Towards Patient-specific Anatomical Model Generation for Finite Element-based Surgical Simulation

M.A. Audette, A. Fuchs, O. Astley, Y. Koseki and K. Chinzei

National Institute of Advanced Industrial Science and Technology- AIST,
Surgical Assist Technology Group
m.audette@aist.go.jp,
http://unit.aist.go.jp/humanbiomed/surgical/
Namiki 1-2, Tsukuba, 305-8564, Japan

Abstract. This paper presents ongoing research on a semi-automatic method for computing, from CT and MR data, patient-specific anatomical models used in surgical simulation. Surgical simulation is a software implementation enabling a user to interact, through virtual surgical tools, with an anatomical model representative of relevant tissues and endowed with realistic constitutive properties. Up to now, surgical simulators have generally been characterized by their reliance on a generic anatomical model, typically obtained at the cost of extensive user interaction, and by biomechanical computations based on mass-spring networks.
We propose a minimally supervised procedure for extracting from a set of CT and MR scans a highly descriptive tissue classification, a set of triangulated surfaces coinciding with relevant tissue boundaries, and volumetric meshes bounded by these surfaces and comprised of tetrahedral elements of homogeneous tissue. In this manner, a series of models could be obtained with little user interaction, allowing surgeons to be trained on a large set of pathologies which are clinically representative of those they are likely to encounter. The application of this procedure to the simulation of pituitary surgery is described. Furthermore, the resolution of the surface and tissue meshes is explicitly controllable with a few simple parameters. In turn, the target mesh resolution can be expressed as a radially varying function from a central point, in this case coinciding with a point on the pituitary gland.
A further objective is to produce anatomical models which can interact with a published finite element-based biomechanical simulation technique which partions the volume into separate parent and child meshes: the former sparse and linearly elastic; the latter dense, centered on the region of clinical interest and possibly nonlinearly elastic.

1 Introduction

Surgical simulation is a software implementation enabling a user to interact, through virtual surgical tools, with an anatomical model representative of relevant tissues and endowed with realistic constitutive properties. The anatomical

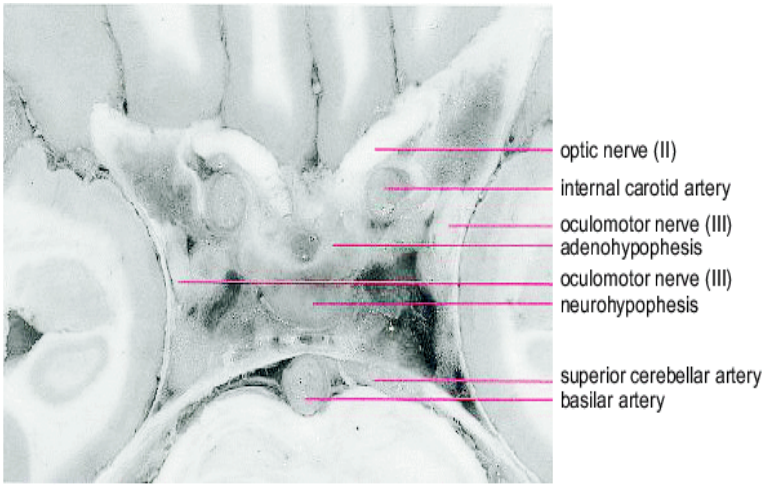

Fig. 1. Illustration of anatomy relevant to pituitary surgery: pituitary gland and surrounding critical tissues (reproduced with permission from [8]).

model must feature tissues whose consideration is essential to the clinical situation, such as critical vasculature and nerves, so as to appropriately simulate and penalize any damage to them. This is particularly true in our clinical application, the simulation of transnasal pituitary surgery [8][10], where the target is surrounded by the optic and oculomotor nerves and by cranial arteries (see figure 1).

Up to now, virtual anatomical models for surgical simulation have been highly task-specific and obtained by elaborate computer-user interaction, including segmentation and meshing. This paper presents on-going research on a minimally supervised procedure for extracting from a set of MR and CT scans a highly descriptive anatomical model, leading not merely to one generic model, but to a family of patient-specific models. This procedure integrates a distortion-tolerant mutual information MR-CT registration, a new tissue classification exploiting global spatial cues, a simplex-based surface meshing model to identify and triangulate the relevant anatomical boundaries, and an automatic almost-regular tetrahedralization of tissue volumes surrounded by these boundaries. This procedure is designed to be extensible to other surgical applications.

Lastly, a further objective of our research is to apply a new finite element (FE) software architecture, proposed by Astley [1][2] and specifically designed for surgical simulation, to a real clinical problem. As shown on figure 2 (a)-(c), this architecture partitions the underlying volume into one or more dense child meshes and a sparse parent mesh. The static FE equation is represented as follows:

$$\mathbf{Ka} - \mathbf{f} = 0 \ , \tag{1}$$

where $\mathbf{K}$ is the *stiffness* matrix, $\mathbf{a}$ is the vector of *node displacements*, and $\mathbf{f}$ is the set of *node forces*. Each represents the assemblage of *elemental* stiffness matrices $\mathbf{K}^e$, displacements $\mathbf{a}^e$ or forces $\mathbf{f}^e$ [28]. The set of forces $\mathbf{f}^e$ include a concentrated loads term which accounts for user-controlled virtual cutting forces [18] [9]. This FE architecture considers each node shared by the parent and child meshes: for each subregion, parent or child, it expresses the other subregion(s) encountered at each node as *one* equivalent impedance and force. A large stiffness matrix is then reduced to $n+1$ decoupled, significantly smaller stiffness matrices, where n is the number of child meshes. Each mesh can be resolved independently and at different rates over time, but the child system(s) surrounding the surgical tool(s) must be solved at haptic rates, typically of the order of 500 Hz [7]. Decoupling the problem naturally leads to parallelization on $n+1$ processors, which can make optimal use of even a dual-processor Pentium computer, particularly if the haptic [19] and visual [21] rendering can be handled by peripheral hardware. Finally, the FE method can be extended to nonlinearly elastic models [22] [5] and the Astley architecture allows for nonlinearly elastic child meshes and has been demonstrated at haptic rates on non-anatomical geometries, justifying its selection over methods that are constitutively limited or that require extensive precomputation [3].

This perspective imposes on our surface and volumetric meshing stages a requirement of explicit control over mesh resolution. Control is exercised with a smooth, radially varying mesh scale function defined from a user-provided central point (e.g.: on the pituitary gland), which naturally leads to a conformal mesh composed of a dense child and a sparse parent. It should also be emphasized that our procedure identifies anatomical surfaces prior to volumetric meshing, rather than proceed directly from the classification to tissue-guided volumetric meshing, because in general the latter approach will not produce a mesh boundary which is smooth and which closely agrees with the anatomical boundary. From a haptic and visual rendering standpoint, an anatomical model with jagged and badly localized boundaries would detract from the realism and clinical relevance of a surgical simulator.

2 Materials and Methods

2.1 Locally Weighted Mutual Information Registration

Pituitary surgery imposes requirements on the surgeon of a highly accurate trajectory through both bone and soft tissue, which entails their accurate resolution near the pituitary gland [8]. To do so, we must co-register CT and MR volumes, and resample MR data in a manner compatible with CT sampling. However, this registration is complicated by MR distortion caused by magnetic susceptibility variations near the pituitary gland and sinus cavities [25]. The prevalent co-registration technique for MR and CT is the global mutual information (MI) procedure [27]. It registers two volumes in manner which iteratively minimizes the dispersion of their joint intensity distribution, is widely used and provides

