0.6650	0.5650	0.6225	0.8750	0.8750	0.9500	WT-MED
Mean	0.7078	0.7072	0.7059	0.7347	0.7481	0.8812	0.9369	

Table 2: Area under the ROC curves in Figure 5. Highest accuracy acquired for each muscle is colored in red. The last column shows, for each muscle, the shape descriptor(s) that resulted in the highest AUC.

Muscle	Precision of the classifier							Selected Descriptor(s)
	GLOB1	GLOB2	WT-XYZ	WT-Z	WT-CYL	WT-SPH	WT-MED	
RF	0.7143	0.5930	0.6594	0.8947	0.6942	0.9091	0.7804	WT-SPH
VL	0.5955	0.5544	0.7500	0.5291	0.7148	0.7647	0.9524	WT-MED
BS	0.8619	1.0000	0.7826	0.9444	0.5846	0.9524	0.9524	WT-MED
SS	0.5919	0.6515	0.5455	0.5873	0.8261	0.8571	0.9091	WT-MED
Mean	0.6909	0.6997	0.6844	0.7389	0.7049	0.8708	0.8986	

Table 3: Precision values resulting from using the optimum setting of the parameters (C, γ) in SVM-classifier.

6 Discussion and Conclusion

A computer-aided method for diagnosing muscle atrophy in people with COPD could facilitate targeting of interventions such as strength training or gene therapy. In order to differentiate 4 individual thigh muscles in the healthy versus COPD groups, we first applied a state-of-the-art 3D shape descriptor; the WT-based shape descriptor proposed by Papadakis et al. [14] resulting in cylindrical projections. A comparison between the classification accuracies obtained by the

Muscle	Recall of the classifier							Selected Descriptor(s)
	GLOB1	GLOB2	WT-XYZ	WT-Z	WT-CYL	WT-SPH	WT-MED	
RF	0.7000	0.6000	0.6250	0.9000	0.6250	0.9000	0.7500	WT-SPH
VL	0.5250	0.4500	0.7500	0.4500	0.6500	0.8000	0.9500	WT-MED
BS	0.8500	1.0000	0.7500	0.9500	0.6167	0.9500	0.9500	WT-MED
SS	0.6500	0.6500	0.6000	0.5000	0.8000	0.8500	0.9000	WT-MED
Mean	0.6813	0.6750	0.6813	0.7000	0.6729	0.8750	0.8875	

Table 4: Recall values resulting from using the optimum setting of parameters (C, γ) in SVM-classifier.

aforementioned descriptors and the global shape descriptors by Ward et al. [20] and HajGhanbari et al. [10] shows that, averaged over all the 4 muscles, the WT-based shape descriptors outperformed the global shape descriptors.

We extended the WT-based descriptors by introducing medial-based mesh projections rather than the cylindrical projections, in order to follow the natural geometry of the muscle more accurately, and to rely on a 1D projection rather than 3D projections. The experimental results showed that we achieved improved classification results for all the 4 muscles using the extended descriptors. These results support the shape- and axis-specific use of shape descriptors for diseased muscles. Because different chronic diseases can have a differential impact on certain muscles or specific muscle regions, muscle-specific shape descriptors should be applied to better discriminate muscle anomalies.

Although the presented descriptors were applied to differentiate thigh muscles, they might have a widespread application for other conditions and chronic diseases that result in muscle atrophy such as chronic heart diseases, AIDS, cancer, and osteoarthritis [1].

One direction for future work is to extend the medial-based, mesh-projection descriptors to complex anatomical shapes that exhibit medial branches in their skeletons (in contrast to the single medial curve that we used here for the tube-like thigh muscles). Possible future work on muscle shape analysis include establishing point correspondence between meshes and building statistical shape models [5], examining alternative shape representation via harmonic analysis [8,13], and studying shape spaces on non-linear manifolds [11,7].

Acknowledgements

We would like to thank Oliver Van Kaick at Simon Fraser University for contributing his classification code.

References

1. J. Argiles, S. Busquets, A. Felipe, and F. Lopez-Soriano. Muscle wasting in cancer and ageing: cachexia versus sarcopenia. *Adv Gerontol*, 18:39–54, 2006. 11
2. S. A. Broughton and K. M. Bryan. *Discrete Fourier Analysis and Wavelets: Applications to Signal and Image Processing*. Wiley-Interscience, 2008. 7

3. S. Chen, P. Lestrel, W. Kerr, and J. McColl. Describing shape changes in the human mandible using elliptical Fourier functions. *Europ. J. of Orthodontics*, 3:205–216, 2000. [2](#)
4. C. Clark, L. Cochrane, E. Mackay, and B. Paton. Skeletal muscle strength and endurance in patients with mild COPD and the effects of weight training. *Eur Respir J*, 15(1):92–97, 2000. [1](#)
5. R. Davies, C. Twining, T. Cootes, and C. Taylor. Building 3D statistical shape models by direct optimization. *IEEE TMI*, 29(4):961–981, 2010. [11](#)
6. S. Durrleman, X. Pennec, A. Trounev, and N. Ayache. Statistical models of sets of curves and surfaces based on currents. *Medical Image Analysis*, 13(5):793–808, 2009. [2](#)
7. P. Fletcher, C. Lu, S. Pizer, and S. Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE TMI*, 23(8):995–1005, 2004. [11](#)
8. G. Gerig, M. Styner, D. Jones, D. Weinberger, and J. Lieberman. Shape analysis of brain ventricles using SPHARM. 1:171–178, 2001. [11](#)
9. B. Gutman, Y. Wang, L. Ming Lui, T. E. Chan, and P. M. Thompson. Hippocampal surface analysis using spherical harmonic function applied to surface conformal mapping. *IEEE ICPR*, 3:964–967, 2006. [2](#)
10. B. HajGhanbari, G. Hamarneh, N. Changizi, A. Ward, and W. D. Reid. MRI-based 3D shape analysis of thigh muscles: Patients with chronic obstructive pulmonary disease versus healthy adults. *Academic Radiology*, 18(2):155–166, 2011. [1](#), [2](#), [3](#), [4](#), [9](#), [11](#)
11. S. Kurtek, E. Klassen, Z. Ding, and A. Srivastava. A novel riemannian framework for shape analysis of 3D objects. *IEEE CVPR*, pages 1625–1632, 2010. [11](#)
12. S. Mathur, K. Takai, D. MacIntyre, and D. Reid. Estimation of thigh muscle mass with magnetic resonance imaging in older adults and people with COPD. *American Physical Therapy*, 88:219–230, 2008. [1](#)
13. D. Nain, S. Haker, A. Bobick, and A. Tannenbaum. Multiscale 3D shape analysis using spherical wavelets. *MICCAI*, pages 459–467, 2005. [11](#)
14. P. Papadakis, I. Pratikakis, T. Theoharis, and S. Perantonis. PANORAMA: A 3D shape descriptor based on panoramic views for unsupervised 3D object retrieval. *IJCV*, 89:177–192, 2010. [2](#), [3](#), [4](#), [6](#), [8](#), [9](#), [10](#)
15. S. M. Pizer. Guest editorial-medial and medical: A good match for image analysis. *IJCV*, 55(2):79–84, 2003. [7](#)
16. K. Siddiqi and S. Pizer. *Medial Representations: Mathematics, Algorithms and Applications*. Springer, 2007. [7](#)
17. O. V. Kaick, G. Hamarneh, A. Ward, M. Schweitzer, and R. Zhang. Learning Fourier descriptors for computer-aided diagnosis of the supraspinatus. *Academic Radiology*, 17(8):1040–1049, 2010. [2](#)
18. V. Vapnik. *Statistical Learning Theory*. Wiley, 1998. [9](#)
19. S. Wang, M. K. Chung, K. Dalton, A. Alexander, and R. Davidson. Automated diagnosis of autism using Fourier series expansion of corpus callosum boundary. *Human Brain Mapping Conf*, 1:219–230, 2007. [2](#)
20. A. Ward and G. Hamarneh. 3D shape analysis of the supraspinatus muscle: A clinical study of the relationship between shape and pathology. *Academic Radiology*, 14:1229–1241, 2007. [2](#), [3](#), [9](#), [11](#)
21. B. Yeo, P. Yu, P. Grant, B. Fischl, and P. Golland. Shape analysis with overcomplete spherical wavelets. *MICCAI*, 11:468–476, 2008. [2](#)

Geometric calibration between PET scanner and structured light scanner

Martin Kjer¹, Oline V. Olesen¹²³, Rasmus R. Paulsen¹, Liselotte Højgaard²,
Bjarne Roed³, and Rasmus Larsen¹

¹ Informatics and Mathematical Modelling, Technical University of Denmark
Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark

<http://imm.dtu.dk/>

² Department of Clinical Physiology, Nuclear Medicine & PET, Rigshospitalet,
Copenhagen University Hospital, University of Copenhagen

³ Siemens Healthcare, Siemens A/S, Denmark

Abstract. Head movements degrade the image quality of high resolution Positron Emission Tomography (PET) brain studies through blurring and artifacts. Many image reconstruction methods allow for motion correction if the head position is tracked continuously during the study. Our method for motion tracking is a structured light scanner placed just above the patient tunnel on the High Resolution Research Tomograph (HRRT, Siemens). It continuously registers point clouds of a part of the patient's face. The relative motion is estimated as the rigid transformation between frames.

A geometric calibration between the HRRT scanner and the tracking system is needed in order to reposition the PET listmode data or image frames in the HRRT scanner coordinate system. This paper presents a method where obtained transmission scan data is segmented in order to create a point cloud of the patient's head. The point clouds from both systems can then be aligned to each other using the Iterative Closest Point (ICP) algorithm.

Keywords: HRRT, PET, structured light, calibration, motion tracking, motion correction

1 Introduction

Technological improvement of the different medical imaging modalities leads to diagnostic images with increasing spatial resolution. As a consequence the techniques also become more vulnerable to the effects of patient motion during image acquisition.

In Positron Emission Tomography (PET) patient movements can cause both artifacts and blurred images [1]. The increased spatial resolution gained by technological advancement is thus countered to a certain degree by the increased sensitivity to motion, unless patient fixation and motion correction is utilised.

Even with fixation methods such as vacuum cushions and head restraints motion still occurs albeit to a lesser degree [2]. The magnitude of motion tends

to increase with the duration of the study, and acquisition times for PET images can be up to several hours. Typically the patient’s head drifts slowly to one side, or at some point the patient repositions themselves to lie more comfortably. The resolution of the Siemens High Resolution Research Tomograph (HRRT) is below 2 mm, and since the described movements can be larger motion correction becomes a necessity [3].

One approach for motion tracking is the Polaris Vicra system from *Northern Digital Incorporated*. It registers a tool with reflective markers attached to the patient’s head. The positions of the markers are relayed to a tool tracker through infrared light. The main issue with such system is to ensure that the tool stays attached and do not move relative to the patient’s head. Further to maintain line of sight between tool and tool tracker, which is troublesome in the narrow patient tunnel of the scanner.

Our approach is a structured light system. Two cameras on both sides of a DLP pico projector from *Texas Instruments* are mounted on the HRRT scanner as shown in Figure 1(a). A series of cosine patterns are projected onto the object, and these patterns are imaged by the cameras. We use phase shift interferometry to obtain a 3D point cloud of the object - in this case a part of the patient’s face as shown in Figure 1(b) [4]. The relative motion between image frames is estimated as the rigid transformation with six degrees of freedom that best aligns the two point clouds. The iterative closest point (ICP) algorithm can be used to find the transformation [11] [12]. In comparison to the tool tracking approach, this approach avoids the use of an optical tool, and can potentially be integrated into future scanners.

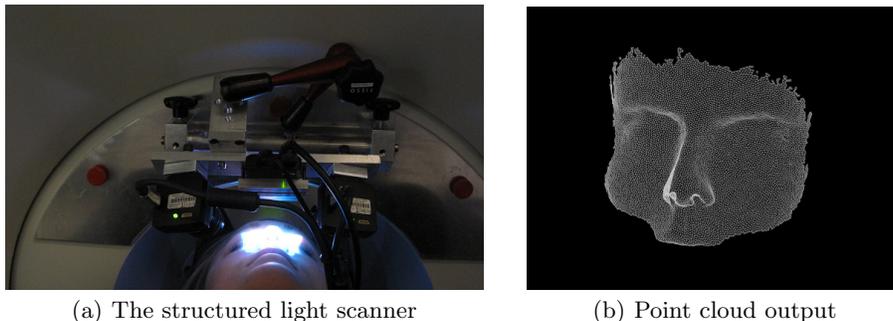


Fig. 1. The structured light scanner mounted on the HRRT scanner and an example of the 3D point cloud it produces.

The issue addressed in this paper is the geometric calibration between the HRRT scanner and the structured light system. Movements observed by the motion tracking system have to be translated into movements in the coordinate system of the HRRT scanner. Then it is possible to reposition all the detected Line Of Responses (LOR) into the position of the head at the given time.

The method for calibration should not add any extra radiation dose to the patient. Furthermore it is undesirable to increase the total duration of the scanning session and to alter the normal workflow. Ideally the method should only employ the data that is already obtained - either the emission (EM) data or the transmission (TX) data.

Previous calibration methods use either a number of EM scans or TX scans of a calibration object and have both successfully been used for finding a calibration transformation [7] [8]. In both cases the motion tracking system was the Polaris Vicra system or a system very similar to it. In the EM approach a positron emitting point source added to the tracking tool allowed for measurements of the tracking tool position in both systems. Multiple independent measurements were required in order to determine all six degrees of freedom. In the TX scan approach retroreflective markers with a sufficient density allowed for detection of the tracking tool in both systems. Both of these methods find relatively few points with a known point-to-point correspondance from which a transformation can be estimated. The measurements must be performed in preparation of the patient scan, and the calibration is preserved as long as the tooltracker is not moved.

For the purpose of calibration with a structured light motion tracking system there are no markers to be detected with either EM scans or TX scans. From the TX data it is however possible to extract the iso-surface of the patient's head, thus producing a point cloud similar to what is obtained from the motion tracking system. The point correspondance is not known, however a large amount of points are available. The best rigid alignment between the two point clouds can be found with the ICP algorithm, and the transformation serves as the calibration. The measurements are a part of the normal scanning procedure, which is very advantageous in terms of time and simplicity. This also allows for adjustments of the motion tracking system or even completely detaching it from the HRRT scanner between scans.

A common approach for extracting iso-surfaces from volume data is the Marching Cubes Algorithm [5]. However, in our case, the iso-levels of the TX scans are not very well defined and would result in a noisy surface. We have therefore investigated an alternative approach to extract the interface between soft-tissue and air from the TX scans. This paper presents a segmentation method using path tracing on the reconstructed TX image.

2 Methods

2.1 Circular resampling

A typical slice from the TX image of a patient is seen in Figure 2(a). The border of the head has to be traced and the procedure repeated for each slice. The used path tracing algorithm finds the optimal path going from one edge of the image to the other. However the boundary of the head in the TX image is located as a circular structure somewhere in the middle, and thus a reshaping is required before path tracing can be applied.

A point within the circular structure is chosen. The point (r_s, c_s) serves as a center from which N spokes of length L radially shoots out from, so that the end points are given as:

$$r_e = r_s + \cos\left(\frac{2\pi n}{N}\right)L, \quad c_e = c_s + \sin\left(\frac{2\pi n}{N}\right)L$$

where $n = [0, 1, \dots, N - 1]$. Each spoke samples L values of the underlying pixels using bi-linear interpolation. The sampled values are placed in a new image, as shown in Figure 2(b).

The center point is chosen as the centroid of the image. Using a fixed point could however be a viable strategy. Patients are always placed in the approximate radial center of the scanner tunnel, since this is where the spatial resolution is highest [6]. The placement is done manually so a slight inter-scan variation is expected. Using a fixed point saves a little computational time, however it is more likely to give a faulty resampling, if the spoke length is chosen too short.

The radial resolution - the amount of spokes - is also a consideration. The choice is dependent on the resolution of the TX image. Too few spokes correspond to an undersampling and lead to loss of information about the true curvature and important small features primarily the nose and ears. Oversampling increases computational time without much extra information being gained.

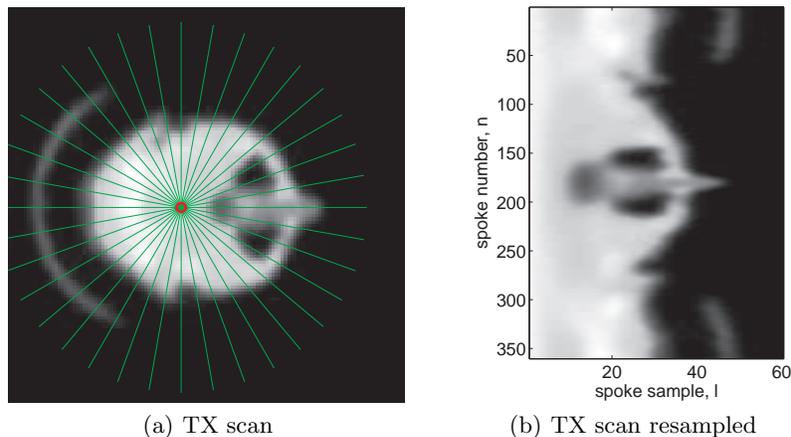


Fig. 2. Circular resampling of a transversal TX slice using 360 spokes. Only every 10th spoke is displayed on Figure 2(a).

2.2 Path tracing

The chosen path tracing algorithm is a simplified version of Dijkstra's algorithm [9]. It is based on dynamic programming and designed to find optimal paths

between the top and bottom of a grayscale image $\mathbf{I}(r, c)$ as illustrated in Figure 3.

Each pixel holds a cost value $C(r, c)$, and for the purpose of edge detection the first derivative or gradient of the image is used. The optimal path \mathbb{P} is then defined as the list of pixels with the lowest (or highest) accumulated cost going from the top of the image to the bottom:

$$C_{tot} = \sum_{(r,c) \in \mathbb{P}} C(r, c)$$

It is possible to calculate the total cost for all paths, however the number of computations would quickly increase beyond reasonable with increasing image size. The algorithm is therefore limited, so that the path is only allowed to move down and up to two pixels to either side.

The algorithm operates with two images of the same dimension as the original image. A value in the accumulator image $\mathbf{A}(r, c)$ is the lowest possible (optimal) accumulated cost required to get to that pixel from the top. Path information is stored in the backtracing matrix \mathbf{T} . The value stored in $\mathbf{T}(r, c)$ is the column index of the previous path entry, $\mathbb{P}(r - 1, \mathbf{T}(r, c))$. The path is therefore read backwards, and the last path entry is the index that holds the lowest value in the last row of \mathbf{A} . This is illustrated in Figure 3.

An important aspect of the algorithm is the ability to wrap around small features. This is highly dependent on the restrictions imposed on the path. Allowing it to sidestep more pixels allows for more sharp features to be traced, although it could also lead to a much less smooth border. In this particular case an alternate possibility is to increase the radial resolution of the resampling, at the expense of more computations.

The traced edge points from each slice are transformed to points in the original image and combined to the resulting point cloud which is shown in Figure 4.

2.3 Transformations and ICP

The calibration between the motion tracking system and the HRRT scanner can be taken as the rigid transformation of the source point set $\mathbb{P} = (p_1, p_2, \dots, p_N)$ that gives the best alignment to the target set $\mathbb{Q} = (q_1, q_2, \dots, q_M)$. The mathematical measure of goodness of fit is the sum of squared errors after the transformation - which is the squared distance from points in \mathbb{P} to their corresponding point in \mathbb{Q} :

$$E = \sum_{i=1}^N \|\mathbf{R}p_i + \mathbf{T} - q_i\|^2 \quad (1)$$

where \mathbf{R} is an orthogonal rotation matrix and \mathbf{T} a translation vector.

Minimization of Equation 1 requires knowledge of the point correspondance between the two point sets. When such information is not given, the problem can be solved with an iterative approach - the ICP. The algorithm iterates through the following steps:

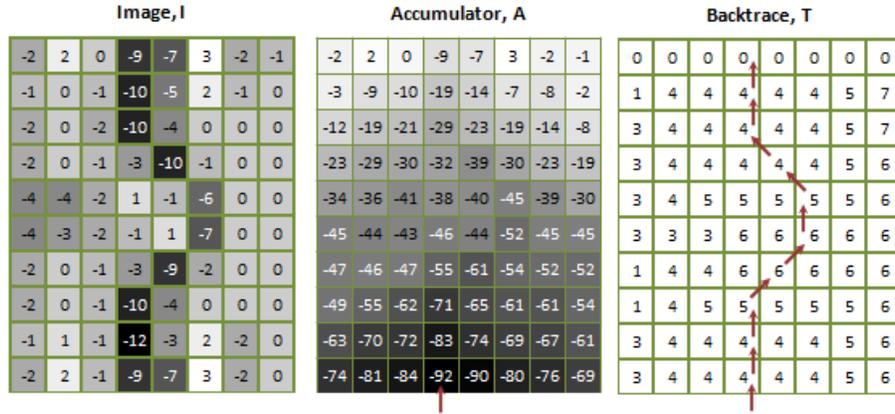


Fig. 3. Illustration of pathtracing on the resized version of Figure 2(b).

- **Matching:** Points in \mathbb{P} are matched with their nearest neighbor in \mathbb{Q} , and this is assumed to be the point correspondance.
- **Minimization:** Equation 1 is minimized using Singular Value Decomposition (SVD).
- **Transformation:** The transformation is applied to the points in \mathbb{P} , and the steps are repeated.

The algorithm can be modified further and improved by adding more steps, such as a selection of only a subset of the points or inclusion of a neighbor pair weighting [10].

When using the ICP there is a consideration of the designation of target and source point cloud. The point cloud from the TX image represents the entire head whereas the points from the motion tracking system only represent a subset of this. Many of the points from the TX image does therefore not have a meaningful nearest neighbor correspondance, and consequently the TX point cloud is chosen as the target. Otherwise the ICP algorithm requires some kind of rejection scheme.

3 Results and Discussion

3.1 Segmentation

Transmission data from four patient studies was available and two of them were motion tracked with the structured light system. The settings and parameters for segmentation of the TX data are based on the two untracked TX studies. The result is illustrated in Figure 4, and while it is recognisable as a face, it does have the issues to be addressed: Which slices to segment and edge tracing.

Which slices to segment: The data contains 207 transversal slices, however the initial slices contain nothing except noise and attenuation from the headrest.

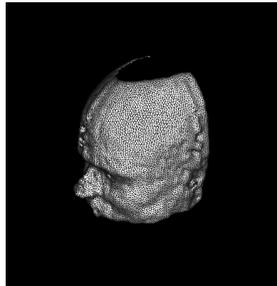


Fig. 4. The output from the TX segmentation. The point cloud was surface reconstructed [13] for the purpose of illustration. Notice that it is possible to see the start of the ear canal, however the ear is missing. Small indentations and bumps are present and the very top and the back-side of the head has been excluded.

The strategy with circular resampling and path tracing assumes that the image contains a circular structure, and otherwise the result is highly unpredictable. A simple threshold strategy was chosen as sanity check. The upper part of the skull is excluded, as seen in Figure 4. The loss is acceptable, since the excluded part is neither seen by the motion tracking system. More slices could be included with a more sophisticated sanity check.

The edge tracing: The cost values used to determine the optimal path are based on the derivative of the γ -ray attenuation. For the most part of the head there is a superficial thin layer of skin followed by the bone layer. Since bone is highly attenuating compared to soft tissues, there is a gradient going from air to skin, but an even steeper gradient going from skin to bone. The method thus favors segmenting the bone edges, which for the most part of the face is very similar to the skin surface. However it is paramount that the method segments important features - especially the nose. The upper part of the nose is quite dense and thus segmented, whereas the lower soft part is not. The resulting faulty segmentation is illustrated in Figure 5(a), and therefore it was necessary to include a weighting of the cost values to improve the segmentation.

A strategy was chosen where all cost values outside the bone edge region were enhanced. This generally improved the skin and nose segmentation as shown in Figure 5(a). However the lower most part of the nose and the ears were still unsegmented. Further the enhancement introduced some artifacts. The gradient from the headrest attenuation was also amplified, resulting in partly segmentation of this as shown in Figure 5(b). The issue was most profound in slices where the head was in direct contact with the headrest. Since the back of the head has no real interest it could be cut-off. Notice also in Figure 5(b) that even though the enhancement improved the tracing of the skin-border, the path occasionally went back to the bone-border resulting in bumps and indentations in the point cloud.

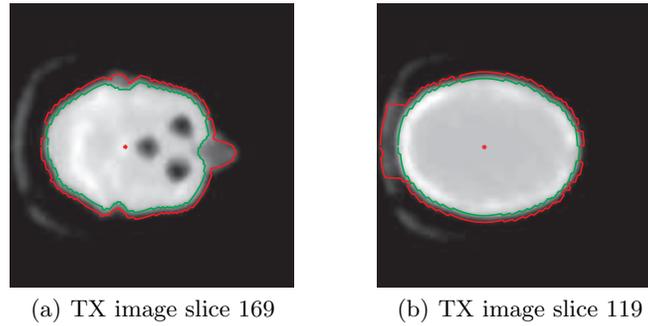


Fig. 5. Illustration of the effect of the added weighting scheme for segmentation. The green and the red lines show the result before and after application of the cost value weighting respectively.

3.2 Calibration

With one point cloud measured from the HRRT scanner and one from the motion tracking system the calibration could be performed. The initial situation is seen in Figure 6(a). A fixed transformation was applied in order to bring the two point clouds into a decent starting position. This was chosen as a rotation of 180° around the y -axis, a rotation of -36° around the x -axis and a (x,y,z) -translation of 160 mm, 385 mm and 340 mm. The resulting situation before the application of the ICP is shown in Figure 6(b). The ICP converged to a minimum with less than 20 iterations, and the final total transformation is shown in Figure 6(c).

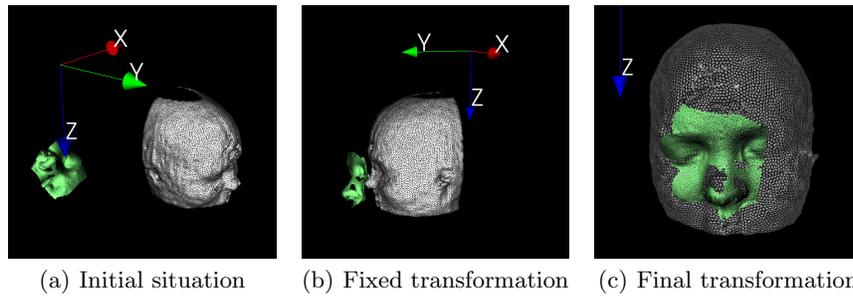


Fig. 6. Illustration of the transformations. The point cloud from the TX scan is shown in white and has been surface reconstructed [13] for illustration purpose. The point cloud from the motion tracking system is shown in green.

Visually it appears to be a decent alignment, however a more quantitative measure is required for validation of the result. One approach could be to use landmarks, however the error would then be correlated with the ability to place

landmarks correctly. A more precise and valuable validation would be a comparison between PET images with and without motion correction, however the motion correction of PET images is beyond the scope of this paper.

3.3 Sources of error

The segmentation method assumes that the two point clouds are measured simultaneously. However the cameras' image capture is instantaneous compared to the time that the TX scan requires. The assumption is that the patient is motionless in this period. This is not likely completely true and could be a potential source of error.

Arguably a minor error for cooperative subjects, and the assumption must be generally accepted - since the TX data is deemed good enough to be used in the emission image reconstruction.

A great perspective for motion tracking is however to use it with patients that are less cooperative - children and patients suffering from disorders that affect motor-function. In that case the error in the TX data increases, and the chosen segmentation method for calibration would perform worse.

4 Conclusion

A method for geometric calibration between the HRRT scanner and the structured light motion tracking system has been presented. The method exploits the data from the transmission scan, and thus it does not alter the scanning procedure or prolong the study. This is in key with the advantages of using the structured light system.

The presented approach segments the edges of the transmission scan slices using a path tracing algorithm, from which a point cloud of the patient's head is obtained. It was shown that the approach suffers from some difficulties due to the attenuation properties of the different tissues. An enhancement strategy for improved segmentation was proposed, and it was shown to give a better segmentation of the nose at the expense of more problematic segmentations in regions that have no interest. The method is not able to produce perfect point clouds of the entire head, however it can produce decent point clouds of the face, and thus the method seems suitable for the purpose of calibration with the structured light system.

The motion tracking system produces a similar point cloud, and it was shown that the two point clouds can be aligned to each other using the ICP algorithm with a visually good result.

References

1. Anton-Rodriguez, J.M. , Sibomana, M., Walker, M.D., Huisman, M.C., Matthews, J.C., Feldmann, M., Keller, S.H., Asselin, M.: Investigation of Motion Induced Errors in Scatter Correction for the HRRT Brain Scanner. IEEE Nuclear Science Symposium Conference Record (MIC), IEEE, (2010)

2. Green, M.V., Seidel, J., Stein, S.D., Tedder, T.E., Kempner, K.M., Kertzman, C., Zeffiro, T.A.: Head Movement in Normal Subjects During Simulated PET Brain Imaging with and without Head Restraint. *Journal of Nuclear Medicine* 35(9), pp. 1538–1546 (1994)
3. Olesen, O.V., Sibomana, M., Keller, S.H., Andersen, F., Jensen, J., Holm, S., Svarer, C., Højgaard, L.: Spatial resolution of the HRRT PET scanner using 3D-OSEM PSF reconstruction. *Nuclear Science Symposium Conference Record 2009*, pp. 3789–3790 (2009)
4. Olesen, O.V., Paulsen, R., Højgaard, L., Roed, B., Larsen, R.: Motion tracking in narrow spaces: a structured light approach. *Medical image computing and computer-assisted intervention : MICCAI. International Conference on Medical Image Computing and Computer-Assisted Intervention*, 13(Pt 3), pp. 253-60 (2010)
5. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. *SIGGRAPH Comput. Graph.* 21(4), pp. 163-169 (1987)
6. Knöß, C.: PhD Thesis: Evaluation and Optimization of the High Resolution Research Tomograph (HRRT). RWTH Aachen Univeristy (2004)
7. Fulton, R.R., Meikle, S.R., Eberl, S., Pfeiffer, J., Constable, C.J. and Fulham, M.J.: Correction for head movements in positron emission tomography using an optical motion-tracking system. *IEEE Transactions on Nuclear Science* 49(1), pp. 116–123 (2002)
8. Bühler, P., Just, U., Will, E., Kotzerke, J. and Hoff, J.: An accurate method for correction of head movement in PET. *IEEE Transactions on Medical Imaging* 23(8), pp. 1176-1185 (2004)
9. Paulsen, R.R. and Moeslund, T.B.: *Introduction to Medical Image Analysis*. DTU Informatics (2011)
10. Rusinkiewicz, S. and Levoy, M.: Efficient Variants of the ICP Algorithm. *Third International Conference on 3-D Digital Imaging and Modeling*, pp. 145 (2001)
11. Chen, Y., Medioni, G.: Object Modeling by Registration of Multiple Range Images. *1991 IEEE International Conference on Robotics and Automation* 3, pp. 2724–2729 (1991)
12. Besl, P.J., McKay, N.D.: A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(2), pp. 239–256 (1992)
13. Paulsen, R.R., Bærentzen, J.A., Larsen, R.: Markov Random Field Surface Reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, pp. 636–646 (2009)

CLARCS, a C++ Library for Automated Registration and Comparison of Surfaces: Medical Applications

Alexandre Abadie^{1,2,3}, Benoît Combès⁴, Claire Haegelen^{1,2,3,5},
Sylvain Prima^{1,2,3}

¹ INSERM, U746, Faculty of Medicine CS 34317, F-35043 Rennes, France

² INRIA, VisAGeS Unit/Project, F-35042 Rennes, France

³ University of Rennes I, CNRS, UMR 6074, IRISA, F-35042 Rennes, France

⁴ INRIA, VisAGeS Unit/Project, F-35042 Rennes, France

⁵ Department of Neurosurgery, Pontchaillou University Hospital, F-35043 Rennes, France

Abstract. In this paper we present the methods implemented in the CLARCS (C++ Library for Automated Registration and Comparison of Surfaces) library. This library allows some basic and high level processing on free-form surfaces, represented as point sets or meshes. Three methods are the “building bricks” of CLARCS; they allow (i) the rigid/affine/non-linear registration of two point sets, (ii) the computation of the mid-sagittal plane of one point set, (iii) the computation of a mean point set from several point sets, and the variability around this mean. These methods are all based on a common methodological framework, in which the point sets/meshes are represented either as a Gaussian mixture model or as a draw of such a model. We propose some applications of the methods implemented in CLARCS on different sets of medical data.

1 Introduction

In medical image processing, the most widely used methods are voxel-based, which means that their required input data must be (most often) 3D arrays of grey values. Some important issues with this kind of data include the large memory needed to store them, the large run time of standard algorithms to process them, the choice of data type to code the grey value of a voxel or the sometimes problematic orientation of the volumes. Thanks to the increase of 3D rendering and computational capacities of computers in the last few years, new surface based processing methods have emerged as an alternative to standard voxel-based techniques. In this context, the VisAGeS team at INRIA/INSERM (<https://www.irisa.fr/visages>) has proposed a set of computational tools that have been implemented in a software library called CLARCS (C++ Library for Automated Registration and Comparison of Surfaces).

In this paper, we outline the methods implemented in CLARCS (Section 2) and we provide some potential medical applications of these (Section 3): assessment of dysmorphology in craniosynostosis and plagiocephaly and construction of a statistical shape model of thalami for deep brain stimulation (DBS).

2 Overview of CLARCS

In the first subsection below, we briefly describe some of the best known libraries or software working on point sets, meshes or surfaces. In the second subsection, we describe the functions implemented in CLARCS.

2.1 Related software for surface processing

There exists a bunch of commercial software allowing for some basic (*e.g.* import/export, visualisation) and higher level (*e.g.* registration) processing on surfaces, such as Amira (<http://www.amira.com>), or Rapidform (<http://www.rapidform.com>), but it is often difficult to know what methods are implemented therein, and always costly to acquire the software.

Some freeware are specifically dedicated to high level geometric computation (*e.g.* CGAL, <http://www.cgal.org>) or high level visualisation (*e.g.* ParaView, <http://www.paraview.org>, and VTK, <http://www.vtk.org>) of surfaces, but we restrict our brief overview here to software which allow surface registration/comparison in a broad sense. The Point Cloud Library (or PCL, <http://pointclouds.org>) and MeshLab (<http://meshlab.sourceforge.net>) are two such software/libraries, but they only implement rigid registration. FreeSurfer (<http://surfer.nmr.mgh.harvard.edu>) allows non-linear registration of surfaces, but is limited to brain (cortical) data (using a specific atlas for this purpose).

Some freely available state-of-the-art software with a larger range of applications include:

- TPS-RPM [4]: <http://noodle.med.yale.edu/~chui/tps-rpm.html>
- TPS-L2 [15]: <http://code.google.com/p/gmmreg>
- CPD [20]: <https://sites.google.com/site/myronenko/research/cpd>
- SPHARM [22]: <http://www.enallagma.com/SPHARM.php>
- weighted-SPHARM [5]: <http://www.stat.wisc.edu/~mchung/software/software.html>

We have experimentally found the first three methods to be most often unable to cope with large (more than 100,000 points) point sets, while the two others impose strong topological constraints on the surfaces to register (they must be closed). The methods implemented in CLARCS allow the processing of large point sets (often needed to encode highly convoluted/complex anatomical structures) without topological constraints, in an efficient and fast way. From a methodological point of view, a strong advantage of these methods over most of

the abovementioned ones is the unified computational framework they are based on, as outlined below. Finally, to our knowledge, there is no freely available software allowing the computation of the symmetry plane of approximately bilateral surfaces, and the assessment of asymmetries thereof.

2.2 CLARCS

CLARCS is a C++ library for surface processing and analysis which has been developed at IRISA (<http://www.irisa.fr>) since 2007. CLARCS is mainly composed of three basic algorithms allowing (i) the rigid/affine/non-linear registration of two point sets, (ii) the computation of the mid-sagittal plane of one point set, (iii) the computation of a mean point set from several point sets, and the variability around this mean. These algorithms are building bricks that can be combined to allow higher-order surface processing such as computation of pointwise asymmetry fields and statistical analysis thereof within or between populations. The three algorithms rest on a common methodological framework, that we briefly outline here. We refer the reader to the corresponding papers for a deeper insight into these algorithms. Let us first define the following pseudo-distance between point sets X^1 and X^2 :

$$\delta^2(X^1, X^2) = \min_{A, T} \left[\sum_{x_i \in X^1} \sum_{x_j \in X^2} A_{i,j} \|x_i - T(x_j)\|^2 + 2\sigma^2 \sum_{i,j} A_{i,j} \log(A_{i,j}) + \alpha L(T) \right]$$

$$\text{with } \forall i \sum_j A_{i,j} = 1 \text{ and } \forall i, j A_{i,j} \geq 0$$

Computing $\delta^2(X^1, X^2)$ is a minimisation problem involving the unknown transformation T linking the two point sets, and the unknown $\text{card}(X^1) \times \text{card}(X^2)$ matrix A . The constraints on this matrix allows it to encode the point-to-point correspondences between the point sets in a *fuzzy* way. The cost function can be seen as the sum of (i) a data-attachment term, (ii) a barrier (smoothing) function, which convexifies the cost function, and (iii) a regularity constraint on the unknown transformation. The relative strength of the three terms is weighted by the positive scalars $2\sigma^2$ and α .

It is enlightening to notice that this minimisation problem is actually the maximum a posteriori (MAP) problem consisting of finding the transformation T best superposing the two point sets, when one makes the hypotheses that (i) each point x_i of X^1 is independently drawn from a Gaussian mixture model (GMM) (whose means are the points $T(x_j)$ of $T(X^2)$, whose covariance matrices are equal to $\sigma^2 I$, I being the identity 3×3 matrix, and whose mixture weights are equal) and (ii) T is a random variable with a distribution of the form $\propto \exp(-\alpha L(T))$. With this probabilistic view, X^1 can be interpreted as a *noised* version of $T(X^2)$, σ being the standard deviation of this noise.

This MAP problem can typically be solved using the Expectation-Maximisation (EM) algorithm, which leads to the following iterative two-step algorithm:

- **E-step:** $\forall i, j \tilde{A}_{i,j} = \frac{\exp[-\|x_i - \tilde{T}(x_j)\|^2 / (2\sigma^2)]}{\sum_k \exp[-\|x_i - \tilde{T}(x_k)\|^2 / (2\sigma^2)]}$
- **M-step:** $\tilde{T} = \arg \min_T \sum_{i,j} \tilde{A}_{i,j} \|x_i - T(x_j)\|^2 + \alpha L(T)$

This EM algorithm can be shown to converge monotonically to an (at least) local maximum of the MAP criterion or, equivalently, to an (at least) local minimum of the aforementioned cost function. Actually, the two steps of the EM algorithm are exactly the same as those obtained when minimising this cost function with respect to A and T in turn.

The formulation of the E-step helps to understand why the matrix A encodes the point-to-point correspondences: when the transformation T is known, $\tilde{A}_{i,j}$ is the posterior probability that the point x_i has been drawn from the mixture component with centre $T(x_j)$. The E-step simply consists in computing these $\text{card}(X^1) \times \text{card}(X^2)$ probabilities. On the contrary, solving the M-step is highly dependent on the type of transformations considered.

Below, we show how this generic algorithm can be instantiated to lead to the three basic methods implemented in CLARCS.

Rigid/affine/non-linear registration. When T is set to be rigid and $\alpha = 0$ (*i.e.* no prior on the rigid transformation), the MAP problem boils down to a maximum likelihood (ML) problem. Several closed-form solutions exist for the M-step, using typically the unit quaternions or the singular value decomposition. This results in what was termed the EM-ICP algorithm by Granger & Pennec [14]. An earlier variant of this algorithm was devised by Rangarajan *et al.* and termed the RPM algorithm [21].

We built on the EM-ICP algorithm to propose some adaptations for non-linear registration. In previous works, we showed how to use the normals in addition to the point coordinates [19], and how to enforce some constraints on the point-to-point correspondences [9]. We also showed how to solve the M-step when considering a model of locally affine transformations [19, 9] or when using the Reproducing Kernel Hilbert Space (RKHS) theory and the Fourier analysis to build the model [10]. Finally, we showed how to obtain symmetric consistency when using this last transformation model [10].

Symmetry plane computation. When T is set to be a reflection, $\alpha = 0$ (*i.e.* no prior on the reflection) and when $X^1 = X^2$, the EM algorithm allows to compute the plane best superposing the left and right parts of $X^1 = X^2$ [7]. Needless to say, this assumes that $X^1 = X^2$ is (at least) approximately *bilateral*. Our specific contribution was to propose a closed-form for the M-step, relying on the parametrization of the unknown reflection plane using its unit normal and distance to the origin of the coordinate system [7]. Non-linear registration of the point set with its flipped image with respect to the approximate symmetry plane allows the pointwise mapping of asymmetries [8].

Atlas construction. The problem here is to find a point set M best representing a set of n point sets X^1, \dots, X^n . We defined it as:

$$M = \arg \min_X \sum_{i=1, \dots, n} \delta(X, X^i)$$

If T is set to be a similarity (rigid transformation plus uniform scaling), it is possible to devise an iterative algorithm converging to an (at least) optimum of this criterion, in which the point-to-point correspondences, the mean point set (which turns out to be the mean shape in this case where T is defined as a similarity) and the unknown transformations linking the n point sets and the mean point set are estimated in turn. However, in such an approach, point-to-point correspondences are likely to be meaningless; this is why we proposed to establish these correspondences using non-linear transformations while computing the mean shape using similarities. The resulting iterative algorithm can no longer be shown to converge, but behaves well in practice [6]. Once the algorithm has converged, it is straightforward to perform PCA on the residuals.

Implementation details. These three algorithms were implemented within a **multiscale** framework. As previously noted by Granger & Pennec, the σ parameter allows to deal with the correspondences in a fuzzy way, leading to a smoother criterion to minimise. When σ is infinitely small, their EM-ICP algorithm is simply the ICP algorithm of Besl & McKay, hence the name [2]. This leads to the idea of devising a scheme in which several EM algorithms are successively run with decreasing σ values, with a large starting value when the point sets to register are far from each other. We also used *kd*-trees, for **increased speed**, and a cut-off distance between the points x_i and $T(x_j)$, above which they are eliminated from the estimation of the transformation, for increased speed and **robustness to outliers**.

3 Applications

3.1 Quantification of skull asymmetries in craniosynostoses

Synostosis is the union of two or more separate bones to form a single bone (Merriam-Webster). Cranial synostosis, or craniosynostosis is a rare congenital disease which consists in the premature fusion of one or several cranial sutures. The last medical condition is typically met in conjunction with a hundred of syndromes, among which are Apert or Crouzon. On the contrary, the etiology of isolated (nonsyndromic) craniosynostosis is largely unknown. Early detection and treatment of craniosynostosis is crucial, as hindered skull growth can lead to increased intracranial pressure and thus alteration of normal brain development [23].

It is expected that improved characterisation of the dysmorphology of the skull associated with craniosynostosis could help understanding its cause(s), improving its diagnosis (e.g. lambdoid synostosis may be easily confounded with

deformational plagiocephaly) and even improving its treatment. In this context, of particular interest are the unilateral coronal synostoses, as it is not clear whether the left and right coronal synostoses are due to the same factors.

We propose to characterise the dysmorphology of one skull with unilateral left coronal synostosis by assessing its pointwise asymmetry. This assessment is based on the computation of an *approximate* symmetry plane, from which the left-right differences (asymmetries) can be evaluated (cf. Fig. 1). The outer surface of the skull was computed using manual grey-level thresholding, mathematical morphology and the marching cube algorithm from a CT scan. The point set is a triangular mesh of a complete skull made up of about 140,000 cells and 82,000 points.

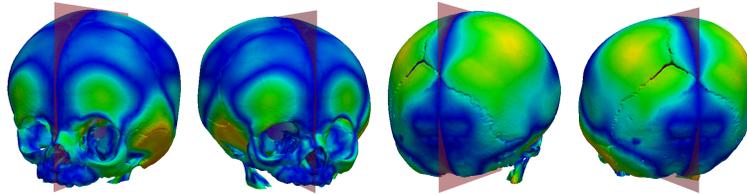


Fig. 1. Evaluation of pointwise asymmetries on a skull with coronal synostosis. We display the norm of the asymmetry field. Strong asymmetries are visible on the temporal bone, the posterior part of the parietal bone and the supraorbital part of the frontal bone.

3.2 Quantification of skull asymmetries in deformational plagiocephaly

Since the inception of the "back to sleep" recommendations in the early 90s to reduce sudden infant death syndrome, the incidence of positional (i.e. nonsynostotic) plagiocephaly has drastically increased [18, 16]. Plagiocephaly consists of the flattening of one side of the head, with aesthetics consequences, and potential altered brain development [17]. Objective assessment of plagiocephaly should help diagnosis and follow-up of this condition [3].

We propose to characterise the deformational plagiocephaly of one skull using the same computational tools as in the previous section (cf. Fig. 2). The outer surface of the skull was computed using manual grey-level thresholding, mathematical morphology and the marching cube algorithm from a CT scan. The point set is a triangular mesh of a complete skull made up of about 200,000 cells and 137,000 points.

3.3 Building statistical shape models of deep grey nuclei

Deep brain stimulation (DBS) was initially targetted to the ventral intermediate thalamus to reduce tremor in patients with Parkinson's disease [1]. Since

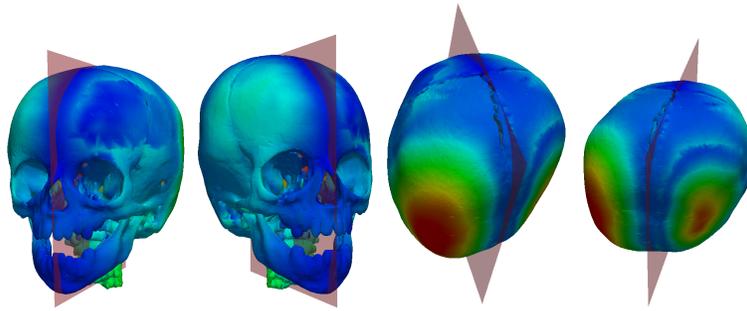


Fig. 2. Evaluation of pointwise asymmetries on a skull with deformational plagiocephaly. We display the norm of the asymmetry field. Strong asymmetries are visible on the posterior part of the parietal bone and on the superior part of the occipital bone.

then, alternative targets have emerged, such as the subthalamic nucleus (STN), globus pallidus interna, which appear to be comparatively efficient in this context [12]. The indications for DBS have also been extended to drug-resistant epilepsy, dystonia, Tourette syndrome or even obsessive compulsive disorders, using the caudate and accumbens nuclei (for instance) as targetted structures. An improved knowledge of the anatomy of these subcortical structures is key to optimise pre-operative planning and to assess treatment efficacy.

In this context, statistical shape models (SSMs) are extremely useful as they allow to help the segmentation of the structures of interest in MR images, in which typical pulse sequences make it hard to distinguish these nuclei based on the grey values alone [11].

As an illustration, we propose to build a statistical shape model of the thalami (Fig. 4) which were manually segmented together with six other deep brain structures (cf. Fig. 3) by a trained neuroanatomist using itk-SNAP (<http://www.itk-snap.org>) in 10 patients with Parkinson’s disease. For each patient, the seven pairs of segmented structures were stored in a 3D image with discrete labels. The surfaces of the structures were computed using the marching cube algorithm. The points sets representing the left and right thalami are triangular meshes with sizes of about 10,000 cells and 5,000 points.

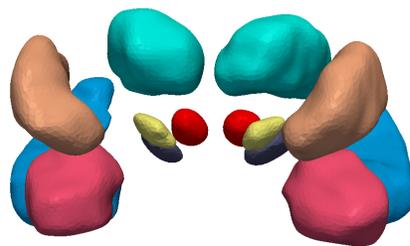


Fig. 3. Front view of seven deep brain structures. Cyan: thalamus, orange: putamen, magenta: amygdala, blue: hippocampus, red: red nucleus, yellow: STN, purple: substantia nigra.

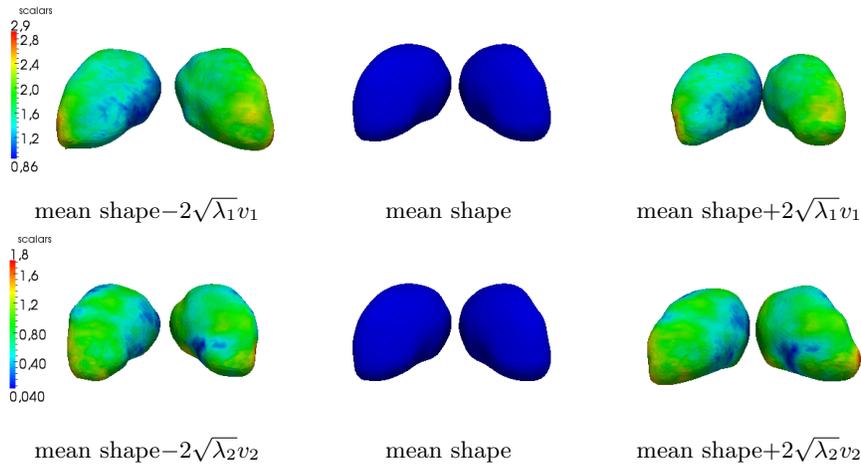


Fig. 4. Mean shape and first two modes of variation on left and right thalami. The colour maps the norm of the displacement (in mm) of each point along the first (top row) and second (bottom row) mode of variation ($\pm 2\sqrt{\lambda_i}v_i$) around the mean shape.

4 Discussion & Conclusion

In this paper, we presented the CLARCS library and its three building bricks, namely: (i) the rigid/affine/non-linear registration of two point sets, (ii) the computation of the mid-sagittal plane of one point set, and (iii) the computation of a mean point set from several point sets. We also gave some possible applications of CLARCS. Some other applications can be found in the full papers describing the methods implemented in CLARCS, for instance:

- the construction of a statistical shape model of the osseous labyrinth [6] and the caudate nuclei [10]
- the estimation of the mean pointwise brain asymmetry in male right-handed subjects and the comparison with situs inversus subjects [6] and with chimpanzees [13]
- the estimation of the mean pointwise facial asymmetry in males and females, and the comparison thereof [8]

With its new, robust and fast methods for surface processing, CLARCS opens interesting perspectives with medical applications. For instance, the non-linear registration of two point sets of about 200,000 points each runs in less than 5 minutes on a recent standard personal computer (3GHz).

The implementation of the framework of CLARCS is compatible with the well-known visualization toolkit VTK (<http://www.vtk.org>), thus it is possible to insert CLARCS specific methods into more general VTK pipelines.

We plan to distribute CLARCS as an open-source library, but this step will require some code refactoring, thus CLARCS will be first available as an external plugin of the future version of MedInria (<http://med.inria.fr>) that will be released in early September 2011.

5 Acknowledgements

We thank Guillaume Captier (CHRU Montpellier, Hôpital Lapeyronie, France) for providing the CT scan of the skull with plagiocephaly, and Gérard Subsol (ICAR Project-Team, LIRMM, CNRS, Montpellier, France) and Gilles Gesquière (LSIS, Université de Provence, Arles, France) for segmenting and reconstructing the outer surface of this skull. We thank Yann Heuzé and Joan T. Richtsmeier (Department of Anthropology, Pennsylvania State University, University Park, USA) for providing the CT scan of the skull with craniosynostosis (through the NIH grant #R01DEO18500), and for segmenting and reconstructing the outer surface of this skull. We thank Pierre Jannin (INSERM/INRIA VisAGeS Unit/Project U746, Rennes, France) for his input on deep brain stimulation through the ACouStiC project (<http://anr-acoustic.inria.fr>).

References

1. Benabid, A.L., Pollak, P., Gervason, C., Hoffmann, D., Gao, D.M., Hommel, M., Perret, J.E., de Rougemont, J.: Long-term suppression of tremor by chronic stimulation of the ventral intermediate thalamic nucleus. *Lancet* 337(8738), 403–406 (Feb 1991)
2. Besl, P.J., McKay, N.D.: A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.* 14(2), 239–256 (1992)
3. Captier, G., Leboucq, N., Bigorre, M., Canovas, F., Bonnel, F., Bonnaf, A., Montoya, P.: Plagiocephaly: morphometry of skull base asymmetry. *Surg Radiol Anat* 25(3-4), 226–233 (2003)
4. Chui, H., Rangarajan, A.: A new algorithm for non-rigid point matching. In: *CVPR*. pp. 2044–2051 (2000)
5. Chung, M.K., Dalton, K.M., Davidson, R.J.: Tensor-based cortical surface morphometry via weighted spherical harmonic representation. *IEEE Trans. Med. Imaging* 27(8), 1143–1151 (2008)
6. Combès, B., Fournier, M., Kennedy, D., Braga, J., Roberts, N., Prima, S.: EM-ICP strategies for joint mean shape and correspondences estimation: applications to statistical analysis of shape and of asymmetry. In: *8th IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'2011)*. pp. 1257–1263 (Mar 2011)
7. Combès, B., Hennessy, R., Waddington, J., Roberts, N., Prima, S.: Automatic symmetry plane estimation of bilateral objects in point clouds. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'2008)*. Anchorage, United States (2008)
8. Combès, B., Prima, S.: New algorithms to map asymmetries of 3D surfaces. In: *11th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'2008)*. vol. 11 (Pt 1), pp. 17–25. New-York, United States (2008)

9. Combès, B., Prima, S.: Setting priors and enforcing constraints on matches for nonlinear registration of meshes. In: 12th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'2009). vol. 12, pp. 175–183 (2009)
10. Combès, B., Prima, S.: An efficient EM-ICP algorithm for symmetric consistent non-linear registration of point sets. In: 13th International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'2010). vol. 13, pp. 594–601. Beijing, China (2010)
11. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models-their training and application. *Computer Vision and Image Understanding* pp. 38–59 (1995)
12. Follett, K.A., Weaver, F.M., Stern, M., Hur, K., Harris, C.L., Luo, P., Marks, W.J., Rothlind, J., Sagher, O., Moy, C., Pahwa, R., Burchiel, K., Hogarth, P., Lai, E.C., Duda, J.E., Holloway, K., Samii, A., Horn, S., Bronstein, J.M., Stoner, G., Starr, P.A., Simpson, R., Baltuch, G., Salles, A.D., Huang, G.D., Reda, D.J., Group, C.S.P.S.: Pallidal versus subthalamic deep-brain stimulation for parkinson's disease. *N Engl J Med* 362(22), 2077–2091 (Jun 2010)
13. Fournier, M., Combès, B., Roberts, N., Keller, S., S., Crow, J., T., Hopkins, D., W., Prima, S.: Surface-based method to evaluate global brain shape asymmetries in human and chimpanzee brains. In: 8th IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI'2011). pp. 310–316. Chicago, United States (Mar 2011)
14. Granger, S., Pennec, X.: Multi-scale EM-ICP: A fast and robust approach for surface registration. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) Seventh European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 2353, pp. 418–432. Springer, Copenhagen, Denmark (May 2002)
15. Jian, B., Vemuri, B.: A robust algorithm for point set registration using mixture of Gaussians. In: IEEE International Conference on Computer Vision. pp. 1246–1251. Beijing, China (Oct 2005)
16. Kinney, H.C., Thach, B.T.: The sudden infant death syndrome. *N Engl J Med* 361(8), 795–805 (Aug 2009)
17. Kordestani, R.K., Patel, S., Bard, D.E., Gurwitsch, R., Panchal, J.: Neurodevelopmental delays in children with deformational plagiocephaly. *Plast Reconstr Surg* 117(1), 207–18; discussion 219–20 (Jan 2006)
18. Moon, R.Y., Horne, R.S.C., Hauck, F.R.: Sudden infant death syndrome. *Lancet* 370(9598), 1578–1587 (Nov 2007)
19. Münch, D., Combès, B., Prima, S.: A modified ICP algorithm for normal-guided surface registration. In: Medical Imaging 2010: Image Processing. p. 76231A. San Diego, United States (Mar 2010)
20. Myronenko, A., Song, X.B.: Point set registration: Coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.* 32(12), 2262–2275 (2010)
21. Rangarajan, A., Chui, H., Mjolsness, E., Pappu, S., Davachi, L., Goldman-Rakic, P., Duncan, J.: A robust point-matching algorithm for autoradiograph alignment. *Med Image Anal* 1(4), 379–398 (Sep 1997)
22. Shen, L., Farid, H., McPeck, M.A.: Modeling three-dimensional morphological structures using spherical harmonics. *Evolution* 63(4), 1003–1016 (Apr 2009)
23. Tamburrini, G., Caldarelli, M., Massimi, L., Santini, P., Rocco, C.D.: Intracranial pressure monitoring in children with single suture and complex craniosynostosis: a review. *Childs Nerv Syst* 21(10), 913–921 (Oct 2005)

Automatic Statistical Shape Analysis of Local Cerebral Asymmetry in 3D T1-weighted magnetic resonance images

Antonietta Pepe¹, Lu Zhao², Jussi Tohka¹, Juha Koikkalainen³, Jarmo Hietala^{4,5}, and Ulla Ruotsalainen¹

¹ Tampere University of Technology, P.O. Box 553, FIN-33101 Tampere, Finland

² McConnell Brain Imaging Center, McGill University, 3801 University Street, Montréal, Québec, Canada

³ VTT Technical Research Centre of Finland, P.O. Box 1300, FIN-33101 Tampere, Finland

⁴ Turku PET Centre, P.O. Box 52, FIN-20521 Turku, Finland

⁵ Department of Psychiatry, University of Turku, FIN-20700 Turku, Finland

Abstract. The study of the structural brain asymmetries can prove useful for the understanding of the functional brain lateralizations, and to examine the progression and assist the early diagnosis of various neuropsychiatric disorders. This work introduces a novel automatic method for the local (vertex-level) statistical shape analysis of gross cerebral hemispheric surface asymmetries from three dimensional (3D) T1-weighted Magnetic Resonance Images (MRI). In the proposed method, after the brain extraction, the cerebral hemispheric volume images were segmented and spatially normalized to a common stereotaxic space. Smooth mesh representations of the cerebral hemispheres were extracted from the spatially normalized cerebral hemispheric masks. Customized measures of vertex-level (local) inter-hemispheric shape asymmetry were finally evaluated at 2562 corresponding surface vertices and tested for statistical significance. The developed method was tested on scans obtained from a small sample of first-episode neuroleptic-naïve female (7) and male (11) schizophrenic subjects.

1 INTRODUCTION

The cerebral hemispheres of a healthy human brain are asymmetrical in width, volume, weight, density, and shape [1]. The evidence of disturbed structural and functional asymmetry in schizophrenia [2,3] as compared to the asymmetry pattern observed in normal population has led to the interpretation of schizophrenia as an abnormal hemispheric lateralization [4,5]. The study of neuroanatomic asymmetries could therefore prove useful for the understanding of the functional hemispheric specializations [1], and for the tracking and early detection of schizophrenia and other psychiatric disorders which are believed to have a neurodevelopmental etiology [6].

Even though the whole brain asymmetry has been widely investigated in classical MRI-based literature by conventional volumetric methods (e.g. [7]), *in vivo* studies performing an accurate, automatic and robust assessment of the structural shape asymmetries of the entire brain are far fewer. Among the existing automatic methods for shape analysis, the Surface Based Morphometry (SBM) analyzes the shape of structural boundaries between different tissue types. In the SBM based studies of the anatomical asymmetries across the whole brain, a detailed cortical surface representation is created, typically through meshing. Based on the geometrical properties of such a representation, one or multiple indices of brain morphology (e.g. the shape asymmetry [8] and the cortical thickness asymmetry [9] of the gyral and sulcal patterns) are computed in the matching homologous locations between the hemispheres and analyzed.

In this study, we propose a novel automatic algorithm named as *Statistical Shape Analysis of Local Asymmetry* (SSALA) for *in vivo* local (vertex-level) statistical shape analysis of the cerebral hemispheric (CH) surface asymmetry from 3D T1-weighted MRI. The CH surfaces follow smoothly the mean features of the hemispheric outlines and do not model the pattern of sulcal and gyral folds in order to capture the gross CH shape. Compared to the SBM based studies of the cortical surface asymmetry, SSALA focuses on coarser level asymmetries exhibited in smooth surfaces approximating the outer shapes of CHs.

In this work, we applied SSALA to study and visualize the CH shape asymmetry in a relatively small sample of the neuroleptic-naïve schizophrenia patients. The use of antipsychotic medications has been shown to affect the brain anatomy [10] and it could therefore act as a confounding element in the study of the brain shape. However, the non-medicated subjects diagnosed with schizophrenia are difficult to recruit [11] and there are only few published studies performing a fully automatic analysis of the CH surface shape asymmetry in the neuroleptic-naïve schizophrenic subjects.

2 Materials and Methods

In SSALA, the left and right CHs were segmented from the 3D MR head volume images and spatially normalized into a common template space. The segmented and spatially normalized right CH volume images were then reflected across the mid-sagittal plane of the ICBM152 space to match the corresponding left CH volume images. Mesh representations of the left and reflected right CH surfaces were extracted from the corresponding normalized volume images. The local inter-hemispheric shape asymmetry was assessed at each surface vertex and tested for statistical significance within groups by non-parametric hypothesis testing. Finally, the permutation testing on extreme statistics was performed to determine thresholds that controlled for the multiple comparisons. The corrected p-values were visualized along with uncorrected effect sizes on a 3D average CH surface. The SSALA method is fully automatic.

2.1 Subjects and Image Acquisition

In this study, clinical 3D T1-weighted MRI head scans of 18 first-episode neuroleptic-naïve schizophrenic subjects, 7 females (*SF group*) and 11 males (*SM group*), and of 19 healthy controls, 7 females (*HF group*) and 12 males (*HM group*) were used [20]. All subjects were right-handed according to the Edinburgh inventory [12]. The MRI scans (voxel size $1.5 \times 1.5 \times 1.0 \text{ mm}^3$, image dimensions from $256 \times 256 \times 150$ to $256 \times 256 \times 170$ voxels) were acquired in the Turku University Central Hospital (Turku, Finland) using a 1.5 T Siemens Magnetom (Erlangen, Germany) with a three-dimensional magnetization-prepared rapid acquisition gradient echo sequence (TR 10 ms, TE 4 ms, 1 acquisition, matrix 256×92 , slice sagittal thickness 1.5 mm and 0 interslice gap).

2.2 Image Preprocessing and Brain Extraction

The skull, scalp and other extraneous tissues (except the cerebrospinal fluid) were initially removed from the original MR images using the Brain Surface Extractor routine of the BrainSuite package [13] (<http://www.loni.ucla.edu/Software/BrainSuite>). The skull stripped volumes were corrected for the intensity non-uniformity artifact using the Bias Field Corrector routine [13], which is also implemented in the BrainSuite package. The skull-stripped and intensity non-uniformity corrected brain volumes were then processed with an automatic partial volume estimation technique to estimate the relative amounts of the white matter (WM), grey matter (GM) and cerebrospinal fluid (CSF) in each brain voxel [14]. This technique also produced the partial volume voxel classification for the CSF/background, CSF/GM, and GM/WM mixtures. We used the information on the partial volume mixture to remove from the skull-stripped volume all the voxels belonging to the pure CSF, CSF/background and the partial volume voxels containing more than 30% of CSF tissues.

The employed partial volume estimation technique, which is based on the trimmed minimum covariance determinant parameter estimation and the Markov random fields based tissue classification, produces more accurate results than the partial volume estimation technique implemented in the BrainSuite package [13]. The correction for the intensity non-uniformity artifact and the accurate partial volume estimation was here important for the accuracy of the brain extraction and hence of the following cerebral hemisphere segmentation.

2.3 Cerebral Hemisphere Segmentation and Spatial Normalization

Once that the non-brain tissues were eliminated, each MR head image was decomposed in the native space into the cerebrum, cerebellum, and the brain stem compartments based on the tissue fraction information by the algorithm described in [15]. An extended partial differential equations based shape bottle-necks algorithm [15] was then used to cut the inter-hemispheric connections of the extracted cerebrum. As a result, the cerebrum compartment was segmented into the left and right CHs with no need for image registration.

As shown in [15], our automatic cerebral hemisphere segmentation algorithm produces more accurate hemisphere segmentation if compared to the traditional shape bottlenecks [16], linear and non-linear registration based brain hemisphere segmentation methods. The accurate CH segmentation was a crucial prerequisite for the veracity of the later CH shape analysis.

The segmented CH volume images were spatially normalized into the stereotaxic ICBM152 space (voxel size $2 \times 2 \times 2$ mm³, dimension $91 \times 109 \times 91$ voxels) using the SPM5 package [17] (<http://www.fil.ion.ucl.ac.uk/spm/software/spm5/>).

The human brain is characterized by complex convolutions whose morphology varies greatly among subjects and between hemispheres [18]. The affine registration can reduce the effects of such normal inter-subject anatomical variability. For quantifying the CH shape asymmetry, the role of the stereotaxic registration in SSALA was two-fold: (1) It reoriented the segmented CH volume images into a common space to allow for the inter-subject comparisons; (2) It compensated for gross (up to affine transformations) anatomical inter-subjects biological variations.

The spatially normalized right CH volume images were then reflected with respect to the mid-sagittal plane of the stereotaxic space to match the left CH volume images. The reflection of the CH volume images was needed for the achievement of the point correspondence.

2.4 Surface Extraction

The surface of each segmented and spatially normalized CH was extracted using a global optimization based deformable surface model automatically initialized with spherical surface mesh [19]. The employed deformable surface model algorithm has been previously demonstrated to be robust against local minima of the cost function, which is a problem often associated with the deformable models [19]. Each hemispheric surface was represented with a triangulated mesh consisting of 5120 triangles (2562 vertices). The geometry of the initial surface meshes was chosen such to extract the salient morphological features and not modeling the sulcal and gyral patterning of the hemispheres. The smoothness of the extracted surface meshes reduced part of the normal inter-subject and inter-hemispheric variability in brain morphology which remained from the spatial normalization. The point correspondence, assuring that the pairs of corresponding vertices represent a same biological location, is fundamental in SSALA for meaningful comparisons between CH surfaces. The point correspondence between the triangular mesh representations of the CH shapes was naturally achieved in SSALA due to: (1) Mirroring of each right CH volume image with respect to the mid-sagittal plane of the ICBM152 space; (2) In the surface deformation algorithm, each surface vertex is constrained to move approximately toward the centre of mass of the initial surface mesh. When the surfaces to be extracted have approximately the same shape and size (as guaranteed by the stereotaxic registration) and the initial surface mesh is the same for all the surfaces to be

extracted, the deformation algorithm produces an approximate point correspondence for all hemispheric surfaces of each subject. The point correspondence was also empirically confirmed.

2.5 Statistical shape analysis

For the purpose of shape analysis, a CH reference shape was obtained as the mean shape of the all data set consisting of 37 MRI head scans. After that, a vertex-level morphometric measure of shape distance between the left CH surface in the subject j at vertex i ($t_{i,j}^L \in \mathbb{R}^3$) and the reference CH surface at every vertex i ($c_i \in \mathbb{R}^3$) was defined as follows

$$d_{i,j}^L = (t_{i,j}^L - c_i) \cdot n_i, \quad (1)$$

where $i = 1, \dots, 2562$ denotes the index of vertices in the triangulated mesh representations of the CH, $t_{i,j}^L$ denotes the 3D coordinates of the i -th vertex of the left CH surface in the subject j , and where n_i denotes the unit normal of the reference surface at c_i . The vertex-level morphological distances were defined similarly for the reflected right CHs

$$d_{i,j}^{rR} = (t_{i,j}^{rR} - c_i) \cdot n_i. \quad (2)$$

The measures $d_{i,j}^L$ and $d_{i,j}^{rR}$ were evaluated at every vertex and among all subjects of each group to study the L>R ($d_{i,j}^L > d_{i,j}^{rR}$) and the R>L ($d_{i,j}^{rR} > d_{i,j}^L$) asymmetry.

The Gaussianity of the set $d_i = \{d_{i,j}^L, d_{i,j}^{rR} : j = 1, \dots, N_G\}$, where N_G is the number of subjects within a group, was investigated separately at every vertex i using the *chi-squared* test and rejected at the significance level 0.05. Since the set d_i was observed to be non normally distributed in each group, the Wilcoxon signed Rank test was used to evaluate the null hypotheses of $d_{i,j}^L, j = 1, \dots, N_G$ and $d_{i,j}^{rR}, j = 1, \dots, N_G$ having symmetric distributions with equal medians. Finally, the effect sizes of the vertex level asymmetry were evaluated in each group as

$$S_i = r_i / (2N_G) \quad (3)$$

where r_i denotes the signed rank statistic obtained using the Wilcoxon signed Rank test at vertex i . The effect sizes S_i were interpreted as small, medium or large effect sizes if $0.4 \leq S_i < 0.6$, $0.6 \leq S_i < 0.8$ or $S_i \geq 0.8$ respectively [21].

2.6 Multiple comparisons correction

As 2562 tests were performed, one at each surface vertex, the results obtained from the Wilcoxon test had to be corrected for the multiple comparisons. Here the studied groups were small due to difficulties in recruiting neuroleptic-naïve subjects diagnosed with schizophrenia, and the test statistics were expected to be strongly spatially correlated. The multiple comparisons correction was thus

performed via permutation tests on maximal rank statistics which rely on minimal assumptions on the underlying distributions and automatically adapt to the spatial correlations in the data [23]. Under the null hypotheses of no vertex-level asymmetries, the $d_{i,j}^L$ and $d_{i,j}^{rR}$ measures (Eq. (1) and (2)) were fully permuted at each matched vertex i by relabeling the *left* and *reflected right* tags [24]. A two-tailed non-parametric statistic of the maximal (over all vertices) mean of rank difference between $d_{i,j}^L$ and $d_{i,j}^{rR}$ was evaluated for each possible relabeling. The false positive corrected thresholds for L>R and R>L asymmetries were computed by selecting the 95th percentile of the maximum statistics.

3 RESULTS

For the *SM* group, patterns of statistically significant L>R CH shape asymmetries ($p < 0.05$, corrected) were observed in the parietal lobe, in the proximity of the superior parietal gyrus, superior parietal lobule, and inferior parietal lobule; in the occipital lobe, in the proximity of the middle occipital gyrus; and in the medial temporal lobe, in the proximity of the parahippocampal gyrus and paraventricular nucleus (see Fig. 1). None of the uncorrected p-values in the *SF* group and none of the vertex-level R>L asymmetries in the *SM* group survived the correction of the multiple comparisons errors. To assist the interpretation of the maps of statistical significance, maps of effect sizes were also obtained and presented in Fig. 2. Fig. 2 demonstrates that the patterns of the effect sizes of the L>R and R>L asymmetry are remarkably different between *SF* and *SM* groups.

4 CONCLUSIONS

The SSALA method is fully automatic and it is therefore suitable for large database analyses. Due to the more gross nature of the extracted CH surfaces, SSALA is conceptually and methodologically easier than many other SBM based methods. Furthermore, the SSALA method can be used for relatively small sample sizes, where the high inter-subject variability between subjects in position, extent and morphology of the sulci and gyral foldings could impede a meaningful SBM based comparison of cortical convolution patterns.

The shape analysis performed by SSALA revealed various patterns of statistically significant asymmetries which survived the correction for the multiple comparisons problem. Interestingly, patterns of asymmetry were found in the male schizophrenic subjects, especially in the occipital, frontal, and temporal lobes. However, the results of our work should be interpreted as a preliminary study that needs to be tested further with larger sample sizes.

ACKNOWLEDGMENT This work was supported by the Academy of Finland (grant 1302759, application number 129657, Finnish Programme for Centers of Excellence in Research 2006 – 2011).

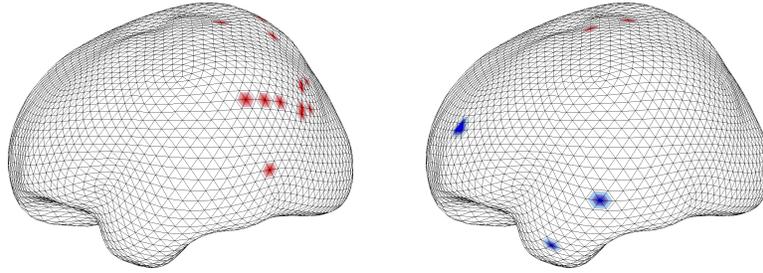


Fig. 1. The significance maps ($p < 0.05$, corrected for the multiple comparisons) of the L>R (red colored areas) and R>L (blue colored areas) CH shape asymmetry in the *SM* group (left) depicted in lateral view. For comparison, the significance maps of the shape asymmetry in the *HM* group (right) is also shown. None of the uncorrected p -values in the *SF* group survived the correction of the multiple comparisons errors.

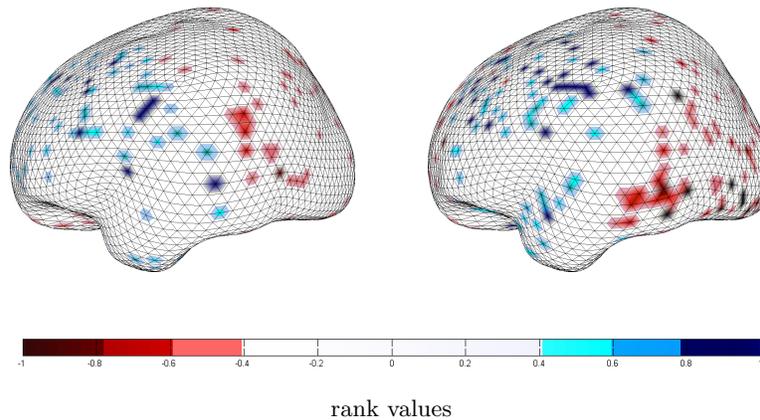


Fig. 2. The effect size maps of the L>R (red color scale) and R>L (blue color scale) CH asymmetry in *SF* (left) and *SM* (right) groups depicted in lateral views.

References

1. Toga, A. W., Thompson, P. M., 2003. Mapping brain asymmetry. *Nat. Rev. Neurosci.* 4 (1), 37–48.
2. Broca, P., 1861. Perte de la parole. Ramollissement chronique et destruction partielle du lobe antérieur gauche du cerveau. *Bull. Soc. Anthropol.* 2: 235–238.
3. Barta, P.E., et al., 1997. Planum temporale asymmetry reversal in schizophrenia: replication and relationship to grey matter abnormalities. *Am. J. Psychiat.* 154 (5), 661–667.
4. Flor-Henry P. *Cerebral Basis of Psychopathology*. Boston: John Wright; 1983.
5. Angrilli, A., et al., L., 2009. Schizophrenia as Failure of Left Hemispheric Dominance for the Phonological Component of Language. *PLoS ONE*, 4(2): e4507.
6. Crow, T., 1990. Temporal lobe asymmetries as the key to the etiology of schizophrenia. *Schizophr. Bull.* 16 (3), 433–443.

7. Bilder, R., et al., 1994. Absence of regional hemispheric volume asymmetries in first-episode schizophrenia. *Am. J. Psychiat.* 151 (10), 1437–1447.
8. Narr, K. L., et al., 2007. Asymmetries of cortical shape: Effects of handedness, sex and schizophrenia. *NeuroImage* 34 (3), 939–948.
9. Hamilton, L., et al., 2007. Asymmetries of cortical thickness: effects of handedness, sex, and schizophrenia. *Neuroreport, Clinical Medicine* 18, 1427–1431.
10. Lieberman, J., et al., 1993. Time course and biological correlates of treatment response in first episode schizophrenia. *Arch. Gen. Psychiat.* 50, 369–376.
11. Steen, R. G., et al., 2006. Brain volume in first-episode schizophrenia: Systematic review and meta-analysis of magnetic resonance imaging studies. *Brit. J. of Psychiat.* 188, 510–518.
12. Oldfield, R., 1971. The assessment and analysis of handedness: The edinburgh inventory. *Neuropsychologia* 9, 97–113.
13. Shattuck, D. W., et al. 2001. Magnetic resonance image tissue classification using a partial volume model. *Neuroimage* 13 (5), 856–876.
14. Tohka, J., et al., 2004. Fast and robust parameter estimation for statistical partial volume models in brain MRI. *NeuroImage* 23 (1), 84–97.
15. Zhao, L., et al., 2010. Automatic cerebral and cerebellar hemisphere segmentation in 3D MRI: Adaptive disconnection algorithm. *Med. Image Anal.* 14(3), 360–372.
16. Mangin, J.-F., et al., 2004. A framework to study the cortical folding patterns. *NeuroImage* 23, 129–138.
17. Ashburner, J., Friston, K., 1999. Nonlinear spatial normalization using basis functions. *Human Brain Mapping* 7(4), 254–266.
18. Tahmasebi, A., 2010. Quantification of Inter-subject Variability in Human Brain and Its Impact on Analysis of fMRI Data. Thesis (Ph.D, Computing)–Queen’s University, 2010-04-29 07:07:55.77.
19. Tohka, J., and Mykkänen, J. M., 2004. Deformable mesh for automated surface extraction from noisy images. *Int. J. Image Graphics* 4 (3), 405–432.
20. Laakso, M. P., et al., 2001. A morphometric MRI study of the hippocampus in first-episode, neuroleptic-naive schizophrenia. *Schizophrenia Research* 50 (1-2), 3–7.
21. Vargha, A., and Delaney, H. D., 2000. A Critique and Improvement of the CL Common Language Effect Size Statistics of McGraw and Wong. *J. Educ. Behav. Strat.* 25, 101–132.
22. Weerahand, S., 2003. Exact Statistical Methods for Data Analysis. Chapter 4.3. Springer Series in Statistics.
23. Pantazis, D., et al., 2004. Statistical surface-based morphometry using a non-parametric approach. *IEEE International Symposium on Biomedical Imaging*, 2, 1283–1286.
24. Nichols, T. E., and Holmes, A. P., 2002. Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Hum. Brain Mapp.* 15, 1–25.

High Quality Remeshing of Closed Genus-0 Surfaces

Anton Semechko, Robert Dony and Michele Oliver,

University of Guelph, Guelph ON N1G 2W1, Canada
{asemechk, rdony, moliver}@uoguelph.ca

Abstract. In this paper, we present a new technique for resampling closed genus-0 surfaces. By combining the properties of conformal maps with the physical principles governing the dynamics of charged particle systems, we develop a method to adaptively sample the parameter domain in a manner that produces nearly uniform sampling of the original surface. Initializing the system with a small number of particles we obtain a base mesh that provides the necessary scaffold to recover the surfaces of increasing complexity and high quality. We implement the proposed algorithm in MATLAB and evaluate its performance on 239 carpal bones as well a number of publically available models.

Keywords: remeshing, spherical parameterization, conformal map, particle system

1 Introduction

Triangle meshes are one of the simplest and perhaps the most widely used representation of surfaces. They are central to various methods in computer graphics, computer assisted design, finite element (FE) simulations and related areas. Meshes are generally subject to two common considerations of quality and complexity. For instance, in applications where meshes are used to discretize boundary value problems, to ensure good convergence and accuracy of the solution, aspect ratios of all triangles faces should fall within a relatively narrow range [1]. Additionally, in applications where efficiency is important, meshes representative of the underlying geometry, reduced complexity and potentially good quality may be desirable. One of the research interests of our group involves creation of FE and statistical shape models of the human wrist joint. Consequently the issues of mesh quality and complexity are equally important to us. In this paper we present a novel, modular algorithm that can be used to address either one of the above issues separately or simultaneously. The proposed method is made possible by combining the equations governing the dynamics of a system of charged particles with the theory of isomorphic maps from differential geometry. We demonstrate the performance of our algorithm on a selected set of anatomical models that can be downloaded from publically available databases [22-24].

2 Related Work

Surface remeshing refers to the procedure of resampling the input mesh to create a mesh that has overall better quality¹ than the original piecewise linear surface while retaining the same topology and as many geometrically salient features as possible. A literature survey of remeshing techniques that aim to optimize mesh quality reveals an enormous amount of research done to address this problem and we refer interested readers to [2] for a more detailed discussion on this subject. In the remainder of this section we provide a brief overview of remeshing techniques most similar to ours, namely the methods that rely on parameterization or physics based principles to obtain high quality surface meshes².

Parameterization of a surface mesh is the problem of establishing a bijective mapping between the original piecewise linear surface and a continuous parameter domain which is a primitive topological analog of the surface. The methods described in this paper are tailored specifically for closed genus-0 surfaces, which are characteristic of various everyday objects as well as many organs in the human body (brain, kidneys, liver, lungs, etc.) including the bones in the wrist. For such surfaces, the unit sphere is the most natural parameter domain [3,8]. Planar parameterizations of closed genus-0 surfaces, however, are also common and may be advantageous in certain situations (e.g. texture mapping) but are less natural from a global point of view, as they tend to produce higher distortions. Remeshing algorithms that rely on parameterizations fall into two categories. One approach consists of tessellating the parameter domain and then mapping the result back to the original surface [4-9]. This operation is very intuitive but requires maps that minimize global area and angle distortions or on other hand some type of mechanism that can correct these distortions. An alternative use of parameterization is to provide surface constraints while locally adaptive optimization procedures modify the mesh and steer it toward some ideal configuration [10-12]. In our implementation we rely on the former approach to resample the surface.

A completely different class of methods used to produce high quality meshes is based on the observation that minimization of the potential energy of physical systems, composed of charged particles or a network of nodes interconnected by springs, tends to produce compact descriptions of the underlying object geometry. In these ground energy states particle/node conformations have ideal (hexagonal) or close to ideal packing [16], producing regular Voronoi cells that in turn generate regular triangular tessellation. For example, uniform sampling of a sphere using a system of charged particles is a well known approach with many practical applications in computer graphics, physical chemistry and numerical analysis [13,14]. Not surprisingly, meshing/remeshing methods described in [15,16] rely precisely on the interaction of charged particle systems to produce high quality triangulations of arbitrary surfaces suitable for FE analysis. By far the most general method in this family is presented in

¹ Mesh quality can refer to any one or a combination of criteria that include the sampling, grading, regularity, size and shape of elements [2].

² Surface mesh serves as input to generate volumetric mesh. Thus by optimizing the quality of the former it is possible to control (to a reasonable extent) the quality of the latter.

[17] and obtains optimal triangular by solving a set of 2nd order (Euler-Lagrange) ODEs:

$$m_i \frac{d^2 \mathbf{x}_i}{dt^2} + b \frac{d\mathbf{x}_i}{dt} = \vec{\mathbf{F}}(\mathbf{x}_i) \quad (1)$$

where \mathbf{x}_i is the position of the i -th particle, m_i is the mass of the particle, b is damping parameter and $\vec{\mathbf{F}}(\mathbf{x}_i)$ is the net force exerted on the particle due to interaction with its neighbours and can be defined using arbitrary physical laws (e.g. Hook’s law, Coulomb’s law, etc.). These equations can be solved using any finite difference scheme and according to [17] “result in evenly distributed points that very closely follow the surface features.” Moreover, $\vec{\mathbf{F}}(\mathbf{x}_i)$ can be modified to include curvature information to obtain adaptive sampling of the surface. The main disadvantage of this method and those presented in [15,16] is the need to consistently enforce surface constraints which is a computationally expensive operation.

The algorithm we describe herein is a hybrid of the above approaches. It is most similar to the remeshing procedure presented in [9] that relies on planar parameterization and sampling density maps. In contrast to [9], however, our technique is based on fundamentally different principles to sample the parameter domain and uses spherical parameterizations (instead of planar), which do not require partitioning of the surface. Overall, our method is conceptually quite simple.

In [18] Gu *et al.* make an important point that for genus-0 surfaces spherical embedding based on minimization of harmonic energy is conformal. The characteristic property of conformal maps is that they preserve angles [18,19]. For closed genus-0 surfaces that are geometrically dissimilar to a sphere, conformal maps produce significant area distortions. Defining area distortion simply as the ratio of areas of the spherical triangle and its planar counterpart, we sample the sphere³ (i.e. parameter domain) adaptively in a manner that promotes dense aggregation of points in regions of low distortion and sparse distribution in zones of high distortion. This type of adaptive sampling of the sphere is equivalent to approximately uniform sampling of the original surface. We control the complexity of the mesh by initializing the system with a small number of particles to produce a base mesh in the parameter domain. Surfaces of increasing complexity are recovered by subdividing the base mesh with any standard subdivision scheme (e.g. triangular quadrisection) and then mapping the result back to the original surface. Subdivision of equilateral triangles produces roughly self-similar triangles at finer scales while conformal maps ensure that the triangles retain their aspect ratios. In the next sections we provide description of our method as well as some practical benefits that stem from its multiresolution construction.

³ Due to perfect symmetry of the sphere the surface constraints can be easily enforced.

3 Methods

Our method consists of four consecutive stages (see Figure 1) that include: (1) calculation of a conformal map of the input mesh, (2) calculation of the distortion field induced by the mapping from the previous step, (3) adaptive sampling of the parameter domain guided by the distortion field from step 2, (4) subdivision and minor regularization. The specific details of these stages are discussed in the following subsections.

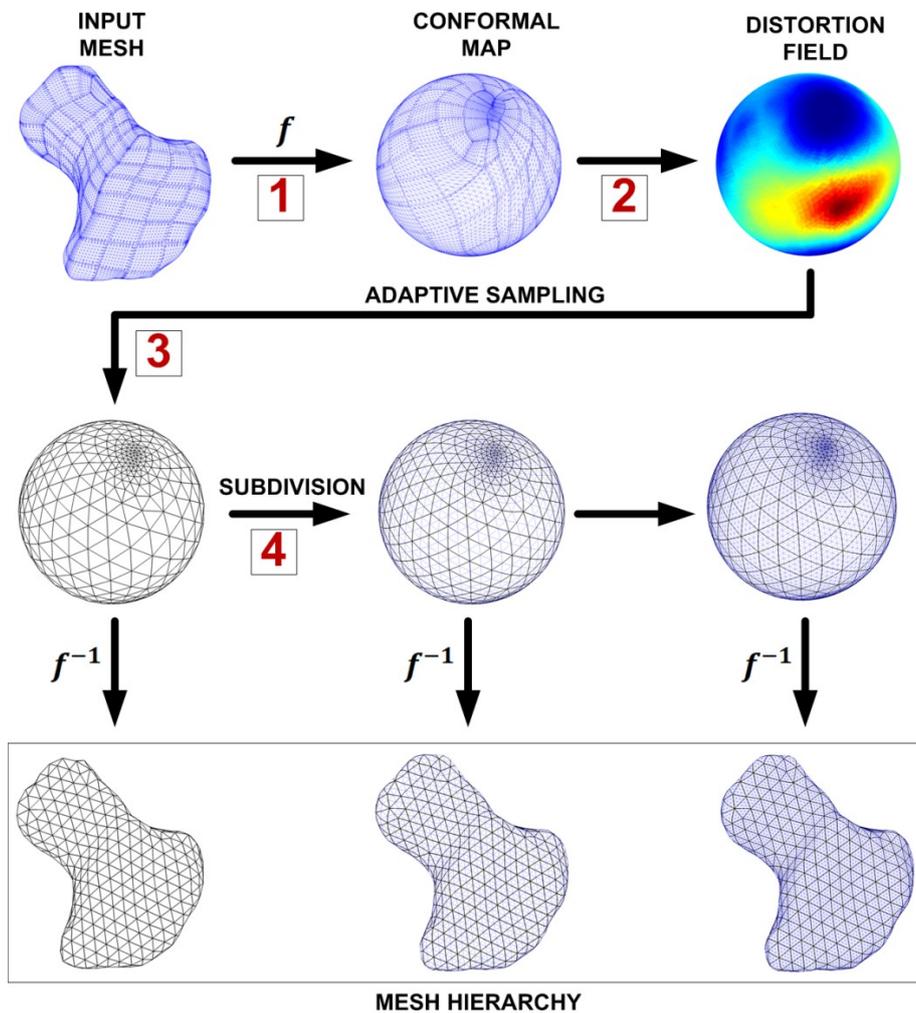


Fig.1. Four main stages of the resampling procedure: (1) conformal parameterization, (2) calculation of the distortion field at the vertices of the input mesh, (3) adaptive sampling of the parameter domain using a system of charged particles, (4) subdivision and optional regularization.

3.1 Conformal Parameterization

Let $M = [V, E]$ be a triangular surface mesh representing the underlying continuous surface S , where $V = \{\mathbf{v}_i \in S \mid i = 1, \dots, N\}$ and E are sets of vertices and edges, respectively. A parameterization of M is defined as an isomorphism $\mathbf{f} : S \rightarrow S^2$, where $S^2 = \{\mathbf{x}_i \in \mathbb{R}^3 \mid \|\mathbf{x}_i\| = 1\}$. We are most interested in conformal maps that by definition minimize the distortion in angles, a property crucial to our method. In [18] Gu *et al.* make an important point that for genus-0 surfaces, harmonic maps are also conformal. A map is harmonic if it is a stationary function of the Dirichlet energy functional [19,20]. For piecewise linear surfaces, the harmonic energy is a quadratic function:

$$\Phi(\mathbf{f}, M) = \frac{1}{2} \sum_{[\mathbf{v}_i, \mathbf{v}_j] \in E} k_{ij} \|\mathbf{f}(\mathbf{v}_i) - \mathbf{f}(\mathbf{v}_j)\|^2 \quad (2)$$

where k_{ij} is the edge weight, derived from the opposing angles (α and β) of the two faces sharing the edge $[\mathbf{v}_i, \mathbf{v}_j]$ as shown below:

$$k_{ij} = [\cot(\alpha) + \cot(\beta)]/2 \quad (3)$$

We optimize $\Phi(\mathbf{f}, M)$ using the method presented in [18]. According to this approach \mathbf{f} is initialized to the Tutte map⁴ of M and then deformed along the negative gradient of $\Phi(\mathbf{f}, M)$ in the tangent space of the sphere followed by periodic enforcement of the zero-mass center and surface constrains (see [18] for more details). In instances where the mesh has a significant number of negative harmonic weights (this occurs if $\alpha + \beta > \pi$) then the map may not be one-to-one. In our experiments we have processed over 240 surfaces and have not observed this effect even for meshes that had as many as 20% negative edge weights.

3.2 Distortion Field

In our algorithm we use conformal maps to establish the connection between the original surface and the parameter domain. Conformal maps do not preserve areas in general and induce large area deformation for surfaces that are very dissimilar to a sphere. To maintain uniform sampling of the original surface, mapping induced distortions have to be taken into consideration. For every face in the mesh, we define area distortion (AD) as the ratio of the normalized areas of the spherical triangle and its planar image. Similarly, for every edge in the mesh, we define length distortion (LD) as the ratio of edge length in the parameter domain and its length on the original surface. Next, for every vertex we define AD_v (LD_v) as the minimum AD (LD) of the faces (edges) incident on that vertex. The net distortion field (DF) used to guide the distribution of points on the sphere is simply the product of LD_v and AD_v . Finally, we use spherical barycentric coordinates to establish the value of the DF at an arbitrary point on the unit sphere.

⁴ Tutte map is computed by minimizing (2) with $k_{ij} = 1$.

3.3 Adaptive Sampling of a Unit Sphere

A system of charged particles provides an efficient means of sampling a surface. We use a slightly generalized version of the electrostatic potential energy, Reisz s-energy (4), to describe dynamics of a particle system with unequal positive charges:

$$R_s = \sum_i^N \sum_{j \neq i}^N \frac{q_i q_j}{d_{ij}^s} \quad (4)$$

where N is the total number of particles q is the particle charge (defined as the product of area and length distortion), d is the geodesic distance between the particles and $s \in (0,1]$ is a free parameter that modulates the strength of particle interactions.

For smooth surfaces such as the bones of the wrist, we observed that the quality of the resulting triangulation is almost independent of s . For surfaces such as the one shown in Figure 4, conformal mapping tends to produce large distortions in isolated regions. In those instances setting $s < 1$ ensures denser packing of particles in those regions. In all our experiments we use $s = 0.1$.

To adaptively distribute the charged particles over the surface of the sphere, we assign particle charges based on particle positions. Let $q_i = q(\mathbf{x}_i)$, where $q : \mathbb{R}^3 \rightarrow \mathbb{R}^+$. For an arbitrary position on the sphere $q(\mathbf{x}_i)$ can be approximated using linear interpolation:

$$q(\mathbf{x}_i) = uq_1^i + vq_2^i + wq_3^i \quad (5)$$

where q_k^i are the charges defined on the mesh vertices and $\mathbf{x}_i = (x_{i1}, x_{i2}, x_{i3})^T$ is the position vector of the i -th particle on the surface of the sphere. Define $\{\mathbf{x}_k^i \mid k = 1,2,3\}$ as the vertex coordinates of the spherical triangle containing \mathbf{x}_i , then u, v and w are the spherical barycentric coordinates of \mathbf{x}_i :

$$\mathbf{A}_i [u \quad v \quad w]^T = \mathbf{x}_i, \quad u + v + w \geq 1, u, v, w \geq 0 \quad (6)$$

where $\mathbf{A}_i = [\mathbf{x}_1^i \quad \mathbf{x}_2^i \quad \mathbf{x}_3^i]$. With this in mind, the derivative of R_s with respect to \mathbf{x}_i can be approximated as:

$$\frac{\partial R_s}{\partial \mathbf{x}_i} \approx 2 \sum_{j \neq i}^{\Omega_i} \frac{q_j}{d_{ij}^s} \left[\text{Adj}(\mathbf{A}_i) [q_1^i \quad q_2^i \quad q_3^i]^T + \frac{q_i s}{d_{ij}} \frac{\mathbf{x}_i \cdot \mathbf{x}_j}{[1 - (\mathbf{x}_i \cdot \mathbf{x}_j)^2]^2} \mathbf{x}_j \right] \quad (7)$$

$$\Omega_i = \{j \mid \mathbf{x}_i \cdot \mathbf{x}_j \geq \cos \theta_{max}\} \quad (8)$$

$\text{Adj}(\mathbf{A})$ is the adjoint matrix of \mathbf{A} and θ_{max} is the angle separation threshold measured in radians (e.g. $\pi/4$). It should be noted that the vector $\text{Adj}(\mathbf{A}_i) [q_1^i \quad q_2^i \quad q_3^i]^T$ does not change in the course of the optimization as it is intrinsic to the conformal parameterization of the input mesh and the distortion field defined at its vertices. In order to speed up the optimization process we precompute and save these vectors.

We initialize the system with N_o particles⁵ randomly distributed across the surface of the unit sphere and iteratively update their positions in a sequential manner using the tangential component of the gradient (9). To keep track of particle positions relative to the mesh faces, we use a ray-triangle intersection algorithm described in [21]. In our experiments using step size of 10^{-3} , optimization converged in 500 to 2000 iterations.

$$\left(\frac{\partial R_s}{\partial \mathbf{x}_i}\right)_{\parallel} = \frac{\partial R_s}{\partial \mathbf{x}_i} - \left(\mathbf{x}_i \cdot \frac{\partial R_s}{\partial \mathbf{x}_i}\right) \mathbf{x}_i \quad (9)$$

3.4 Subdivision and Regularization

Once the particles have been fixed on the unit sphere, their positions determine the vertices of the base mesh⁶ in the parameter domain. Surfaces of increasing complexity can be recovered by iteratively subdividing this mesh with any standard subdivision scheme and mapping the result back to original surface⁷. Triangular quadrisection is the subdivision method used in this study. Figure 2 illustrates one iteration of this procedure. The method works because triangular quadrisection of equilateral triangles produces (approximately) self-similar triangles at finer scales while conformal mapping ensures that the triangles retain their aspect ratios. Every subdivision iteration roughly quadruples the number of vertices thus resulting in exponential increase in mesh complexity.

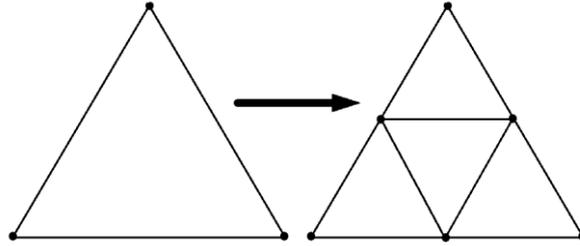


Fig.2. One iteration of triangular quadrisection. New vertices are inserted at edge midpoints.

Regularization is only necessary if the base triangulation is subdivided more than once, because the vertex position of the subdivided mesh do not accurately take into account the distortions induced by conformal mapping. To account for this discrepancy, prior to inverse mapping, the mesh is relaxed with a few iterations of gradient-based optimization of the Tutte energy (same procedure as conformal mapping but with all harmonic weights set to one). As can be seen in Figure 3 this

⁵ To get a mesh with approximately N_k vertices after k subdivisions $N_o \approx \lceil N_k/4^k \rceil$.

⁶ We obtain triangulations using MATLAB's 'convhulln' routine.

⁷ Inverse mapping to the primary domain does not produce any foldovers if the conformal parameterizations are one-to-one (see subsection 3.1).

operation tends to homogenize the distribution of mesh vertices and increases the quality of the resulting mesh. The step size (e.g. 10^{-3}) and the number of iterations of Tutte mapping (e.g. 40) are two free parameters of the regularization procedure. It is important to note that because the mesh is subdivided and regularized in the parameter domain, its vertices are confined to the original piecewise linear surface.

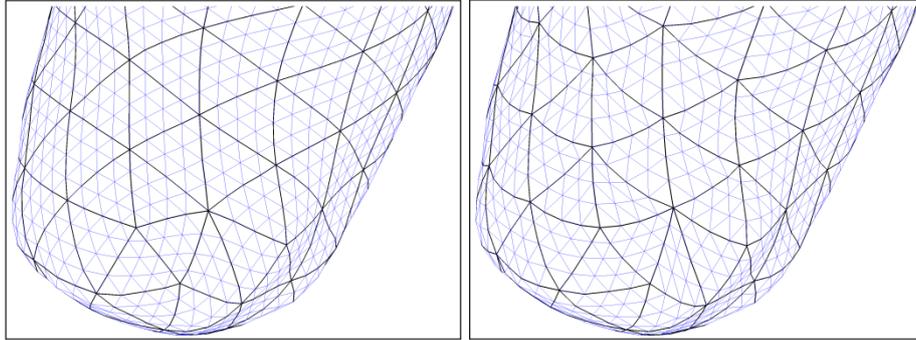


Fig.3. The effect of regularization: left - mesh with regularization, right - mesh without regularization. The thick black lines correspond to the edges of the base mesh. Close-up of the object seen here is the root of the molar shown in Fig.4.

4 Experiments

To demonstrate the effectiveness of our algorithm in generating high quality triangulations, we used it to resample 239 meshes of the carpal bones of the right hand [22] as well as a small number of other anatomical objects that were downloaded from open-access databases [23,24]. Figures 4 and 5 show side by side comparisons of six different anatomical objects before and after resampling. Table 1 provides a corresponding summary of the changes in mesh complexity, quality and approximation errors as measured by the difference in surface area and volume. We assess mesh quality in terms of the minimum and mean triangle aspect ratios (η) [1]. By this definition equilateral triangles are the most desirable and have $\eta = 1$.

$$\eta = 2 \frac{\textit{inscribed radius}}{\textit{circumscribed radius}} \quad (10)$$

The carpal bones shown in Figures 5a-d were obtained from a publically available database [22]. The thin streaks that appear to criss-cross the original surfaces are in fact patches of nearly degenerate triangles that were produced during reconstruction of surfaces with “with a mosaic of individual NURBS patches” [22]. The results summarized in Table 1 show a dramatic improvement in mesh quality at the cost of minor errors (less than 0.25%) in volume and surface area.

All of the mesh processing operations described in this paper were performed in MATLAB running on a desktop with 6 GB RAM and 2.80 GHz x6 AMD processor.

Typical processing time for an average mesh consisting of 13500 vertices and for the base mesh of 500 vertices was 19.5 min, with conformal parameterizations accounting for approximately 60% of the total run time. Although at first this performance may appear inefficient, it should be noted that implementations of iterative algorithms such as the one presented here, are inherently much slower in MATLAB than in 3rd generation programming languages such as C++.

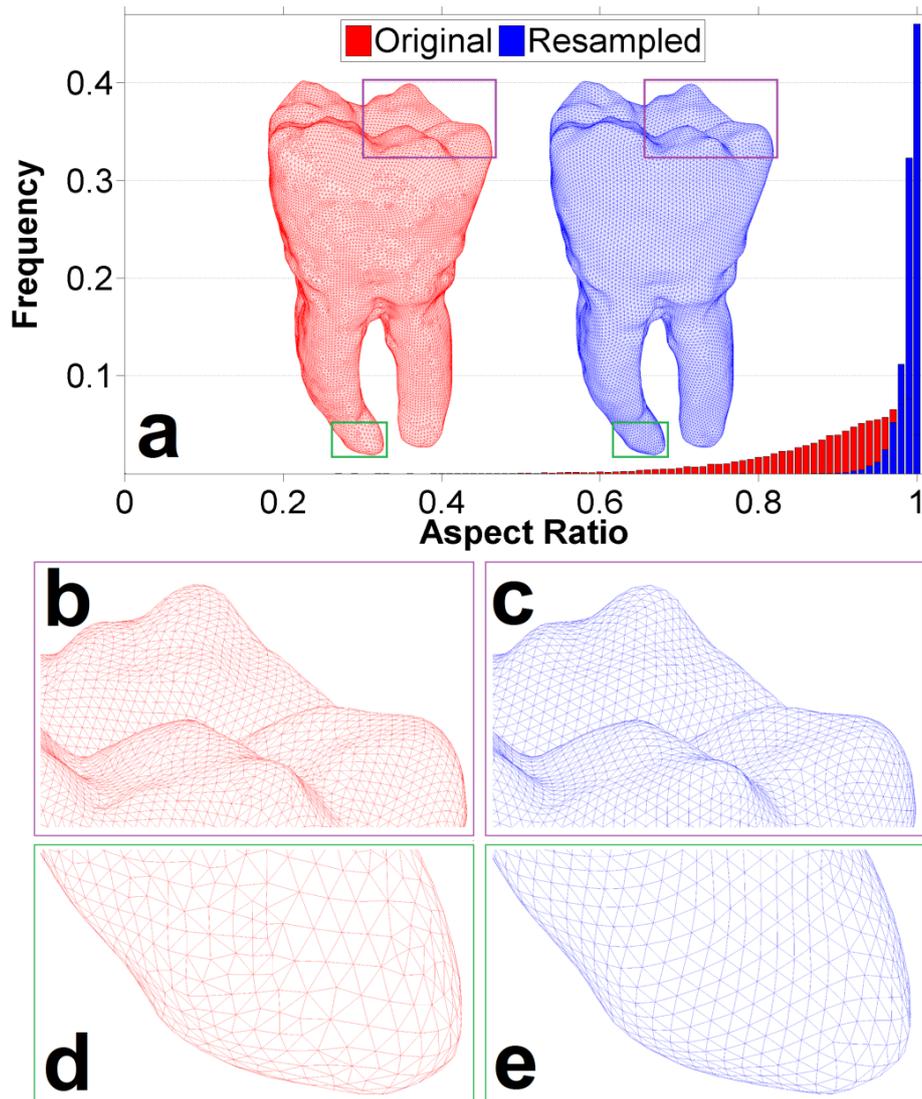


Fig.4. Resampling of the molar. Original mesh is shown on the left and new mesh on the right. Base mesh of the resampled object contained 10^3 vertices. The final mesh was obtained by subdividing the base mesh two times. (b-e) are the close-ups of the rectangular contours in (a).

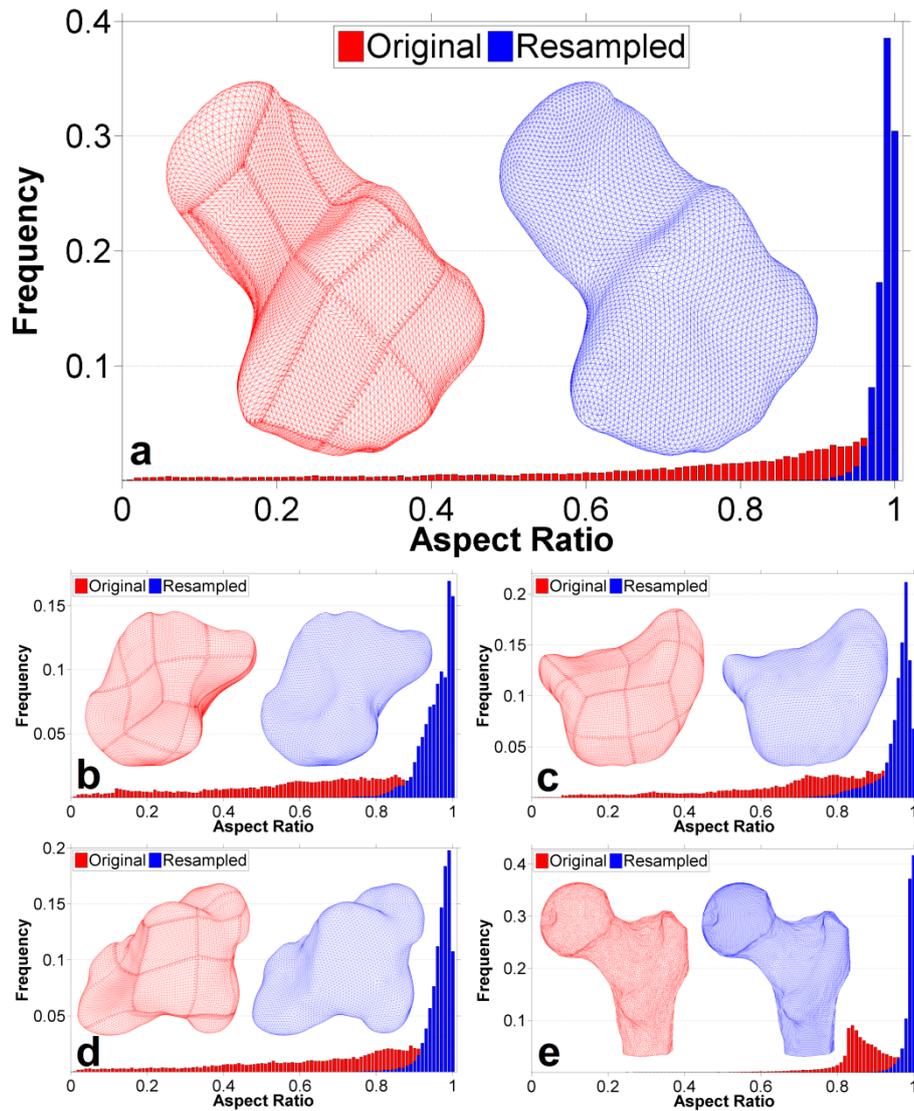


Fig.5. Original and resampled surface models with corresponding triangle aspect ratio distributions; (a) *hamate*, (b) *scaphoid*, (c) *lunate*, (d) *trapezium*, (e) *proximal femur*. Base meshes of the resampled objects shown in (a-d) and (e) contained 500 and 900 vertices, respectively. In all instances the final mesh was obtained by subdividing the base mesh two times. The thin streaks that appear to criss-cross the original surfaces in (a-d) are in fact patches of nearly degenerate triangles that were produced during reconstruction of surfaces with “with a mosaic of individual NURBS patches” [22].

Table 1. Mesh complexity, quality and reconstruction errors before and after resampling for objects shown in Figures 4, 5 and 6. ‘O’ stands for original and ‘R’ for resampled. AE = surface area error = $(A_R/A_O - 1)*100$. VE = volume error = $(V_R/V_O - 1)*100$.

		# verts	# faces	η_{ave}	η_{min}	AE (%)	VE (%)
Molar	O	21900	43796	0.8685	0.2540	-0.3674	-0.0720
	R	15970	31936	0.9890	0.8351		
Hamate	O	12871	25738	0.7466	0.0040	-0.2026	-0.1133
	R	7970	15936	0.9871	0.7951		
Scaphoid	O	12736	25468	0.6666	0.0056	-0.1753	-0.0971
	R	7970	15936	0.9568	0.6912		
Lunate	O	14181	28358	0.7118	0.0129	-0.1647	-0.0919
	R	7970	15936	0.9557	0.6856		
Trapezium	O	13936	27868	0.7173	0.0028	-0.1961	-0.1034
	R	7970	15936	0.9667	0.7438		
Proximal Femur	O	15002	30000	0.8941	0.2678	-0.9194	-0.0788
	R	14370	28736	0.9897	0.7776		

5 Conclusion

We have presented a new method for resampling closed genus-0 surfaces to obtain high quality triangulations. The suitability of the proposed approach was tested on a number of anatomical objects and the preliminary results indicate that the method does indeed greatly improve the quality and visual appearance (see Fig.5a-d) of the input triangulation. Moreover the resulting meshes have the so called subdivision connectivity, a property that can be exploited in compression as well as in mesh processing applications that use hierarchal mesh representations. Future work will involve generalization of the resampling methodology presented herein to surfaces with boundaries and higher genera.

References

1. Shewchuk J.R.: What is a Good Linear Element? Interpolation, Conditioning, and Quality Measures. In: 11th International Meshing Roundtable, pp. 115--126. Sandia National Laboratories, Ithaca, New York (2002)
2. Alliez, G., Ucelli, C., Gotsman, Attene, M.: Recent Advances in Remeshing of Surfaces. In: De Floriani, L., Spagnuolo, M. (eds.) Shape Analysis and Structuring pp. 53--82. Springer, Heidelberg (2008)
3. Gotsman C., Gu X., Sheffer A.: Fundamentals of Spherical Parameterization for 3D Meshes. ACM Trans. Graph. 22, 358--363 (2003)
4. Hormann, K., Labsik, U., Greiner, G.: Remeshing Triangulated Surfaces With Optimal Parameterizations. Comp. Aided Design 33, 779--788 (2001)
5. Praun, E., Hoppe, H.: Spherical Parametrization and Remeshing. ACM SIGGRAPH 22, 340--349 (2003)

6. Marchandise, E., Compère, G., Willemet, M., Bricteux, G., Geuzaine, C., Remacle, J.-F.: Quality Meshing Based on STL Triangulations for Biomedical Simulations. *Int. J. Numer. Method. Biomed. Eng.* 26, 876--889 (2010)
7. Remacle, J.-F., Geuzaine, C., Compère, G., Marchandise, E.: High-Quality Surface Remeshing Using Harmonic Maps. *Int. J. Numer. Meth. Eng.* 83, 403--425 (2010)
8. Zwicker, M., Gotsman, C.: Meshing Point Clouds Using Spherical Parameterization. In: Alexa, M., Rusinkiewicz, S. (eds.) *Eurographics Symposium on Point-Based Graphics*, Zurich (2004)
9. Alliez, P., Meyer, M., Desbrun, M.: Interactive Geometry Remeshing. *ACM SIGGRAPH* 21, 347--354 (2002)
10. Surazhsky, V., Alliez, P., Gotsman, C.: Isotropic Remeshing of Surfaces: A Local Parameterization Approach. In: *Proceedings of 12th International Meshing Roundtable*, pp. 215--224 (2003)
11. Escobar, J.M., Montero, G., Montenegro, R., Rodríguez, E.: An Algebraic Method for Smoothing Surface Triangulations on a Local Parametric Space. *Int. J. Numer. Meth. Eng.* 66, 740--760 (2006)
12. Jiao, X., Bayyana, N.R., Zha, H.: Optimizing Surface Triangulation Via Near Isometry With Reference Meshes. In: Shi, Y., van Albada, G.D., Dongarra, J., Sloot, P.M.A. (eds.) *ICCS 2007, LNCS*, vol. 4487, pp. 334--341. Springer, Heidelberg (2007)
13. Saff, E.B., Kuijlaars, A.B.J.: Distributing Many Points on a Sphere. *Math. Intelligencer* 19, 5--11 (1997)
14. Katanforoush, A., Shahshahani, M.: Distributing Points on the Sphere, I. *Experiment. Math.* 12, 199--209 (2003)
15. Meyer, M., Nelson, B., Kirby, R.M., Whitaker, R.: Particle Systems for Efficient and Accurate High-Order Finite Element Visualization. In: Chen, M., Hansen, C., North, C., Pang, A., van Wijk, J. (eds.) *IEEE Trans. Vis. Comput. Graphics*. vol. 13, pp. 1015--1026. IEEE Press, Sacramento (2007)
16. Shimadat, K., Gossard, D.C.: Bubble Mesh: Automated Triangular Meshing of Non-manifold Geometry by Sphere Packing. In *SMA '95: Proceedings of the 3rd ACM Symposium on Solid Modeling and Applications*, pp. 409-419 (1995)
17. Szczerba, D., McGregor, R., Székely, G.: High Quality Surface Mesh Generation for Multi-physics Bio-medical Simulations. In: Shi, Y., van Albada, G.D., Dongarra, J., Sloot, P.M.A. (eds.) *ICCS 2007, LNCS*, vol. 4487, pp. 906--913. Springer, Heidelberg (2007)
18. Gu, X., Wang, Y., Chan, T.F., Thompson, P.M., Yau, S.-T.: Genus Zero Surface Conformal Mapping and Its Application to Brain Surface Mapping. *IEEE Trans. Med. Imag.* 23, 949--958 (2004)
19. Floater, M.S., Hormann, K.: Surface Parameterization: A Tutorial and Survey. In: Dodgson, N.A., Floater, M.S., Sabin, M.A. (eds.) *Advances in Multiresolution for Geometric Modelling*, pp. 157--186. Springer-Verlag, New York (2005)
20. Duchamp, T., Certain, A., Derosé, A., Stuetzle, W.: Hierarchical Computation of PL Harmonic Embeddings. Technical Report, University of Washington (1997)
21. Möller, T., Trumbore, B.: Fast, Minimum Storage Ray-Triangle Intersection. *JGT* 2, 21--28 (1997)
22. Moore, D.C., Crisco, J.J., Trafton, T.G., Leventhal, E.L.: A Digital Database of Wrist Bone Anatomy and Carpal Kinematics. *J. Biomech.* 40, 2537--2542 (2007)
23. INRIA Gamma, <http://www-roc.inria.fr/gamma/download/ANATOMY/index0.php>
24. Aim@Shape Project - Shape Repository, <http://shapes.aimatshape.net>

Development of Individualized Human Whole-body Thermoregulation Models Using Finite Element Analysis

Tejash Patel¹, Gary P. Zientara², Xiaojiang Xu¹, Reed W. Hoyt¹,

¹ U.S. Army Research Institute of Environmental Medicine (USARIEM),
Biophysics and Biomedical Modeling Division,
Natick, MA 01760-5007, USA

{Tejash.Patel1, Xiaojiang.Xu, Reed.Hoyt } @us.army.mil

² Harvard Medical School and Brigham and Women's Hospital Department of Radiology,
Division of Magnetic Resonance,
Boston, MA 02115, USA
zientara@bwh.harvard.edu

Abstract. Thermoregulatory models capable of predicting human physiological responses to thermal stress are useful tools for heat illness & injury prevention. Extant thermoregulatory models have many geometric simplifications of the physical distribution of heat within the body. The simple nature of these cylinder-based models precludes realistic modeling and prediction of physiological responses of individuals. To address this gap in anatomical accuracy and specificity, a new human thermoregulatory model capable of being individualized is being developed. This paper describes our approach to creating an anatomically-correct and segmented standard model volumetric mesh for finite element analysis, derived from two data sources: the National Library of Medicine's Visible Human Project's whole-body thin section images, and the 'Virtual Family' whole-body voxel labelmaps from the IT'IS Foundation (Zurich, Switzerland). As a first approximation of individual physiology, the standard model multi-component volume meshes will be deformed to match the geometry of an individual's laser surface scan.

Keywords: heat stress, core temperature, medical imaging, finite element analysis, mesh, tetrahedral, COMSOL

1 Introduction

Physically active military personnel can have high rates of metabolic heat production [1]. Physiological mechanisms for dissipating excess metabolic heat are significantly compromised by current protective clothing ensembles, resulting in overheating and diminished work capacity. A number of models based on simplistic cylinder-based body geometry are currently used [2,3,4] to predict human thermoregulatory responses and have proven useful for occupational safety and injury prevention purposes. Human thermal status can be estimated given certain relevant input parameters such as: individual characteristics (height, weight, body fat, heat acclimation, hydration, etc.), clothing characteristics, load weight, work rate, as well as mission parameters (weather, metabolic cost of task) [2]. Model individualization

can extend across these categories, but anatomically the process only includes adjustments for height, weight, and body fat percentage. These types of predictive models show two major functionality gaps. First, existing models simplify human anatomy and individual variability. They employ user input measures such as height and weight, or may employ more advanced schemes to predict the individual variability, but tend to neglect important differences such as fat distribution and tissue composition in specific regions. Second, these models are honed to predict core temperature, a value of high significance, but neglect to depict heat production and storage in local regions of the body. In addition, core temperature itself has been shown to vary widely depending on the mode and anatomic location of measurement.

A new anatomically correct finite-element thermoregulatory model (FETM) is needed to predict the thermo-physiological effects of new protective clothing ensembles on soldiers during activity. Existing cylinder-based thermoregulatory models made many geometric simplifications of the physical distribution of heat within the body [4]. The FETM under development will be capable of simulating regional differences and interactions in thermal physiology and clothing. Further, the finite element (FE) mesh will be used to develop individualized models via mesh warping techniques.

2 Background

Mesh modeling of human anatomy for medical applications is a field growing at a rapid pace with many software tools in development including stand-alone functional software, libraries of primitive functional modules, and various Matlab (Mathworks, Natick MA) applications [5]. Medical imaging data, primarily available as serial digital images of thin sections, magnetic resonance (MR) images and x-ray computed tomography (CT) scans, provide the anatomic information needed to create an anatomically correct mesh for FE analysis. Segmentation, usually laborious, of the image data to form anatomy-specific (i.e., organ, muscle, vessel, bone) volume labelmaps is a necessary step prior to multi-component volume mesh generation. Not insignificant in this combined manual and semi-automated processing pathway is the need for reformatting of data, labelmaps, surface and volume meshes to take advantage of existing software. Much software is available, yet integration demands many daunting data reformatting tasks.

The National Library of Medicine's Visible Human Project is one example of a medical imaging database of anatomically detailed three-dimensional representations of the normal male and female adult human bodies [6]. Cross sectional images were taken at 1mm intervals in the male and at 0.33mm intervals in the female, providing sufficient anatomical information for our FETM model development.

As a time-saving benefit to our project, we have been provided with one computed example of the Visible Human male mesh, courtesy of Dr. Mariette Yvinec at INRIA Sophia Antipolis [7]. These Visible Human male and female meshes will provide one set of 'standard model' volume meshes, to be individualized in a later deformation to a specific soldier's whole-body surface mesh, obtained from a whole-body laser

surface scanner made by Cyberware, Inc. in Monterey, California [8], for our thermal modeling.

As a further aid for us, a research team from the ETH Switzerland has previously segmented other standard human male and female adult imaging data and produce segmented voxel datasets now available through the IT'IS Foundation [9]. Multi-component volume mesh generated from these labelmap datasets provide the input mesh generation for follow-up individualization (through deformation) and thermal modeling, a second research pathway we have exploited.

An exhaustive search revealed a large number of software tools available to accomplish mesh generation, format conversions, mesh refinements/optimization, mesh editing and visualization. The available software spans the range of: well-established and documented versions through crudely simple, undocumented software; source code and/or specific binaries; OS types including Linux, Windows XP, Windows 7, and Mac and 32 bit and 64 bit versions of each of these.

A number of these tools were found to be relevant to our specific processing pathway. 3D Slicer from Brigham and Woman Hospital in Boston MA [10] is a highly developed tool for medical image segmentation (and many other functions) which we use for labelmap production along with ITK-SNAP from the Penn Image Computing and Science Lab [11]. Medit and YAMS developed by Pascal Frey, as well as Gmsh are capable of visualization, refinement, optimization, and reformatting of 3D meshes [12-14]. Other tools such as Meshlab [15], Paraview [16], and MedINRIA [17] allow mesh visualization and reformatting, which allows visual mesh analysis and facilitates data transfer between applications. CGAL is a library of geometry algorithms; specifically, we have used the CGAL 3D tetrahedral mesh generation from the multi-component labelmaps [18]. ImageJ from the National Institutes of Health allows reformatting and visualization of voxel data, an important precursor to mesh building [19]. Tetgen allows the computation of tetrahedral volume meshes from commonly available surface meshes [20].

The ultimate goal of mesh generation, refinement and manipulation with these tools is to import a viable mesh into the FE analysis package COMSOL [21]. At its core, COMSOL is a partial differential equation (PDE) solver which can model the thermal analysis of the human body represented by volumetric FE meshes. Within the COMSOL model, a controlled system will define thermal properties (including metabolic heat production) of the tissues within the body, define their thermal conduction characteristics, as well as define boundary conditions of the tissues. The controlling feedback system will define heat loss through evaporative sweat loss and vasodilatation/constriction as well as metabolic heat production functions. Some tissue thermal properties will be temperature dependant. This analysis profile within COMSOL describes the interaction of each body region's (e.g. extremities, torso, etc.) tissues with the clothing and environment as metabolic heat is generated and dissipated. Probe points are defined to monitor core, extremity, and skin temperature as metabolic heat is transferred through clothing into the environment.

2 Methods

Three mesh processing approaches are attempted to create compatible standard FE meshes that can be individualized to provide input to COMSOL modeling. These are presented in order of increasing complexity, which also represents levels of increasing user control of mesh characteristics.

The Visible Human male adult mesh is a multi-component tetrahedral volumetric mesh (.mesh format) created from segmented tissue data, but it is not in a format readily compatible for FE analysis in COMSOL. To accomplish the mesh format conversion, we first performed mesh refinement in Medit. Then, Gmsh is used to convert the resultant file from .mesh format to COMSOL compatible .vrml format as shown in Figure 1. Although this approach seems direct, it does not produce a viable mesh for our analysis due to the loss of individual tissue segmentation information. Without the segmented tissue information, thermal analysis across the various tissues of the body cannot be performed within COMSOL.

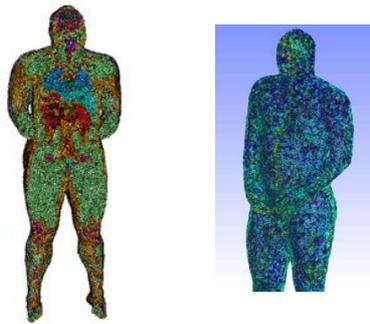


Fig. 1. Screenshot of National Library of Medicine’s Visible Human male. The segmented mesh was originally computed by INRIA (left). Visible Human male mesh refined in Medit and converted to .vrml in Gmsh (right).

Starting from the IT’IS Foundation voxel data, which are volume labelmap datasets of 77 tissue types, the possibility of adding additional structures or combining non-relevant structures into surrounding tissue remains an option (e.g., combining the meniscus with the femur or tibia). ImageJ converts voxel data (.raw format) into its internal Analyze format. The resultant .hdr and .img files are converted to .nrrd format via 3D slicer. Next, MedINRIA converts the .nrrd files to .inr format which is required by the CGAL software as seen in Figure 2.

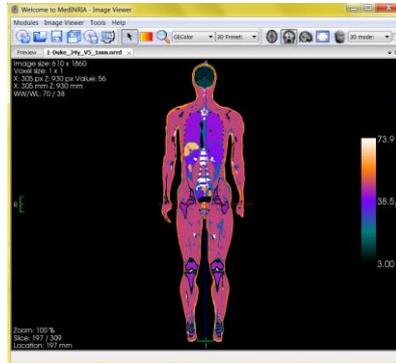


Fig. 2. Screenshot of IT'IS voxel labelmap data within MedINRIA. Voxel size is 1 mm x 1 mm. The image consists of 1860 x 610 voxels.

Thus far in our project, the IT'IS model data has successfully been converted into volumetric meshes. This first required conversion of the labelmaps from .raw to .inr format. Then, the .inr male segmented data (343 MB) was used to create a tetrahedral mesh file in .mesh format with the CGAL Mesh_3 software. The mesh generation processing procedure took 52 seconds on a Dell Precision T3500n Linux workstation with four 3.2 GHz Xeon processors and created a 51 megabyte output mesh. Finally, Gmsh was used to convert this mesh into COMSOL-compatible .vrml format as seen in Figure 3 for the 'Virtual Family' male adult.

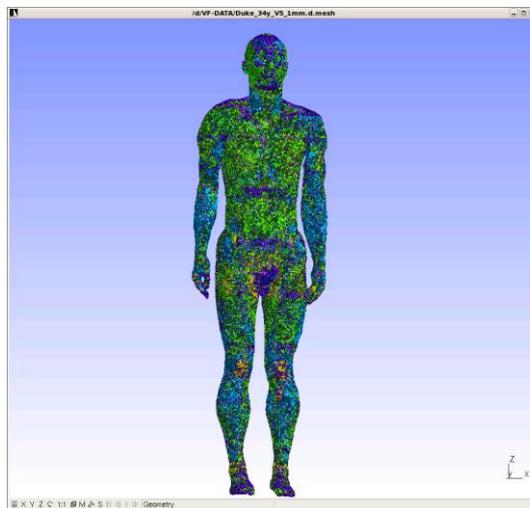


Fig. 3. Screenshot of IT'IS Virtual Family adult male as a tetrahedral mesh within Gmsh which performs the .mesh to .vrmf file type conversion.

Despite the successes of the previous procedures with the Visible Human Male, a freely available segmented version of the Visible Human female in any compatible format does not exist. For the female dataset, we have masked, cropped, and down sampled the original 2048x1530x5190 data (64Gb) and converted to 8 bit grayscale, creating a 492x125x1297 image dataset (79Mb). This whole-body dataset was then separated into overlapping head, thoracic region, abdomen, pelvis, upper and lower leg volumes. These volumes were individually segmented using 3D Slicer (Fig. 4) and ITK-SNAP into specific anatomy represented by labelmaps (.nrrd format). This will go into MedINRIA to create an .inr format file and then into CGAL Mesh_3 program to create a tetrahedral multi-component volume mesh. Currently, we are near completion of the female whole-body segmentation.

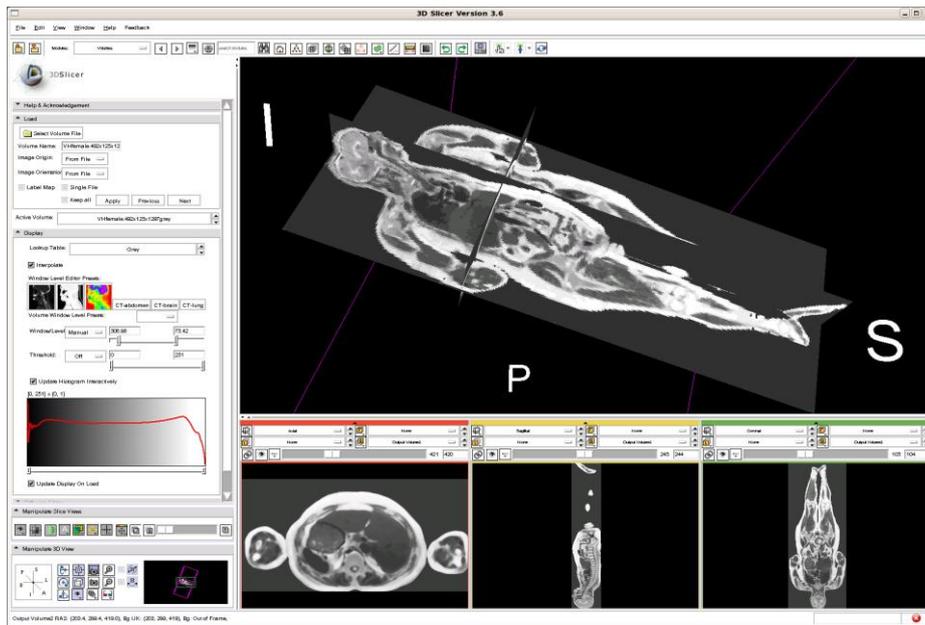


Fig. 4. Screenshot of National Library of Medicine's Visible Human female within 3D slicer. Coronal, sagittal, and transverse planes are shown to aid the user with manual segmentation of anatomical structures.

Future work will focus on individualization of all standard .mesh format meshes originating from VH data and IT'IS data. We will use spline deformation, i.e., elastic warping, of the VH and IT'IS meshes to the surface mesh of the individual soldier acquired with the Cyberware whole-body laser surface scanner to generate a dense

deformation field that will be used to deform the standard model mesh to that of the soldier.

The individualized mesh will then be used for FE analysis within COMSOL. The tissue metabolic heat production, thermal conduction, radiation properties, evaporative properties, and clothing properties are combined with the ambient meteorological conditions. The analysis profile is defined within COMSOL and describes the interaction of each body region's (e.g. extremities, torso, etc.) tissues with the clothing and environment as metabolic heat is generated and dissipated. Probe points are defined to monitor core, extremity, and skin temperature values as metabolic heat is transferred through clothing into the environment.

4 Conclusion

Biomedical modeling is advantageous in that it organizes scientific knowledge and reduces the logistical and financial burdens associated with experimentation with human test volunteers. Creating increasingly higher fidelity anatomically correct models of human thermoregulation may help increase the understanding of the intricate thermal balance in humans. The anatomy embodied in a model can be based on existing medical imaging databases, while future efforts may be able to add individualized anatomic enhancements. Full body MRIs are expensive and time consuming, both in terms of data collection and analysis. We are arguing that warping a standard model to surface scan data provides a rational alternative. In the future could use 3D ultrasound imaging may enable mesh individualization. As the power and functionality of computer hardware and software continues to evolve, it will become more feasible to automate the creation of individual thermoregulatory models.

Our experiences in project computations to date have involved the creation of the standard models and multi-component whole-body meshes of male and female adults from medical imaging data. The step-wise needs to accomplish this processing have demonstrated the pitfalls and problems caused by the multifarious format specifications used by the numerous software programs, libraries and packages in the field. However, on particularly positive note, the state-of-the-art of multi-core PCs and Linux workstations, and the accompanying high capacity compact on-line and off-line storage devices, has brought the capability for manipulation and visualization of whole-body medical data to the scientists desktop. Similarly, the performance of mesh generation and optimization software, and visualization software on desktop computers (now able to include graphics processing units dedicated to high performance computing) is well within practical limits for the ultimate modeling tasks we encounter in thermoregulatory modeling. FETM will be a high-fidelity, individualized, anatomically correct model capable of simulating the regional differences and interactions in thermal physiology, clothing, and environmental conditions in humans.

References

1. Tharion, W.J., Lieberman, H.R., Montain, S.J., Young, A.J., Young, A.J., Baker-Fulco, C.J., DeLany, J.P., Hoyt R.W.: Energy requirements of military personnel. *Appetite*. Vol. 44, Issue 1, pp 47-65. (2005).
2. Kraning, K.K., Gonzalez, R.R., 1997b. Scenario: a military/industrial heat strain model modified to account for effects of aerobic fitness and progressive dehydration. TN97-1 Natick, MA, USARIEM.
3. Xu, X., Berglund, L.G., Chevront, S.N., Endrusick, T.L., Kolka, M.A., 2004. Thermoregulatory model for prediction of long-term cold exposure. *Comput. Biol. Med.* 35, 287–298.
4. Xu, Xiaojiang, Werner Jurgen: A Dynamic Model of the Human/Clothing/Environment-System. *Applied Human Science*. 16, 61-75 (1997).
5. Matlab – The Language Of Technical Computing, <http://www.mathworks.com>
6. U.S. National Library of Medicine – National Institutes of Health – The Visible Human Project ©, http://www.nlm.nih.gov/research/visible/visible_human.html
7. INRIA: Institut National de Recherche en Informatique et en Automatique, <http://www.inria.fr>
8. Cyberware Inc., <http://www.cyberware.com>
9. Foundation for Research on Information Technologies in Society (IT²IS): Human & Animal Models, <http://www.itis.ethz.ch/services/human-and-animal-models/human-models>
10. 3DSlicer: A multi-platform, free and open source software package for visualization and medical image computing, <http://www.slicer.org>
11. itk-SNAP, <http://www.itksnap.org>
12. Medit: OpenGL-based scientific visualization software, <http://www.ann.jussieu.fr/~frey H>
13. YAMS: Optimization of tri/quad surface meshes, <http://www.distene.com/build/optimesh.html>
14. Gmsh: A three-dimensional finite element mesh generator with built-in pre and post-processing facilities, <http://geuz.org/gmsh>
15. MeshLab, <http://meshlab.sourceforge.net>
16. Paraview, <http://www.paraview.org>
17. MedINRIA, <http://www-sop.inria.fr/asclepios/software/MedINRIA>
18. CGAL: Computational Geometry Algorithms Library, <http://www.cgal.org>
19. ImageJ: Image Processing and Analysis in Java, <http://rsbweb.nih.gov/ij>
20. Tetgen: A Quality Tetrahedral Mesh Generator and a 3D Delaunay Triangulator, <http://tetgen.berlios.de>
21. COMSOL Multiphysics, <http://www.comsol.com>

An Atlas-Based Geometry Pipeline for Cardiac Hermite Model Construction

Yongjie Zhang¹, Xinghua Liang¹, Jun Ma¹, Yiming Jing¹, Matthew J. Gonzales², Adarsh Krishnamurthy², Paul Stark³, Sanjiv M. Narayan³, and Andrew D. McCulloch²

¹ Department of Mechanical Engineering, Carnegie Mellon University, USA.
{jessicaz, xliang1, junm1, yjing}@andrew.cmu.edu

² Department of Bioengineering, University of California, San Diego, USA.
{mjgonzales, adkrishnamurthy, amcculloch}@ucsd.edu

³ Division of Cardiology, VA Medical Center, University of California, San Diego, USA.
{psstark, snarayan}@ucsd.edu

Abstract. This paper presents a novel atlas-based geometry pipeline for constructing three-dimensional cubic Hermite finite element meshes of the whole human heart from tomographic patient image data. To build the cardiac atlas, two superior atria, two inferior ventricles as well as the aorta and the pulmonary trunk are first segmented, and then the epicardial and endocardial boundary surfaces are extracted and smoothed. Critical points and skeletons (or central-line paths) are identified, following the cardiac topology. Then, the surface model and the path tree are used to construct a hexahedral control mesh and the final cubic Hermite model via a skeleton-based sweeping method. The thickness of the atria and the ventricles is obtained using segmented epicardial boundaries or via offsetting from the endocardial surfaces in regions where the image resolution is insufficient. Given images from another patient, instead of going through each step again as described above, we choose an optical flow approach to deform the constructed atlas and align it with the new patient’s images. This registration method is fully-automatic, and it avoids tedious manual interactions required by segmentation and path extraction. The constructed Hermite models have been used for finite element analysis of clinically recorded cardiac electrical activity.

Key words: Atlas, geometry pipeline, cubic Hermite finite elements, human heart, atrium, ventricle, image data, registration, optical flow.

1 Introduction

The heart is one of our most vital organs being responsible for delivery of oxygen and nutrients to the whole body. Its dysfunction can commonly be life threatening. As of 2007, heart disease remains the leading cause of mortality, accounting for 25.4% of all deaths in the USA [1]. The heart’s role as a mechanical pump triggered by waves of electrical depolarization, and the dysfunction of the mechanical and electrical systems in many diseases, make it an important target for computational modeling studies. Owing to advances in non-invasive imaging technology, it is now a feasible goal to create patient-specific models to provide insight into complicated disease processes like ventricular arrhythmia, atrial fibrillation and congestive heart failure. Generating an accurate heart geometry from images [18, 19] is well-studied when the resultant geometry

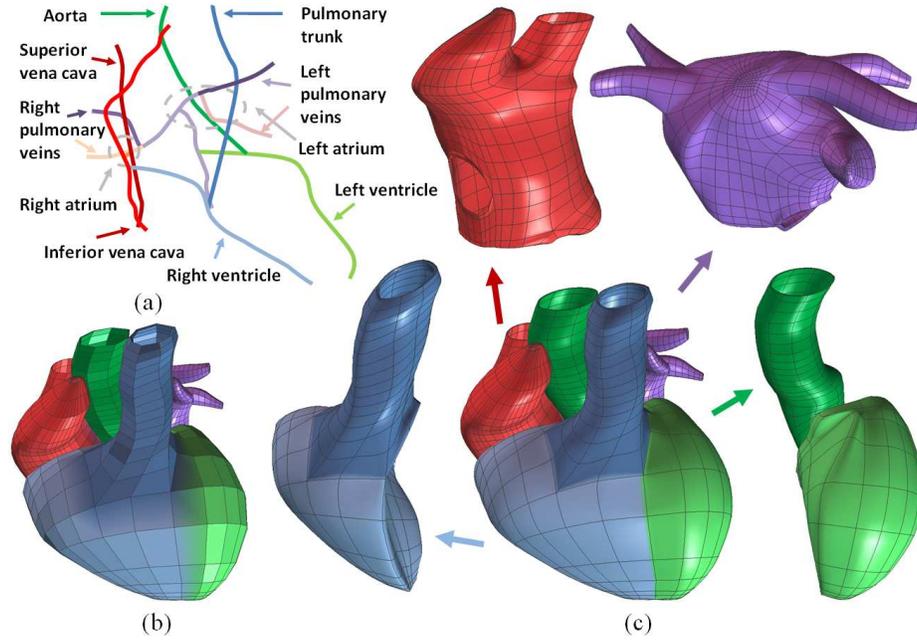


Fig. 1. A whole heart model with four chambers and all major blood vessels. (a) The central-line path tree; (b) the control mesh; and (c) the cubic Hermite model in an exploded view.

is composed of linear elements but is more challenging for high-order C^1 -continuous elements, yet there are many attractive features of a smooth, differentiable geometric model in numerical analysis especially in biomechanics and electrophysiology.

Most work on generation of high-order cardiac geometries from image data has been completed on cubic Hermite meshes, as described by Bradley et al [3]. That study focused on relatively simple geometries such as the chest wall, lungs, and cardiac ventricles, and relied on user-input parameter values that cannot always be anticipated a priori. Additionally, though some cubic Hermite models of the heart have been constructed with more complicated topologies such as the valve annuli [13], construction of such meshes can be labor-intensive. The construction of an accurate four-chamber geometric model based on cubic Hermite splines would require a tremendous amount of user input due to the complicated topology of the atria and the complex connectivity between the four cardiac chambers (see Fig. 1).

In this paper, we describe a novel atlas-based geometry pipeline for constructing three-dimensional cubic Hermite finite element meshes of the human heart from non-invasive medical imaging, taking into account details of the whole heart structure, including the four chambers and all major blood vessels. Our algorithm makes use of minimal user interaction to create the cardiac atlas based on the skeleton of a patient's heart. Once this atlas is obtained, we can create meshes for different patients using deformable registration. To build a cardiac atlas, two superior atria, two inferior ventricles as well as the aorta and the pulmonary trunk were first segmented from cardiac computed tomography (CT) images, and surface models of the epicardial and endocardial boundaries were extracted and smoothed. Critical points and the skeleton (or

central-line path) of each chamber and vessel were identified to represent the topology of the organ. Then, the surface model and the path tree were used to construct a hexahedral control mesh as well as the final cubic Hermite model for the whole heart via a skeleton-based sweeping method. The thickness of the atria and the ventricles was obtained using segmented epicardial boundaries or via offsetting from the endocardial surfaces in regions where the image resolution is insufficient. Given images from another patient, we choose an optical flow approach to construct the new cubic Hermite model by deforming the constructed atlas control mesh to align it with the new patient’s images, instead of going through each step again as described earlier. This registration method is fully-automatic and avoids the tedious manual interactions required by segmentation and path extraction. In addition, it helps construct new cardiac atlases with different anatomic details.

One of the main advantages of our method is that we directly construct cubic Hermite hexahedral meshes. Our Hermite mesh generation algorithm has the following three important features different from other existing methods: (1) a novel atlas-based geometry pipeline to construct cubic Hermite models for the whole heart, including the four chambers and all major blood vessels; (2) a unique 1D center-line path tree to represent the complex cardiac topology and decompose the heart into simple components, which are meshed individually and connected to build the atlas; and (3) an optical flow approach to deform the constructed atlas to match with new patients’ images. In finite element analysis, cubic Hermite meshes have the advantage that the convergence of the solution is faster than using linear meshes that have the same level of detail. In addition, we can guarantee C^1 -continuity for the solution except at extraordinary points. Another advantage of cubic Hermite meshes over linear meshes is that for finite element models of biomechanics [5] and electrophysiology [12], they can achieve equivalently converged solutions not only with substantially fewer elements but also with significantly fewer degrees of freedom. This translates to faster running time for the finite element algorithms. Our constructed Hermite finite element models are suitable for continuum modeling of cardiac electrical and mechanical functions. For example, a bi-atrial Hermite finite element model has been used to simulate atrial electrical activity.

The remainder of this paper is organized as follows. Section 2 overviews the atlas-based geometry pipeline and then the following sections explain details. Section 3 discusses control mesh and cubic Hermite construction for the cardiac atlas. Section 4 describes an optical flow approach to match the constructed atlas with a new patient’s images. Finally, Section 5 shows how clinical measurements of atrial electrical activation can be mapped onto the model geometry, and Section 6 draws conclusions.

2 Atlas-Based Geometry Pipeline

In this project, three patients underwent electrocardiogram-gated 64-slice CT with an iodine-based contrast agent as part of clinical evaluation for atrial fibrillation. Contiguous slices of 1.25-mm were obtained along the axis of the scanner, with in-plane resolutions of 0.5 mm by 0.5 mm (two patients) and 0.7 mm by 0.7 mm (one patient). As shown in Fig. 2, we used images from one patient to construct the atlas. First, four cardiac chambers and all major arteries were manually segmented using ITK-SNAP (www.itksnap.org) with intensity gradients as a guide [17]. From the segmented images, the luminal surface was extracted via isocontouring coupled with geometric flow

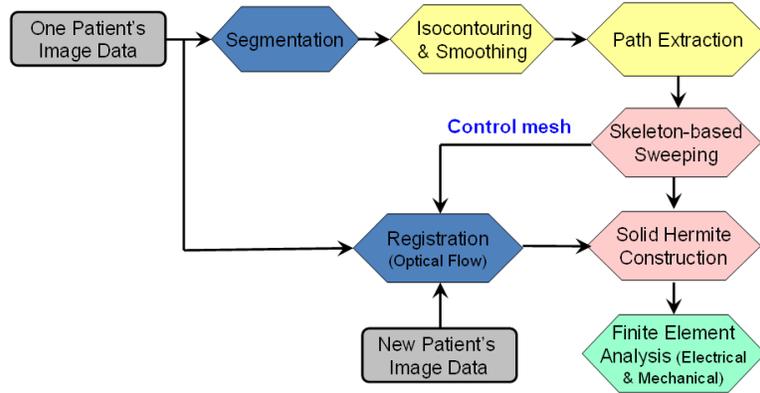


Fig. 2. A schematic diagram of the atlas-based geometry pipeline.

smoothing [20] for two superior atria, the aorta and the pulmonary trunk. It is difficult to obtain their thickness information from images. For the ventricles, the epicardial and endocardial boundary surfaces were extracted and smoothed. After the surface models were obtained, the skeleton (or center-line path) for each component was extracted with some user interaction, defining the cardiac anatomy topology. Then the hexahedral control mesh was constructed via a skeleton-based sweeping method [21], and a cubic solid Hermite model was constructed for finite element analysis of cardiac electrophysiology and biomechanics.

For images from the other two patients, instead of going through each step as described above, here we choose a different and more automated method to construct Hermite models, making use of the atlas model from the first patient. For each new patient, an optical flow approach was used to deform the hexahedral control mesh of the constructed atlas to match it with the new patient's images. In this way, we can avoid the tedious manual interaction required by segmentation and path extraction, and the cubic Hermite model construction becomes fully-automatic. In the following sections, we discuss details of the atlas construction and deformable registration.

3 Atlas Construction via Sweeping

As shown in Fig. 3(a) [1], the human heart consists of multiple connected-components. For the first patient, we obtain a surface model from the segmented images, see Fig. 3(b), and then a cubic Hermite atlas is constructed via the following three steps: path extraction, control mesh, and cubic Hermite construction.

3.1 Path Extraction

Reeb-graphs [10] have been used to analyze surface topologies and construct medial axes or skeletons [7]. As reviewed in [9], there are a lot of techniques developed for skeleton extraction, including topological thinning [2], distance field based methods [22], potential field based methods [4], thinning via medial geodesic function [6] and others [15]. Here we choose the critical point theory of the distance function [21] to define a novel skeleton (or central path) tree for the whole cardiac model, following

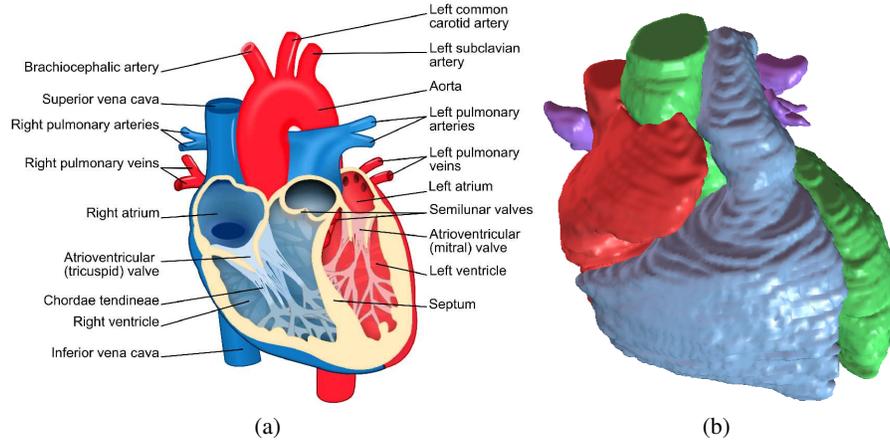


Fig. 3. (a) The human heart anatomy [1]; and (b) the segmented heart with four chambers, the aorta and the pulmonary trunk.

the blood flow path inside the heart and the cardiac anatomy structure. The primary principle we comply with is the “equidistant rule”, which means the skeleton point in a certain cross section should be the centroid of that cross section. Owing to the quasi-tubular anatomical structure of the four chambers, we can conveniently calculate the centroid at an arbitrary cross section given that the corresponding isocontouring surface and normal vector are designated. Sometimes the shape of a cross section is far away from a regular circle, so that adjustments are made to select correct skeleton points. With the collection of skeleton points for each cross section, we obtained the 1D path tree with each color representing the skeleton of one component, see Fig. 1(a).

3.2 Hexahedral Control Mesh Construction

In the sweeping method [21], a templated quadrilateral mesh of a circle was projected onto each cross section along the skeleton, then the corresponding vertices in adjacent cross-sections were connected to form a hexahedral control mesh. A hexahedral cubic Hermite control mesh should satisfy the following three requirements: (1) no intersection is allowed between any two cross sections; (2) each cross section should be perpendicular to the skeleton; and (3) in order to achieve a G^1 -continuous surface around extraordinary nodes, the boundary vertex shared by two patches in the control mesh should be collinear with its two neighbors across the shared boundary, and the boundary vertex shared by three or more patches should be coplanar with all of its neighboring boundary vertices. This is because, for a so-called open knot vector, a cubic Hermite curve is tangent to the control polygon at the first and the last control nodes.

One-to-one sweeping requires that the source and target surfaces have the same topology. The cardiac model has four chambers connecting with surrounding arteries, therefore, we first decompose the cardiac topology into branches or components based on the central-line path tree in Fig. 1(a).

Branch construction: For each branch, we use the one-to-one sweeping method to construct the control mesh. The cross-section template introduced in [21] induces a

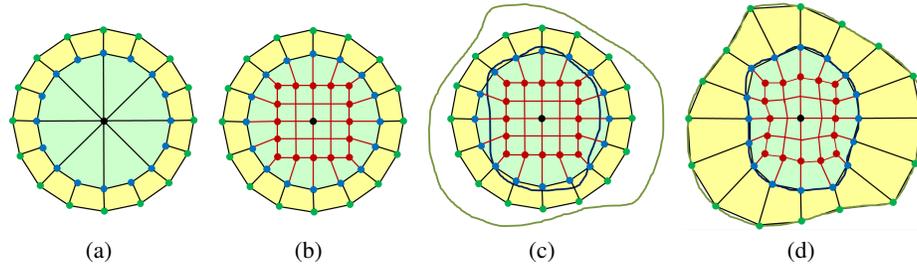


Fig. 4. Branch template. (a) The cross-section template provided in [21]; (b) the new circle-square template; (c) the circle-square template with epicardial and endocardial boundaries; and (d) the circle-square template conformal to the epicardial (green points) and endocardial (blue points) boundaries. The black point is the center, and the red points locate at the square edges.

large valance number at the tip, see Fig. 4(a), which is undesirable in cubic Hermite finite element interpolation. As an improvement, a new circle-square template is introduced in Fig. 4(b). This template limits the valance to be within four, and provides elements with better quality.

The basic idea of the circle-square template is that we use a square mesh to replace the center part in Fig. 4(a). As shown in Fig. 4(b), the number of peripheral nodes can only be 2^n ($n \geq 2$). The center point locates on the skeleton. To generate a conformal mesh, each green node is projected to the epicardial boundary, and each blue node is projected to the endocardial boundary. Red nodes are adjusted proportionally via smoothing, see Fig. 4(d). This procedure is repeated for each cross section along the skeleton. For most cross sections, only the yellow region in Fig. 4 is needed. However, the inner green region is required for enclosing the tips of each chamber. By connecting the corresponding control nodes in adjacent cross-sections, we obtain a hexahedral control mesh for one branch. Notice that in the process of sweeping, we translate the cross-section template to the selected locations on the skeleton, and rotate it to make its normal vector lie along the central-line path.

Ventricle construction: In contrast to one branch construction, the ventricle has two inner surfaces and one outer boundary. To mesh the ventricle muscle structure, we design another template as shown in Fig. 5. The two inner black points are located on the skeleton, and a middle line is defined to separate the two inner surfaces. For each

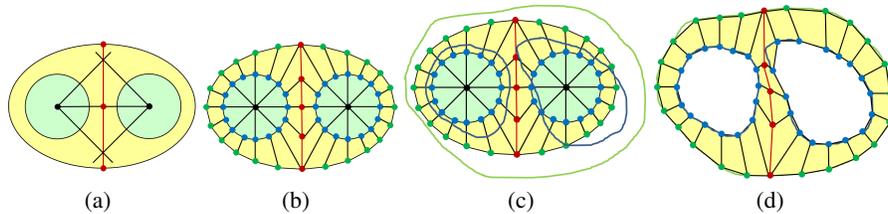


Fig. 5. Ventricle template. (a) Locate two inner center points (black) on the skeleton and determine the middle-line points (red); (b) the sweeping template; (c) the sweeping template with one epicardial (green) and two endocardial (blue) boundaries; and (d) project green points to the epicardial boundary, and blue points to the two endocardial boundaries.

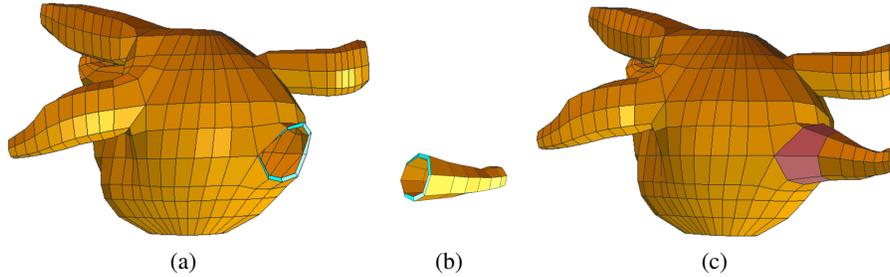


Fig. 6. Branch connection. (a) The chamber body with an open mouth; (b) one slave branch; and (c) the left atrium model after branch connection.

cross section, we first create a sweeping template like Fig. 5(b). Then, taking the black center points as the references, green points are projected to the epicardial boundary, and blue points are projected to the corresponding endocardial boundaries. The middle line is also adjusted by moving the corresponding red points to the middle of the septum. The final control mesh is obtained by connecting the corresponding nodes in adjacent cross-sections.

Branch connection: To obtain a hexahedral control mesh for the whole heart, we need to connect all the separated branches or components one by one following the cardiac topology. For each connecting procedure, a so-called master branch and a so-called slave branch are chosen. Typically, the master branch is a chamber or artery with a larger diameter, and the slave branch is relatively smaller.

For example, let's choose the left atrium in Fig. 6(a) as the master branch and one small artery in Fig. 6(b) as the slave branch. First, we choose one segment on the master branch which is close to the slave branch and has half number of nodes as on the cross-section of the slave branch. Then these nodes are duplicated. Half of them are moved up, and half of them are moved down, resulting in an open mouth on the left atrium, see Fig. 6(a). Now the open mouth has the same topology as the cross-section of the slave branch. Then, we can easily connect the mouth with the slave branch to obtain an integrated control mesh, see Fig. 6(c). This mesh will be considered as the master branch in the next connecting procedure. Following this way, we connect all the components together and obtain a hexahedral control mesh for the whole heart, see Fig. 1(b).

3.3 Cubic Hermite Mesh Construction

One solid Hermite patch is constructed for each hexahedral element in the control mesh: the inner and outer surfaces are 2D bi-cubic Hermite, and the thickness direction is interpolated linearly. First, let's take a look at a 1D Hermite curve

$$P(u) = h_{00}P_0 + h_{10}P_1 + h_{01}P_0'' + h_{11}P_1'', \quad (1)$$

where $h_{00} = 2u^3 - 3u^2 + 1$, $h_{10} = -2u^3 + 3u^2$, $h_{01} = u^3 - 2u^2 + u$, $h_{11} = u^3 - u^2$, and $u \in [0, 1]$ is the variable. Fig. 7(a) shows these four Hermite blending functions. P_0 and P_1 are two endpoints. P_0'' and P_1'' are the first derivatives at the two endpoints. Similarly, the parametric form of a 2D Hermite surface can be written as

$$P(u, v) = U^T M_H C M_H^T V, \quad (2)$$

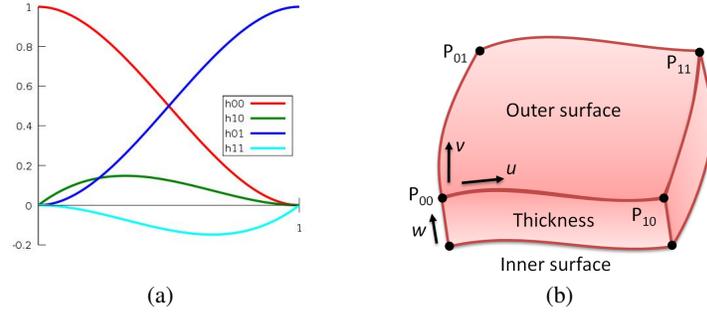


Fig. 7. (a) 1D Hermite blending functions; and (b) one bi-cubic-linear solid Hermite patch.

where

$$U = \begin{bmatrix} u^3 \\ u^2 \\ u \\ 1 \end{bmatrix}, \quad V = \begin{bmatrix} v^3 \\ v^2 \\ v \\ 1 \end{bmatrix}, \quad M_H = \begin{bmatrix} 2 & -2 & 1 & 1 \\ -3 & 3 & -2 & -1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}, \quad C = \begin{bmatrix} P_{00} & P_{10} & P_{00}^u & P_{10}^u \\ P_{01} & P_{11} & P_{01}^u & P_{11}^u \\ P_{00}^v & P_{10}^v & P_{00}^{uv} & P_{10}^{uv} \\ P_{01}^v & P_{11}^v & P_{01}^{uv} & P_{11}^{uv} \end{bmatrix}.$$

Here $P^u = \frac{\partial P}{\partial u}$, $P^v = \frac{\partial P}{\partial v}$, and $P^{uv} = \frac{\partial^2 P}{\partial u \partial v}$. For each element on one cross section as shown in Fig. 7(b), the peripheral direction is referred to u ($u \in [0, 1]$), and the normal direction is referred to v ($v \in [0, 1]$). Linear interpolation is conducted along the wall thickness direction. Fig. 8 shows the bi-cubic-linear solid Hermite constructed from the bi-atrial control mesh. Each element of the control mesh is subdivided into 4x4 sub-elements using the bi-cubic-linear Hermite interpolation for visualization.

The constructed bi-cubic-linear Hermite is C^1 -continuous over the heart surface except the local region around extraordinary nodes, which is C^0 -continuous. These extraordinary nodes are induced when connecting two separated branches. To obtain G^1 -continuity for them, we need to adjust control nodes surrounding each extraordinary node: the boundary node shared by two patches in the control mesh should be collinear with its two neighbors across the shared boundary, and the boundary node shared by three or more patches should be coplanar with all of its neighboring boundary nodes.

4 Deformable Registration using Optical Flow

After constructing one Hermite atlas for the first patient, we set his image data as the static image S (or the reference image), and set a new patient's image data as the moving image M (or the target image). To construct the cubic Hermite model for the new patient, instead of going through each step as explained in Section 3, we deform the atlas control mesh to match it with the moving image M by minimizing the difference between S and M . In this way, we can construct the cubic Hermite model automatically for any new patient, avoiding the tedious manual interaction required by segmentation and path extraction. In addition, this registrations method provides an efficient approach to build an atlas database for the human heart.

Here we choose an enhanced optical flow algorithm, which is primarily based on Thirion's diffusing model, also known as the "demons" algorithm [14, 16]. The "demons"

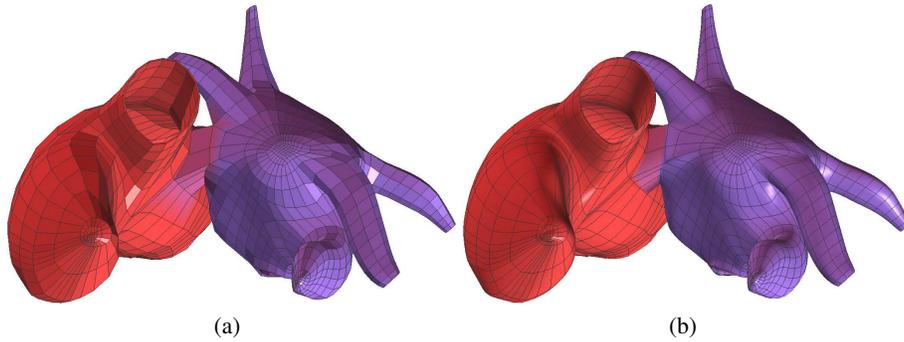


Fig. 8. A bi-atrial model. (a) Control mesh; and (b) bi-cubic-linear Hermite model.

algorithm calculates the demons force, using the gradient value from the static image S , in order to match with the moving image M . Usually, the optical flow formula was applied to calculate the “demons” force at each grid point in a greyscale image,

$$\vec{f}_s = \frac{(m-s)\vec{\nabla}s}{|\vec{\nabla}s|^2 + (s-m)^2}, \quad (3)$$

where s , m are the intensity values in the static image S and the moving image M , respectively. $\vec{\nabla}s$ is the gradient, and $\vec{f}_s = (\vec{u}_x, \vec{u}_y, \vec{u}_z)$ is the displacement vector also called the “passive” force.

The original algorithm may not be efficient, especially when image varies little among neighbouring grid points in a local region. Based on Newton’s third law of motion, Rogelj and Lovaicic introduced a new force [11]. The advantage of this accelerated algorithm was that it made use of the information from both static and moving images, which could speed up the rate of convergence. Another force named as an “active” force is introduced based on the information from the moving image M ,

$$\vec{f}_m = -\frac{(s-m)\vec{\nabla}m}{|\vec{\nabla}m|^2 + (s-m)^2}. \quad (4)$$

The term “passive” force denotes the contribution to the force from the static image S . Similarly, the term “active” force denotes the influence from the moving image M . The reason why the second term was named “active” might be that the equation iteratively calculates the deformation to match with the moving image M and it was active to track the corresponding point in M . Combining both the “passive” force \vec{f}_s in Eqn (3) and the “active” force \vec{f}_m in Eqn (4), the total force at a specific grid point can be calculated as

$$\vec{f} = \vec{f}_s + \vec{f}_m = (m-s) \times \left(\frac{\vec{\nabla}s}{|\vec{\nabla}s|^2 + (s-m)^2} + \frac{\vec{\nabla}m}{|\vec{\nabla}m|^2 + (s-m)^2} \right). \quad (5)$$

To calculate the displacement for each vertex in the atlas control mesh, we use a trilinear interpolation on the 3D regular grids. As shown in Fig. 9, Eqn (5) is calculated iteratively in the “demons” algorithm. The demons set Ds is pre-computed from S . Here we choose the whole image grid as Ds . After each iteration, a stopping criterion is required to determine when the program ends. If, for each mesh vertex, the maximum

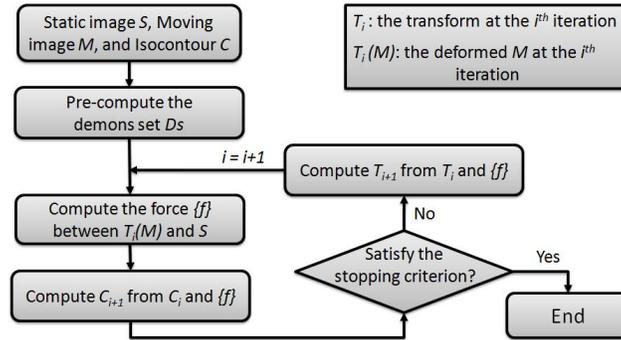


Fig. 9. Flowchart of the optical flow scheme.

difference of its displacement is less than a pre-defined threshold such as 0.01, which is roughly 10% of the minimum span among X, Y and Z coordinates, the program stops. In each iteration, a regulation of the deformation field follows this optical flow calculation, using a Gaussian filter in which the variance of σ^2 is set as 1.0. The regulation plays an essential role as a smoothing operation to remove noise and preserve the geometry continuity, when the displacement is calculated merely using the local information. To obtain precise displacements, a scaling factor needs to be included in the algorithm to compensate the differences in intensities between the static and moving images.

We set Patient 61 as the static image, and Patient 64 or 66 as the moving image. The optical flow method is used to deform Patient 61's control mesh to obtain the control mesh for Patients 64 and 66. As shown in Fig. 10, we can observe that our registration algorithm catches some detailed differences between various patients.

5 Cardiac Electrical Activity Analysis

To facilitate visualization of electrical activation data, the pulmonary veins of the left atrial control mesh were shortened. The local activation times of tissue were projected onto the left atrial shell and fit as a scalar field based on a least-squares problem with Sobolev smoothing weights [3]. This least-squares problem was solved for a smooth

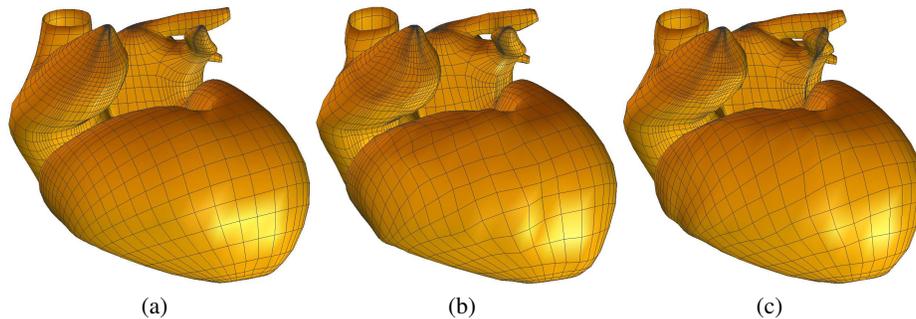


Fig. 10. Results for deformable registration. (a) Control mesh of Patient 61 (reference model); (b) the obtained control mesh for Patient 64; (c) the obtained control mesh for Patient 66.

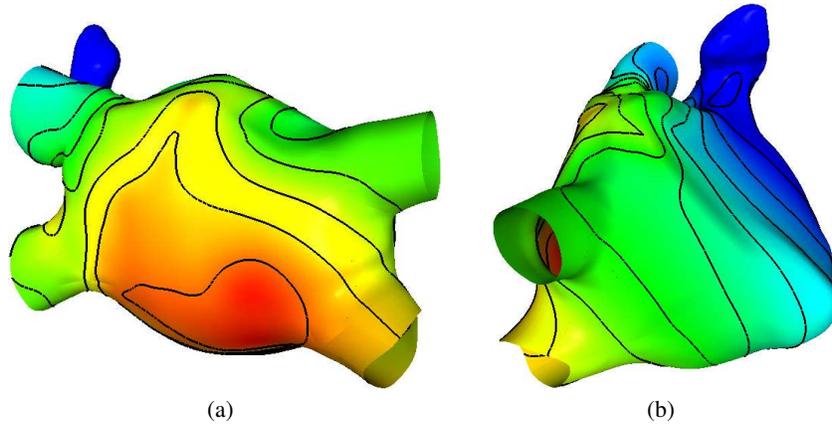


Fig. 11. Result for cardiac electrical activation analysis. (a) Posterior view; and (b) side view.

scalar field representing activation time of that region of the atrium. Regions of the mesh have no measurements, so the values of activation times had to be extrapolated.

Fig. 11 depicts the measured activation sequence of the human left atrium as electrical activity propagates during an in-vivo electroanatomic mapping study. A stimulus was applied on the floor of the left atrium, and local times of activation were identified from potential waveforms recorded by 32 bipoles. Red regions indicate early activation, and blue regions indicate late activation. The black lines indicate isochrones, or regions of identical temporal activation, at a space of 5 milliseconds. The relative slowing of activation near the pulmonary veins on both the left and the right sides is consistent with the abrupt changes in fiber bundles and wall thickness that have been observed near the pulmonary veins in humans [8].

6 Conclusion

In this paper, we have developed a novel atlas-based geometry pipeline to construct cubic Hermite finite element models from image data. From the segmented images of one patient, the surface model and the 1D center-line path tree were extracted, and a hexahedral control mesh was constructed via a skeleton-based sweeping method. Given images of another patient, instead of going through each step again during the atlas construction, an optical flow method was used to deform the constructed atlas to match with the new patient's images. In this way, an atlas-based cardiac database can be built efficiently and automatically. For different patients, the cardiac topology may have different anatomic details. By using this registration method, we can also build new atlases. For a new cardiac topology, we can first search in the database which atlas is the most similar one, then we deform that one to construct a new atlas. As part of our future work, we will test our pipeline on more patients' data and build a larger database by collaborating with VA Medical Center in University of California, San Diego.

Acknowledgments: We would like to thank W. Wang for useful discussion on Hermite construction. This research project was supported by NIH grants P41 RR08605, R01 HL096544 (ADM), R01 HL83359 (SMN) and K24 HL103800 (SMN).

References

1. http://en.wikipedia.org/wiki/human_heart.
2. G. Borgefors, I. Nystrom, and G. Baja. Computing skeletons in three dimensions. *Pattern Recognition*, 32(7):1225–1236, 1999.
3. C.P. Bradley, A.J. Pullan, and P.J. Hunter. Geometric modeling of the human torso using cubic Hermite elements. *Annals of Biomedical Engineering*, 25(1):96–111, 1997.
4. N. Cornea, D. Silver, X. Yuan, and R. Balasubramaniam. Computing hierarchical curveskeletons of 3D objects. *The Visual Computer*, 21(11):945–955, 2005.
5. K.D. Costa, P.J. Hunter, J.S. Wayne, L.K. Waldman, J.M. Guccione, and A.D. McCulloch. A three-dimensional finite element method for large elastic deformations of ventricular myocardium: Part II - Prolate spheroidal coordinates. *J Biomech Eng*, 118:464–472, 1996.
6. T. K. Dey and J. Sun. Defining and computing curve-skeletons with medial geodesic functions. In *Sympos. Geom. Processing*, pages 143–152, 2006.
7. F. Lazarus and A. Verroust. Level set diagrams of polyhedral objects. In *ACM Symposium on Solid Modeling and Applications*, pages 130–140, 1999.
8. V. Markides, R.J. Schilling, S.Y. Ho, A.W. Chow, D.W. Davies, and N.S. Peters. Characterization of left atrial activation in the intact human heart. *Circulation*, 107:733–739, 2003.
9. D. Silver N. Cornea and P. Min. Curve skeleton applications. *IEEE Visualization*, pages 95–102, 2005.
10. G. Reeb. Sur les points singuliers d'une forme de pfaff complètement intégrable ou d'une fonction numérique. *Comptes Rendus de L'Académie des Séances, Paris*, 222:847–849, 1946.
11. P. Rogelj and S. Kovacic. Symmetric image registration. In *Proceedings of SPIE, Vol. 5032, Medical Imaging 2003: Image Processing*, pages 484–493, 2003.
12. J.M. Rogers, M. Courtemanche, and A.D. McCulloch. *Finite element methods for modeling impulse propagation in the heart. Chapter 7 in: Panfilov AV, Holden AV, eds. Computational Biology of the Heart*. Sussex: John Wiley and Sons, Ltd, 1996.
13. C. Stevens and P.J. Hunter. Sarcomere length changes in a 3D mathematical model of the pig ventricles. *Prog Biophys Mol Biol*, 82(1-3):229–41, 2003.
14. J. P. Thirion. Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis*, 2(3):243–260, 1998.
15. A. Verroust and F. Lazarus. Extracting skeletal curves from 3D scattered data. *The Visual Computer*, 16:15–25, 2000.
16. H. Wang, L. Dong, J. O'Daniel, R. Mohan, A. Garden, K. Ang, D. Kuban, M. Bonnen, J. Chang, and R. Cheung. Validation of an accelerated demons' algorithm for deformable image registration in radiation therapy. *Phys in Med and Biol*, 50(12):2887–2905, 2005.
17. P.A. Yushkevich, J. Piven, H.C. Hazlett, R.G. Smith, S. Ho, J.C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability. *NeuroImage*, 31:1116–1128, 2006.
18. Y. Zhang and C. Bajaj. Finite element meshing for cardiac analysis. *ICES Technical Report 04-26, The University of Texas at Austin*, 2004.
19. Y. Zhang, C. Bajaj, and B-S. Sohn. 3D finite element meshing from imaging data. *Computer Methods in Applied Mechanics and Engineering*, 194:5083–5106, 2005.
20. Y. Zhang, C. Bajaj, and G. Xu. Surface smoothing and quality improvement of quadrilateral/hexahedral meshes with geometric flow. *Communications in Numerical Methods in Engineering*, 25(1):1–18, 2009.
21. Y. Zhang, Y. Bazilevs, S. Goswami, C. Bajaj, and T. J.R. Hughes. Patient-specific vascular NURBS modeling for isogeometric analysis of blood flow. *Computer Methods in Applied Mechanics and Engineering*, 196(29-30):2943–2959, 2007.
22. Y. Zhou and A. Toga. Efficient skeletonization of volumetric objects. *IEEE Trans. Visualization and Comp. Graphics*, 5(3):196–209, 1999.

Ultra Fast Optical Sectioning: Signal preserving filtering and surface reconstruction

Rasmus R. Jensen, Mike van der Poel, Rasmus Larsen, and Rasmus R. Paulsen

Informatics and Mathematical Modelling, Technical University of Denmark
Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark
raje@imm.dtu.dk, www.imm.dtu.dk

Abstract. In 3D surface scanning it is desirable to filter away *bad* data without altering the quality of the remaining *good* data. Filtering of raw scanner data before surface reconstruction can minimize the induced error and improve on the probability of reconstructing the true surface. If outliers consist of actual data such as hair, and not just evenly distributed noise, these outliers tend to err smoothing algorithms away from the wanted result. We present a novel algorithm based on a Markov Random Field that uses a distance constraint to robustly classify a 3D scan volume. Through this classification a signal preserving filtering of the data set is done. The remaining data are used for a smooth surface reconstruction creating very plausible surfaces. The data used in our work comes from a newly developed hand held 3D scanner. The scanner is an Ultra Fast Optical Sectioning scanner, which is able to extract high quality 3D surface points from 2D images recorded at over 3000 fps. The scanner has been developed for digital impression taking in the dental area. Our work relates to future in-ear scanning for fitting custom hearing aids without impression taking.

-Keywords: *3D scanning, Markov Random Field, computer vision, surface reconstruction, noise filtering.*

1 Introduction

3D surface acquisition is an established and active research and development area. Novel applications and devices continues to emerge, where the data scale ranges from minuscule in microscopy optical sectioning [7] to large scale aerial surface laser scanning of the earth [13]. A variety of scanners and cameras exist; each with their own strengths and weaknesses. While some scanners produce 3D surface data, scanning an object from one direction is known as 2.5D scanning as it only portrays the object from one side and does not provide a full 3D model. To construct a 3D model several 2.5D scans need to be patched creating a full reconstruction [8].

We have worked with a new scanner, the Ultra Fast Optical Sectioning TRIOS scanner from 3Shape[1], which has been developed to facilitate 3D impression taking in the dental area. Our work is a study on how to reconstruct surfaces in the presence of structured noise. The study is a preliminary study,

which relates to fitting custom hearing aids, where the construction of an in-the-ear scanner would make the ear canal impression step obsolete.

Anticipating the problem of structured noise from hair in the ear canal we want an algorithm that filters the data and leaves only valid surface data. As no scanner has yet been produced that will actually go into the ear, we have used data of hairy arms and bearded chins recorded by the TRIOS scanner. The scanner produces high quality data with some very sparse *salt-and-pepper* type noise and also good scanning of actual hair strands. As the scanner is a dental scanner with both high precision and accuracy, we were faced with a specific problem of removing only the hair without degrading the remaining data.

Simple mean or median filtering can make any surface fair (if one smoothens enough) but these filters also distort the data set and are useable only when the noise is Gaussian (mean filtering) or when the number of outliers are few (median filtering). An adaptive application of such filters used only on outliers can remove noise without too much degrading of the data, but such an approach would break down in areas with a lot of outliers. Generally, local filtering does only preserve local structure; for areas with a lot of hair a global method is needed.

Implicit functions err towards outliers and noise. Splines [10] are a well known tool for creating both smooth curves and surfaces, but if noise is not Gaussian the smoothing will be skewed. This is the case in our data, where outliers are mainly found above the surface. Splines are also continuous and therefore do not handle discontinuity well.

A way of removing outliers in a data set is to use random sampling such as RANSAC [4] and fit to the random sample until a fit matches the data well. Even though this removes the influence of outliers it does not guarantee an optimal match. A RANSAC approach is described in [11], where an algorithm is created that finds basic shapes and structures in noisy data sets.

There exist a large body of literature covering noise properties and handling in direct surface scanners [3]. However, the used scanning device is not directly comparable with devices previously investigated.

To maintain as much of the good data as possible in our scans, we have solved our problem based a *Markov Random Field* (MRF) formulation on a 3D voxel grid. An early description of MRFs in 2D image analysis is on the noise removal in dirty pictures addressed in [2]. The novelty of our work is the use of a 3D MRF with a distance based smoothness prior that classifies the data set into surface data and not surface data. The classification allows for a signal preserving filtering of the data set before any actual surface reconstruction.

2 Data

The data come from an Ultra Fast Optical Sectioning TRIOS scanner, which records a stream of 2D images. Approximately 130 images are collected into a set that constitutes a voxel volume. From known changes in the scanner during the volume acquisition a scan surface is constructed which can be converted

into a real world coordinate system through a calibration step. The scanner computes the voxels that defines the interface between air and solid material as seen from the scanners viewpoint. Since the scanner is known to be above the surface it is possible to label the voxels along the depth axis as either *above* or *below* the sought surface. When the scanner firmware has determined which voxels belong to the interface, the row, column and depth volume is transformed into real world 3D points using the calibrating parameters. The resulting 2.5D point cloud represents the interface between air and solid material as seen from the scanners viewpoint. To perform a full 3D surface scan, the scanner is moved around the object and the partial scans are merged together using a proprietary algorithm.

The scanner is normally used for digital impression taking in dental work, which requires a very high accuracy. Therefore, the quality of scans are very high and the noise levels due to the scanner hardware is minimal. However, real physical objects as for example hair will also be captured by the scanner. In the current application (direct ear-scanning) we are interested in the true surface of the ear and therefore hair is considered noise and should be removed from the scan. In the following we consider hair as being structured noise.

A 2.5D scan can also be considered a depth map, where the pixel value reflect the distance to the object. Figure 1 shows four depth maps of scans from different surfaces. The scans constitute a range from very bad to perfect, and the algorithm should be able to handle all examples. When there is a lot of structured noise in the scan, we need an algorithm that does not break down but does a good filtering leaving only the actual surface data (even if it is very little).

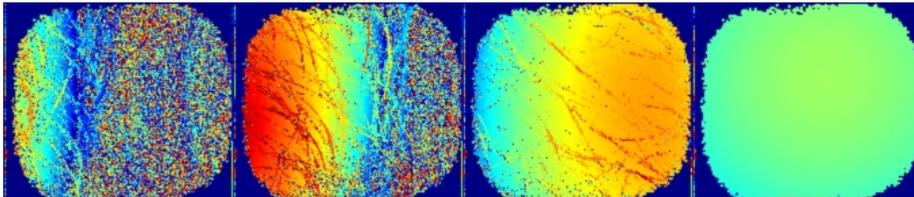


Fig. 1. Four depthmaps of scan surfaces. From left to right: Little surface coverage with hair, half coverage with hair, full coverage with hair, full coverage without hair.

3 Markov Random Field volume classification

The raw output from the scanner (and the scanner firmware) is a voxel set where the scanner is virtually placed above the voxel volume looking down the depth direction. Each voxel is labeled as being either *above* or *under* the *scan surface*. The scan surface, S_{scan} , is the initial surface that can be extracted as the interface between the *above* and *under* voxel sets. However, S_{scan} is noisy (in the

sense that hair is present in the scan) and does not represent the true underlying skin surface. We aim to produce a consistent and locally smooth skin surface, S_{skin} , from the data set. In order to re-label the voxel set and thereby implicitly producing, S_{skin} , a Markov Random Field (MRF) classification/regularisation approach is chosen. In the following, a short introduction to MRFs and a description of the chosen models are given.

3.1 The volume random field

We define a random field with spatial voxel positions $\{v_1, v_2, \dots, v_n\}$ in the volume V with the index set I . In this set each voxel v_i takes a value x_i from the binary label set $L = \{\textit{under}, \textit{above}\}$, where *under* is under and *above* is over the skin surface. Notice that we make a distinction between scan surfaces S_{scan} and skin surfaces S_{skin} ; the classification relates to the latter. All values of x_i are represented by the vector \mathbf{x} , which is the configuration of the random field.

A neighbourhood system to v_i is defined as $N = \{N_i | i \in I\}$ for which it holds that $i \notin N_i$ and $i \in N_j \Leftrightarrow j \in N_i$. A random field is said to be a Markov field, if the probability P of any configuration of \mathbf{x} satisfies the positivity property:

$$P(\mathbf{x}) > 0 \quad \forall \mathbf{x} \in L \quad (1)$$

And the Markovian property:

$$P(x_i | \{x_j : j \in I \setminus \{i\}\}) = P(x_i | \{x_j : j \in N_i\}) \quad (2)$$

Or in other words the probability of x_i given the index set $I \setminus \{i\}$ is the same as the probability given the neighbourhood of i . Our use of neighbourhood is limited to the direct 6-neighbours in the volume. The goal is to compute the configuration of the field that maximizes the probability.

3.2 Defining the Markov Random Field

We aim to produce a labelling \mathbf{x} of the voxel volume V , such that the boundary of the labelling coincides with the skin surface. A likelihood term and two priors are defined on the following:

- Under scan surface is likely to be under skin surface (*under*), while above scan surface is likely to be above skin surface (*above*).
- Skin surface points are in the vicinity of scan points.
- Skin surface is locally smooth.

Here *surface points* and *scan points* are defined as the local interface between *above* and *under* labelled voxels. Using these priors a MRF is created for which a minimal energy problem is defined using the following:

Likelihood term: This term is based on the relative position of a voxel and the scan surface:

$$\Phi(v_i | x_i) = -\log P(v_i | x_i) \quad (3)$$

Generally, we expect the scan and skin surfaces to coincide and therefore make a simplification of the likelihood function, such that voxels below the scan surface S_{scan} have low energy if labelled *under* and high energy if labelled *above* and vice versa for voxels above the scan surface. The energy then becomes:

$$\begin{aligned}\Phi(v_i < S_{\text{scan}}|x_i) &= \begin{cases} 1 & x_i = \textit{under} \\ 0 & x_i = \textit{above} \end{cases} \\ \Phi(v_i > S_{\text{scan}}|x_i) &= \begin{cases} 0 & x_i = \textit{under} \\ 1 & x_i = \textit{above} \end{cases}\end{aligned}\quad (4)$$

This somewhat loosely defined term alone would just produce the scan surface.

Vicinity prior: The skin surface should be close to scan surface points. This is induced by adding an energy penalty to changes in label, which relates to the distance from the voxel v_i to the nearest scan point S_{scan} :

$$\lambda(x_i, x_j) = \begin{cases} K_{\text{dist}} \cdot \text{dist}(v_i, S_{\text{scan}}) & x_i \neq x_j \\ 0 & x_i = x_j \end{cases}\quad (5)$$

The distance is approximated using an Euclidean distance transform (EDT) [5]. This is a fast linear time algorithm that approximates distance in a number of sweeps. This penalty should only affect voxels that are not in the immediate neighbourhood of scan surface points, which is why 1 is subtracted from the distances, such that both an actual surface point and its direct neighbours have distance 0. The vicinity constraint mainly effects areas with high discontinuity but it also forces the resulting surface to be true to the data in areas with continuity.

Smoothness prior: Neighbouring voxels are expected to have the same label with higher probability than having different labels, therefore an energy penalty is given to adjacent voxels with different label:

$$\psi(x_i, x_j) = \begin{cases} K_{ij} & x_i \neq x_j \\ 0 & x_i = x_j \end{cases}\quad (6)$$

This is a general smoothness constraint.

Combining the likelihood term with the vicinity and smoothness prior, we get the following energy minimization problem for the scan volume:

$$E(\mathbf{x}) = \sum_{i \in I} \left(\Phi(v_i|x_i) + \sum_{j \in N_i} (\lambda(x_i, x_j) + \psi(x_i, x_j)) \right)\quad (7)$$

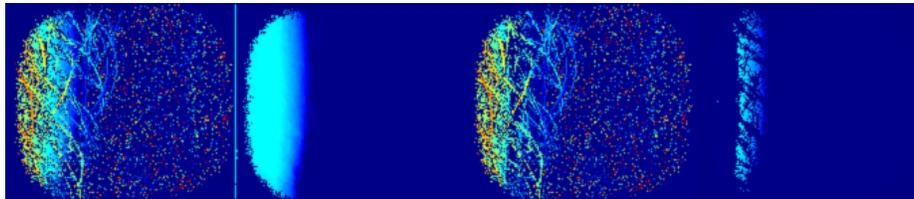
Where \mathbf{x} is the classification of the whole volume. The vicinity constant K_{dist} and the smoothness constant K_{ij} relates to each other and the likelihood term, which we defined to be 0 or 1. The solution to the *MRF* ensures maximum probability with the constraint that the labelling is both highly consistent with the data and smooth. To solve the minimization problem the *Graph Cut* algorithm [6] is used. This algorithm is an efficient way to find the optimal solution for such binary problems. From the resulting MRF classification a new surface S_{MRF} is extracted as the interface between the *under* and *above* labelled voxels.

4 Filtering based on the MRF solution

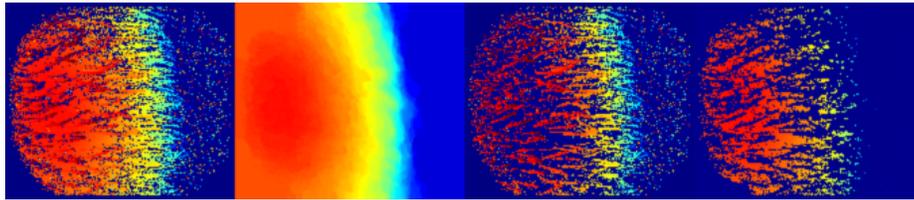
With a solution to the MRF and a new surface estimate S_{MRF} the changes compared to the original surface S_{scan} can be analyzed. As the MRF is set up so create a surface that coincides with the skin surface S_{skin} , we set up a filtering based on the following:

$$S_{\text{skin}} = S_{\text{scan}} \cap S_{\text{MRF}} \quad \wedge \quad S_{\text{hair}} = S_{\text{scan}} \setminus S_{\text{MRF}}$$

This results in a pure skin surface estimate and an estimate of the structured noise (hair) removed from the original scan to create the skin surface. We have used the strict definition but to give some flexibility one might add a threshold on how much is considered a change when comparing the original data with the MRF surface. Figure 2 shows examples of the resulting depthmaps of the *MRF* based filtering. Even though it is difficult to quantify the result, clearly both hair and noise are removed leaving only the smooth skin surface data.



(a) Very little skin with hair



(b) Half skin coverage with hair

Fig. 2. From left to right: depth map of original scan, surface based on volume classification, changes made in new surface, unchanged data in new surface. It is especially worth noting that the first scan is of very poor quality. The scanner is only focusing on very little skin surface and on top of that there is a lot of hair. In spite of this the algorithm returns the small amount of actual skin surface present in the scan.

Setting the relations between the likelihood term, the vicinity and smoothness prior is not trivial and based on *trial-and-error*. However, a reasonable approach is (at least for this type of data) to set the smoothness prior K_{ij} to 0, while incrementing the vicinity prior K_{dist} until a good result is achieved (Fig. 3(a)). This parameter effects areas with discontinuity such as skin to hair,

while it actually forces the algorithm to be true to the input data in areas with continuity such as skin to skin. Setting the parameter removes most of the hair strands leaving only a little stubble, which can then be removed by adding the smoothness constraint. Figure 3(b) shows the difference between surface data found using only the vicinity prior and using both vicinity and smoothness, the difference is seen as green, and the blue is the filtered surface data.

Even though the *MRF* solution actually creates a surface without holes, this surface tends to have bumps where the surface is closed below the hair strands. This is why only the unchanged part of the MRF surface should be kept. These remaining data are a much better starting point for a smooth surface reconstruction, as these data belong to the actual surface. Figure 3(c) shows a surface reconstruction based on the filtered point set. The surface has been reconstructed with the Markov Random Field Surface Reconstruction [9]. This algorithm uses a Markov Random Field to regularize a point distance field and it creates plausible hole filling, where data is missing. If computation time is a factor a simple 2D Delaunay triangulation [12] (considering the data as a 2D height map) would produce a reasonably fair surface.

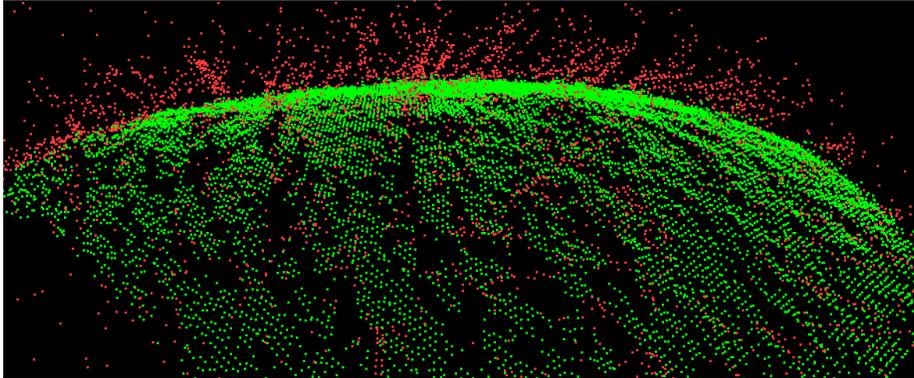
5 Conclusion

In this paper data from a novel surface scanner has been used in an approach to remove structured noise from raw scanner data. The approach is based on relabelling a voxel set using a Markov Random Field classification and extracting the sought surface as the interface between two voxel labels.

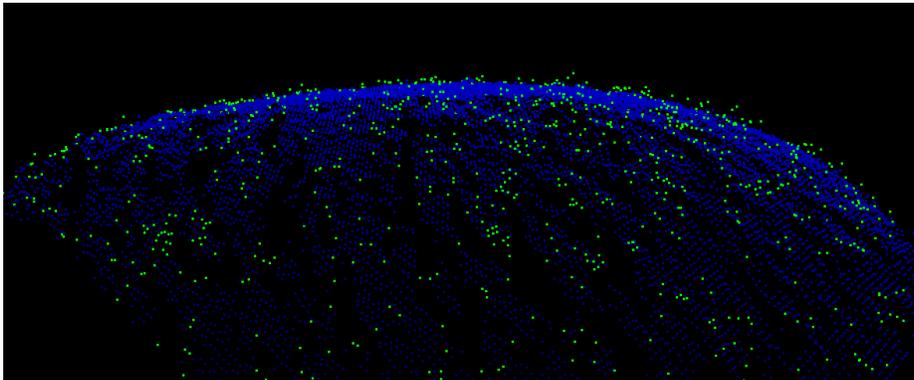
We have achieved good results on filtering noise and hair from our surface scans. Even though it is hard to quantify the ability to filter both noise and hair, qualitative visual inspections of the results are very promising. Because of the high quality expectation of the scans it is a strong point that the data is filtered such that the remaining data is unchanged.

Our algorithm has been implemented in Matlab with crucial parts in C++ .mex-files and it filters in a matter of a few seconds on a 2.8 GHz Intel processor laptop. There is good reason to believe that the surface filtering could be done in realtime with a full implementation in a precompiled programming language.

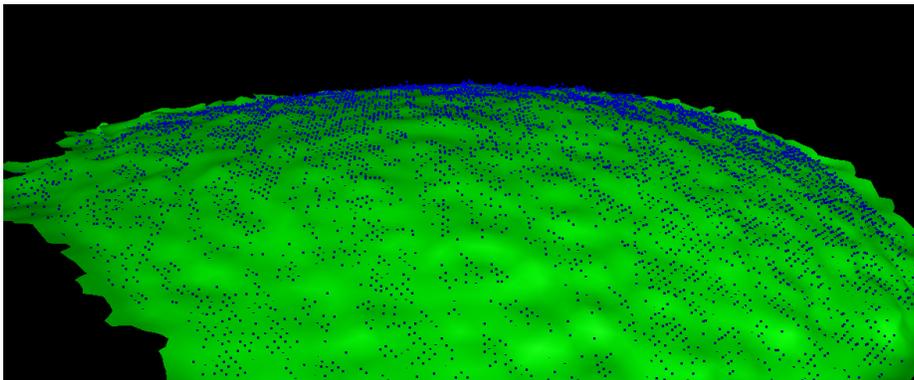
One could argue that a full volume classification as part of the filtering is not necessary. This is meant in the sense that large parts of the scan volume will be *far* from the actual surface and therefore a smarter way of classification close to the surface could speed up the process. This would however complicate the approach and is a topic for future research.



(a) Filtering done using only the vicinity prior



(b) Filtering after adding the smoothness prior



(c) Filtered data and surface reconstruction

Fig. 3. The figure shows a side view of a hairy surface. In the first image the effect of adding the vicinity prior alone is shown. Hair strands are coloured red and the surface green. The second image shows the difference between the surface using only vicinity prior and the surface, where both vicinity and smoothness prior has been applied. The difference is shown in green, while the resulting surface points of both vicinity and smoothness prior filtering are shown in blue. The third image shows a surface reconstruction using the filtered point set.

Bibliography

- [1] 3Shape. Trios scanner. Website, 2011. <http://www.3shapedental.com/restoration/dentist/digital-impression-taking.aspx>.
- [2] Julian Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, 48(3):259–302, 1986. ISSN 00359246.
- [3] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *Proceedings of ACM SIGGRAPH*, pages 303–312, 1996.
- [4] M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. ISSN 0001-0782.
- [5] M.W. Jones et al. 3D distance fields: A survey of techniques and applications. *IEEE Transactions on Visualization and Computer Graphics*, pages 581–599, 2006. ISSN 1077-2626.
- [6] V. Kolmogorov and R. Zabini. What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159, 2004. ISSN 01628828.
- [7] MAA Neil, R. Juskaitis, and T. Wilson. Method of obtaining optical sectioning by using structured light in a conventional microscope. *Optics Letters*, 22(24):1905–1907, 1997. ISSN 1539-4794.
- [8] R. R. Paulsen and R. Larsen. Anatomically plausible surface alignment and reconstruction, sep 2010. URL <http://www2.imm.dtu.dk/pubdb/p.php?5893>.
- [9] R.R. Paulsen, J.A. Baerentzen, and R. Larsen. Markov random field surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics*, pages 636–646, 2010. ISSN 1077-2626.
- [10] C.H. Reinsch. Smoothing by spline functions. *Numerische Mathematik*, 10(3):177–183, 1967. ISSN 0029-599X.
- [11] R. Schnabel, R. Wahl, and R. Klein. Efficient RANSAC for Point-Cloud Shape Detection. In *Computer Graphics Forum*, volume 26, pages 214–226. Wiley Online Library, 2007.
- [12] J. Shewchuk. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. *Applied Computational Geometry Towards Geometric Engineering*, pages 203–222, 1996.
- [13] A. Wehr and U. Lohr. Airborne laser scanning—an introduction and overview. *ISPRS Journal of Photogrammetry and Remote Sensing*, 54(2-3): 68–82, 1999. ISSN 0924-2716.

Approach-guided controlled resolution brain meshing for FE-based interactive neurosurgery simulation

Michel Audette¹, Denis Rivière², Matthew Ewend³, Sébastien Valette⁴

¹ Kitware

² Service Hospitalier Frdric Joliot

³ University of North Carolina, Dept. Neurosurgery

⁴ CREATIS; CNRS; Université de Lyon, France

Abstract. This paper presents an adaptation of a two-stage tetrahedralization, featuring boundary-faithful controlled-resolution surface and volumetric meshing, suitable for multi-grid finite-element based neurosurgery simulator, whose architecture reflects the choice of neurosurgical approach and the nature of the pathology. Moreover, the first application of our approach-guided methodology is a simulation of the pterional approach, which entails a brain mesh that is separable about the left or right Sylvian fissure, providing access to deep-seated lesions. In addition to this sulcal separability constraint, the emphasis of this procedure on lesions in the skull base also suggests that the meshing should model ventricles faithfully. However, another requirement that conflicts with the faithfulness to these sulcal and ventricular boundaries is the need to limit the number of tetrahedra, at the coarsest level, in order to produce an interactive yet materially and geometrically nonlinear response, which will be achieved by GPU-accelerated Total Lagrangian Explicit Dynamics. Finally, we view a controlled resolution as a cornerstone of a multi-grid approach to the simulation, which in turn we view as a necessity, given also the conflicting requirements of having to compute a deformation on a relatively large organ such as the brain, while also capturing finely detailed critical tissues, such as blood vessels, deleterious gestures to which must be penalized by the simulation.

1 Introduction: clinical need and conflicting requirements

CBTRUS has calculated a worldwide estimate of 186,678 newly diagnosed primary non-malignant brain and central nervous system tumors per annum for 2002 [4]. The five-year relative survival rate following diagnosis of a primary malignant brain tumor is 33.8% for males and 37.5% for females. It has been recognized that the main factor in long-term patient survival, other than the nature of the tumor, is completeness of the excision [5][22] underscoring the importance of surgical efficacy.

Surgical efficacy can be improved by simulation-based training, but despite advances in other areas and an important clinical need, simulation has not made major inroads into neurosurgery. Surgery simulation can be broadly categorized into predictive and interactive techniques, respectively used by expert surgeons for planning, through high-quality finite element (FE) computations, and by surgical residents for developing their skill, through interactive biomechanics and haptics. Correction of scoliosis of

the spine [1] and maxillofacial bone repositioning [13] exemplify predictive simulation. Meanwhile, validation studies have shown that residents having trained with an interactive simulator outperformed residents without such training in performing laparoscopic cholecystectomies [18]. The two types of simulation can be seen as poles of a spectrum characterized by the conflicting requirements of interactivity and faithfulness to tissue response, which research increasingly attempts to bridge through new efficiencies [10].

The clinical need that justifies pursuing interactive neurosurgery simulation with a strong predictive aspect lies in two tendencies in modern surgery: the compression of training schedules of surgical residents and the constant influx of new therapeutic technologies available to expert surgeons. Training schedule compression coincides with a recent movement to limit resident duty hours [14]. Interactive surgery simulation can provide a means for accelerating the training of residents, allowing them to take a more active role than in the traditional framework, whereby residents observe senior surgeons and gradually assume increased responsibility, and can result in measurable improvements to both skill and patient outcome. Meanwhile, the influx of new technology in neurosurgery is characterized by the increasing use of physics- and computer-based technology in modern practice [6]: stereotactic and minimally invasive neurosurgery, modalities such as ultrasound and digital atlases, as well as recent therapies such as the ultrasonic surgical aspirator, lasers, radiotherapy, and molecular therapies. The result is an armamentarium that is broad and technically complex, where simulation could play a pivotal role for expert surgeons assimilating or refining new treatments.

Our main objective is a neurosurgery simulator that can effectively model skull base surgery. Its architecture is organized according two nearly orthogonal parameters: i) the neurosurgical approach (e.g.: pterional, trans-nasal, and so on), which has implications for relevant critical tissues and the subvolume of clinical interest; and ii) the nature of the pathology, which dictates the choice of surgical tools and treatment. A requirement of skull base surgery simulation is depicted in figure 1 (a): one of the techniques available to the neurosurgeon for resecting deep-seated lesions involves parting the brain tissue about the Sylvian fissure. This technique characterizes pterional approach, which consists of approaching the brain and the anterior skull base through their lateral aspect after removing the frontal and temporal bones and the greater wing of the sphenoid [12]. The extensive spectrum of the neural and the vascular structures within reach includes insula, basal ganglia, lateral ventricle, middle cerebral artery, temporal operculum, frontal and parietal opercula, uncus, orbit, anterior cranial fossa, optic nerve, internal carotid artery and branches, lamina terminalis, and interpeduncular fossa.

This paper presents on-going research that features an adaptation of controlled-resolution clustering-based surface meshing, which is sulcal-separable if the surgical approach requires it of the anatomical model, followed by controlled-resolution tetrahedralization suitable for multi-grid FE [24] based neurosurgery simulator. Currently, the second stage is the Almost-regular Tetrahedralization method [9], which uses a mesh size objective consistent with that of the surface mesh. Moreover, simulation of the pterional approach entails a brain mesh that is separable about the left or right Sylvian fissure, providing access to deep-seated lesions. In addition to this sulcal separability constraint, the emphasis of this procedure on lesions in the skull base also suggests that the meshing should model ventricles faithfully. However, another requirement that con-

flicts with the faithfulness to these sulcal and ventricular boundaries is the need to limit the number of tetrahedra, to roughly 2000 or fewer, at the coarsest level, in order to produce an interactive yet materially and geometrically nonlinear response, which will be achieved by GPU-accelerated Total Lagrangian Explicit Dynamics [20][10]. We view a multi-grid approach to the simulation as a necessity, given the conflicting requirements of modeling a large organ such as the brain and finely detailed critical tissues, however the scope of this paper is restricted to the computation of the coarse-level mesh. Last, this model is used in conjunction with the SOFA platform, which implies that this tetrahedral mesh can be significantly coarser than the surface-based visual model, as long as the main structures relevant to the biomechanics, e.g. ventricles, are represented.

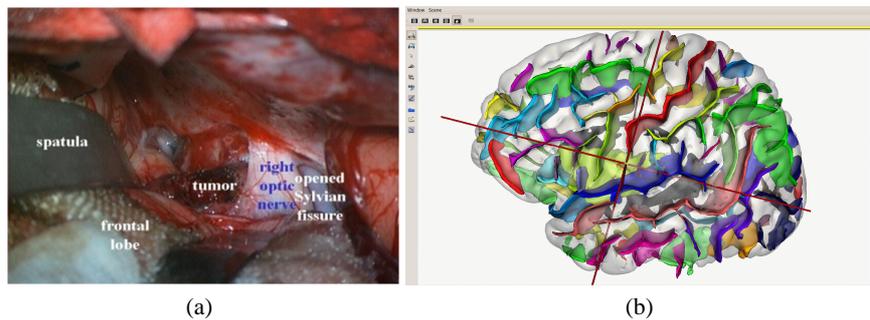


Fig. 1. Relevance and implementation of sulcal separability. (a) illustration of parting of Sylvian fissure to provide access to deep-seated tumor, and depiction of critical tissues; (b) illustration of BrainVISA results for labeling sulci.

2 Preprocessing: segmentation and sulcal labeling

The first stage of our current processing method to produce descriptive anatomical models of the brain is the image analysis pipeline that segments T1-weighted MR data according to white and grey matter (WM, GM) as well as corticospinal fluid (CSF), based on scale-space analysis, which is also the first step of the BrainVISA image analysis pipeline [3]. An alternate segmentation method is an atlas-based Expectation Maximization [15], if sulcal labeling is not required. With both methods, a separate tumor segmentation [16] will eventually be integrated for practical patient-specific neurosurgery simulation. The latter method identifies tumors based on their statistical outlying characteristic in T_1 - T_2 space, in conjunction with a probabilistic digital brain atlas that incorporates T_1 - T_2 features vectors of WM, GM and CSF, and with a level-set based region growing algorithm. In the BrainVISA pipeline, preliminary classification in terms of WM, GM and CSF is followed by morphological analysis to separate hemispheres and detect sulci on the basis of a white matter skeleton, which are subsequently expressed as nodes within relational graphs, which nodes coincide with simple surface

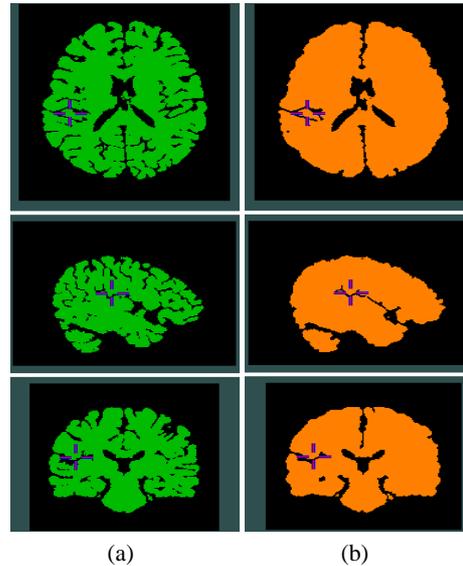


Fig. 2. Post-segmentation sulcus-aware morphological processing of brain binary image: (a) raw brain binary image; (b) result of eliminating irrelevant sulci, while emphasizing clinically relevant left Sylvian and hemisphere fissures, while preserving the ventricles.

primitives of various topological types [17]. The last computational stage of BrainVISA inputs this relational graph, within which each node is labeled according to its underlying surface type, to a multi-layer perceptron, which can deal with topological variability of sulci across individuals. The output of the sulcal labeling pipeline is shown in figure 1 (b), and is postprocessed to produce a suite of VTKPolydata files, which serves as input to the next stage, in addition to the brain tissue segmentation produced in the preprocessing.

Prior to the computation of triangular surface mesh, the sulci are used to focus the morphological filtering of the binary brain volume, so as to prevent clinically irrelevant sulci from appearing in the anatomical model, while also imposing separability where the clinically relevant sulcus appears, namely the left Sylvian fissure. The basic idea is to ensure that the mesh does not bridge the sulcus about which separability is required, while limiting the element count in the rest of the brain. This stage is implemented with a series of morphological operations on each hemisphere:

- a small morphological dilation of a voxel-based representation of all sulci;
- a small morphological closing of the brain binary volume OR'ed with the dilated irrelevant sulci;
- a small morphological opening of the closed brain binary volume from step ii on the basis of the clinically relevant sulcus;

- a blurring with a small kernel to produce brain hemispheres and subcortical structures from which stable results can be obtained by the application of Marching Cubes;
- a MAX operation of the blurred hemispheres and subcortical structure.

In addition, our experience suggests that rather than producing surface and volume meshes that are of constant mesh size, it is more feasible to achieve a mesh size objective that is smaller at surfaces of high curvature, such as the ventricles and the relevant Sylvian fissure. To this effect, the ventricles are isolated by morphological filtering, and a distance map is computed from them as well as from the relevant Sylvian fissure. Last, this surface model will also emphasize the space between hemispheres, and the extraction of this space is improved by exploiting prior labeling of left and right hemispheres produced by BrainVISA. This process is illustrated in figure 2.

3 Surface mesh simplification with topological guarantees

3.1 Motivation for a two-stage approach

We adopt a two-stage method for achieving a controlled-resolution, boundary-faithful mesh, in the spirit of the method developed by us in [2]. We have found by experience that a number of one-stage tetrahedralization methods tend to fuse two boundaries that are close together, which inevitably annihilates the Sylvian fissure imbedded in the mesh. The surface meshing stage begins with a Marching Cubes computation on the preprocessed binary brain volume of image in figure 2 (b). This Marching Cubes result is shown in figure 3.1 (b).

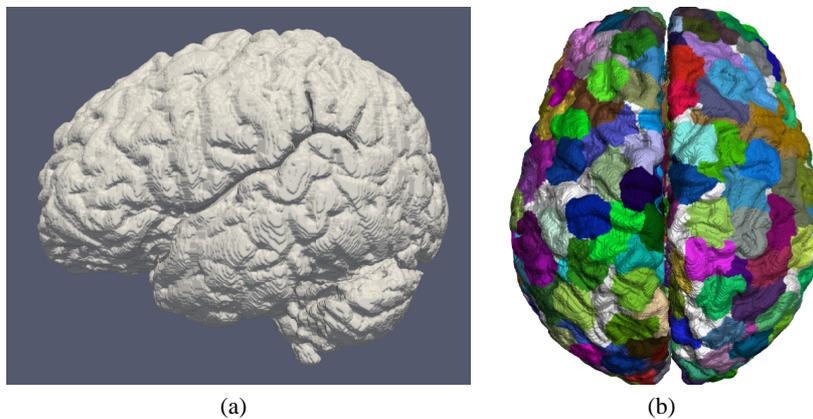


Fig. 3. Preliminary steps in surface meshing: (a) Marching Cubes; (b) first iteration of clustering.

3.2 Adaptive coarsening

The simplification of the surface resulting from the Marching Cubes algorithm is carried out using a modified version of the ACVD software. This algorithm, based on [23], constructs an approximate Centroidal Voronoi Tessellation (CVT) [8] on the input mesh by clustering its vertices v_j in n clusters C_i , in order to minimize the following energy:

$$E_{CVT} = \sum_{i=1}^n \left(\sum_{v_j \in C_i} \rho(v_j) \|v_j - c_i\|^2 \right) \quad (1)$$

where $\rho(v)$ is a density function defined for each vertex v_j , and allows adaptive clustering. In our experiments, we used a linear combination of the previously computed distance map to generate more vertices near the ventricles and the Sylvian fissure. Once the clustering is constructed, the set of cluster sites c_i defines a point set that can be easily triangulated by analysing the clusters connectivity. The resulting triangulation is a good candidate for numerical simulations.

Unfortunately, this algorithm does not guarantee that the resulting mesh will be a manifold surface, as some regions might locally require more clusters to be represented by a manifold surface in the simplified mesh. To overcome this issue, after the triangulation is constructed, we check that each simplified mesh vertex ring is equivalent to a topological disk. If some vertices do not respect this constraint, we insert one cluster near every failing vertex, and re-perform the clustering step. For this clustering step, we freeze all clusters except those that are in the 1-neighbourhood of the problem vertices. Freezing those clusters prevents the clusters to redistribute too much on the mesh, otherwise the clusters insertion would have little impact on the topology correction. This process is repeated until the resulting mesh is watertight. Since the input mesh is a watertight mesh, the convergence of this algorithm is guaranteed.

Note that even if the resulting mesh is watertight, its genus might be lower than the genus of the original surface. As an example when a small handle fits entirely in one cluster, it will be repaced by a single vertex in the simplified mesh. This allows our approach to filter out topological noise that would forbid good coarse simplifications.

Figure 4 shows a closeup view of a simplification example. The mesh shown in (a) contains relatively thin tubular regions where the clusters density is not big enough. The resulting triangulation is highly non-manifold, and vertices need to be inserted. After two insertion steps, the resulting mesh is manifold (e).

3.3 Further simplifications

Although the resulting mesh is watertight, it can happen that many vertices were inserted in thin regions, which can penalize the performance of the simulation toolchain. In order to further reduce the number of vertices in the mesh, we perform a topology-preserving simplification step, in spirit with [11]: we perform successive edge collapse operations on the mesh until the desired number of vertices is reached. Similarly to [11], the ordering of the edge collapses is computed using Quadric Error Metrics (QEM), and we forbid any edge collapse that would alter the mesh topology. Final results are shown in section 5.

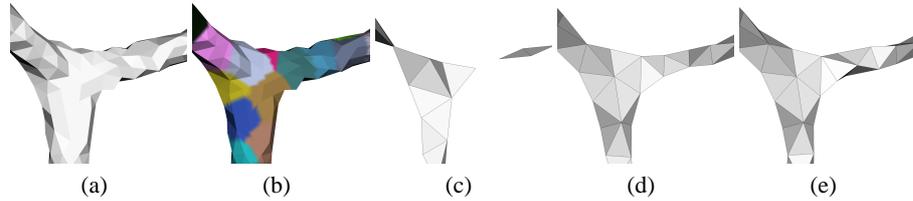


Fig. 4. Enforcing the correct topology. (a) : close-up view of a tubular input surface. (b) initial clustering. (c) resulting triangulation. This triangulation is not manifold, hence we introduce more clusters in the vicinity of non-manifold vertices. (d) after one iteration. (e) after two iterations, the result is manifold

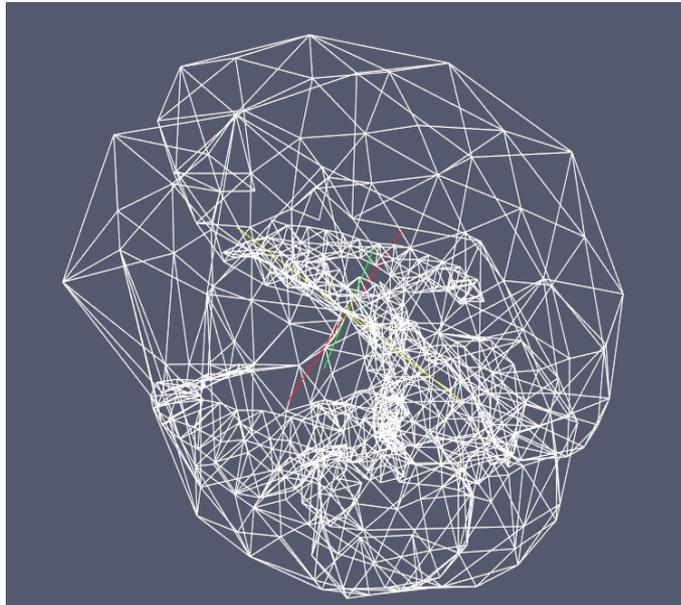
4 Surface-mesh constrained Tetrahedral Meshing

The tetrahedralization stage uses as input the controlled-resolution surface meshing results from the previous stage. Currently this stage is still underway. We are experimenting with open-source tetrahedralization software such as TetGen [21] that is faithful to boundaries and with the Almost-regular Tetrahedralization (ART) [9] used by us in past implementations where resolution control and boundary faithfulness were emphasized. The latter method admits an explicit mesh size objective, which is used to govern insertion of Steiner points. The final tetrahedral mesh is converted to an input format supported by the Simulation Open Framework Architecture (SOFA) platform [19], namely the VTK legacy file format.

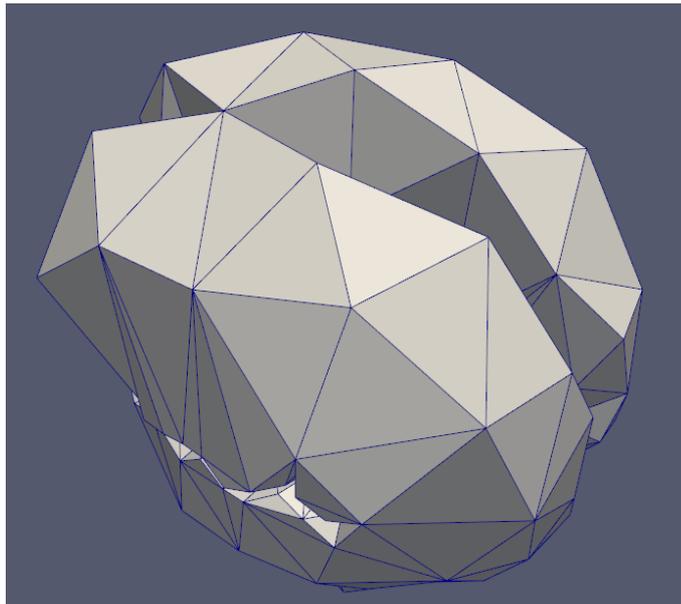
Preliminary results illustrating tissue parting on SOFA have so far been achieved with the Co-rotational FEM, which demonstrate small-deformation parting of the brain tissue about the left Sylvian fissure. The deformation tends to be localized, which motivates our on-going integration of haptic interactivity with Total Lagrangian Explicit Dynamics finite elements, which we feel will better resolve the large deformation required in this type of simulation.

5 Results

Figure 5 illustrates typical results of the surface meshing, whose wireframe rendering clearly shows the left Sylvian fissure and ventricles, as well as extremely coarse elements away from these structures. Figure 6 displays an exploded view of the tetrahedralization, and a preliminary mouse-based interaction via SOFA's Co-RotationalFEM biomechanics engine. These results also illustrate that a suitably sparse and coarse surface mesh can be achieved, while the development of the tetrahedral mesh that mirrors this edge size function is still work in progress. However we feel that the application of the ART method, which affords precise control over internal edge size, and the adaptation of the tetrahedral clustering method [7] in a manner that maintains the existing boundary, show promise of producing volumetric mesh with that adheres to the density function of the surface mesh.

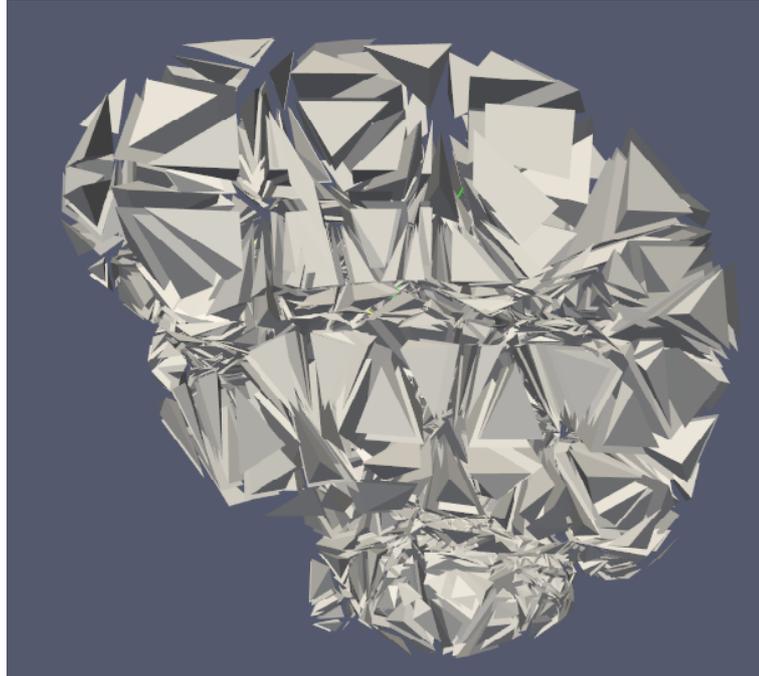


(a)

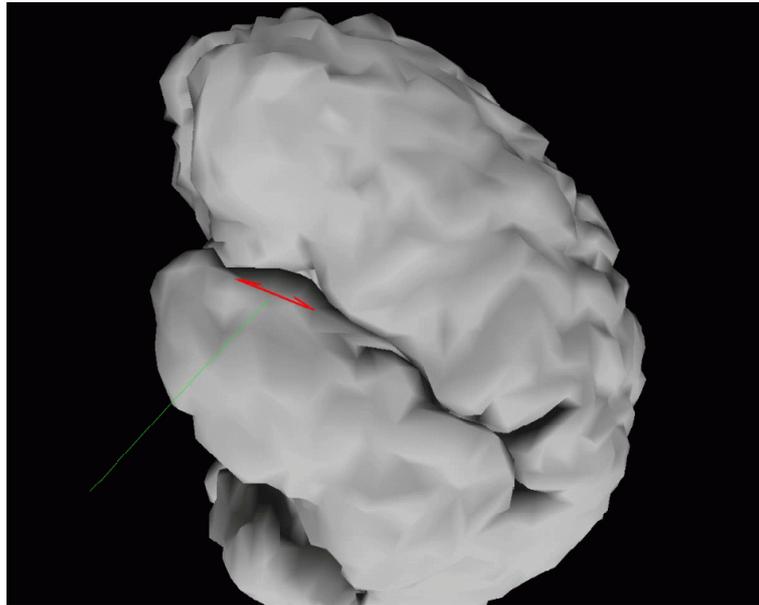


(b)

Fig. 5. Surface mesh simplification results: (a) wireframe and (b) surface and edge-based renderings. The mesh displayed here contains fewer than 600 elements.



(a)



(b)

Fig. 6. Volume meshing results: (a) exploded view and (b) preliminary interactive biomechanics results via SOFA co-rotational tetrahedral FEM class.

6 Conclusions

This paper presented a two-stage method for producing controlled-resolution brain mesh for neurosurgical simulation applications, where in particular the meshing strategy reflects an approach-specific adaptation. Moreover, we have shown that a mesh that is separable about a Sylvian fissure is both desirable and achievable for the specific case of the pterional approach. This method was restricted to the computation of the coarse-level representation of the brain in a multi-grid framework, but the same techniques can be applied to computing medium and fine representations of a subvolume of the brain.

References

1. Aubin, C.E., Petit, Y., Stokes, I.A., Poulin, F., Gardner-Morse, M., Labelle, H.: Biomechanical modeling of posterior instrumentation of the scoliotic spine. *Computer Methods in Biomechanics and Biomedical Engineering* 6(1), 27–32 (2003)
2. Audette, M., Delingette, H., Fuchs, A., Burgert, O., Chinzei, K.: A topologically faithful, tissue-guided, spatially varying meshing strategy for computing patient-specific head models for endoscopic pituitary surgery simulation. *Comput Aided Surg.* 12(1), 43–52 (2007)
3. BrainVISA: Brainvisa / anatomist webpage (June 2011), <http://brainvisa.info>
4. CBTRUS: Cbtrus fact sheet (April 2011), <http://www.cbtrus.org/factsheet/factsheet.html>
5. Chan, R.C., Thompson, G.B.: Morbidity, mortality, and quality of life following surgery for intracranial meningiomas a retrospective study in 257 cases. *J. Neurosurg.* 60(1), 52–60 (June 1984)
6. Ciure, A.V., Lencean, S.M., Brehar, F.M.: Actualities and perspectives in neurosurgery. *J Med Life* 1(1), 23–9 (Jan-Mar 2008)
7. Dardenne, J., Valette, S., Siauve, N., Prost, R.: Variational tetrahedral mesh generation from discrete volume data. *The Visual Computer (proceedings of CGI 2009)* 25(5), 401–410 (May 2009)
8. Du, Q., Faber, V., Gunzburger, M.: Centroidal voronoi tessellations: applications and algorithms. *SIAM Review* (41(4)) (1999)
9. Fuchs, A.: Automatic grid generation with almost regular delaunay tetrahedra. In: *Proc. 7th Int. Mesh. Round.* pp. 133–48 (1998)
10. G. R. Joldes, A. Wittek, K.M.: Real-time nonlinear finite element computations on gpu - application to neurosurgical simulation. *Computer Methods Appl. Mech. Engin.* 199(49-52), 3305–3314 (Dec 2010)
11. Garland, M., Heckbert, P.S.: Surface simplification using quadric error metrics. *Computer Graphics* 31(Annual Conference Series), 209–216 (1997)
12. Jr, A.R.: *Cranial Anatomy and Surgical Approaches.* Lippincott Williams and Wilkins (2008)
13. Luboz, V., Chabanas, M., Swider, P., Payan, Y.: Orbital and maxillofacial computer aided surgery: patient-specific finite element models to predict surgical outcomes. *Comput Methods Biomech Biomed Engin.* 8(4), 259–265 (Aug 2005)
14. Miulli, D.E., Valcore, J.C.: Methods and implications of limiting resident duty hours. *J Am Osteopath Assoc.* 110(7), 385–95 (2010)
15. Pohl, K., Bouix, S., Shenton, M., WE, W.G., Kikinis, R.: Automatic segmentation using non-rigid registration. In: *Med Image Comput Comput Assist Interv.* pp. 1201–1212 (2007)
16. Prastawa, M., Bullitt, E., Ho, S., Gerig, G.: A brain tumor segmentation framework based on outlier detection. *Med Image Anal.* 8(3), 275–283 (2004)

17. Riviere, D., Mangin, J., Papadopoulos-Orfanos, D., Martinez, J., Frouin, V., Regis, J.: Automatic recognition of cortical sulci of the human brain using a congregation of neural networks. *Med Image Anal.* 7(4), 403–16 (2003)
18. Seymour, N.E., Gallagher, A.G., Roman, S.A., O'Brien, M.K., Bansal, V.K., Andersen, D.K., Satava, R.M.: Virtual reality training improves operating room performance: results of a randomized, double-blinded study. *Ann Surg.* 236(4), 458–463 (Oct 2002)
19. SOFA: Sofa simulation open framework architecture webpage (June 2011), <http://www.sofa-framework.org/>
20. Taylor, Z., Cheng, M., Ourselin, S.: Factors affecting morbidity and mortality following surgical intervention in patients with intracranial meningioma. *IEEE T Med Imaging* 27(5), 650–663 (2008)
21. TetGen: Tetgen - a quality tetrahedral mesh generator and a 3d delaunay triangulator (June 2011), <http://tetgen.berlios.de/>
22. Turgut, M., Ozcan, O.E., Benli, K., Ozgen, T., G urcay, O., Bertan, V., Erbenli, A., Saglam, S.: Factors affecting morbidity and mortality following surgical intervention in patients with intracranial meningioma. *Aust. N.Z. J Surg.* 66(3), 144–50 (Mar 1996)
23. Valette, S., Chassery, J., Prost, R.: Generic remeshing of 3d triangular meshes with metric-dependent discrete voronoi diagrams. *IEEE Trans Visu Comp Grap* 14(2), 369–381 (2008), <http://hal.archives-ouvertes.fr/hal-00537025>
24. Wu, X., Tendick, F.: Multigrid integration for interactive deformable body simulation. In: *Medical Simulation: International Symposium, ISMS 2004*. pp. 92–104 (2004)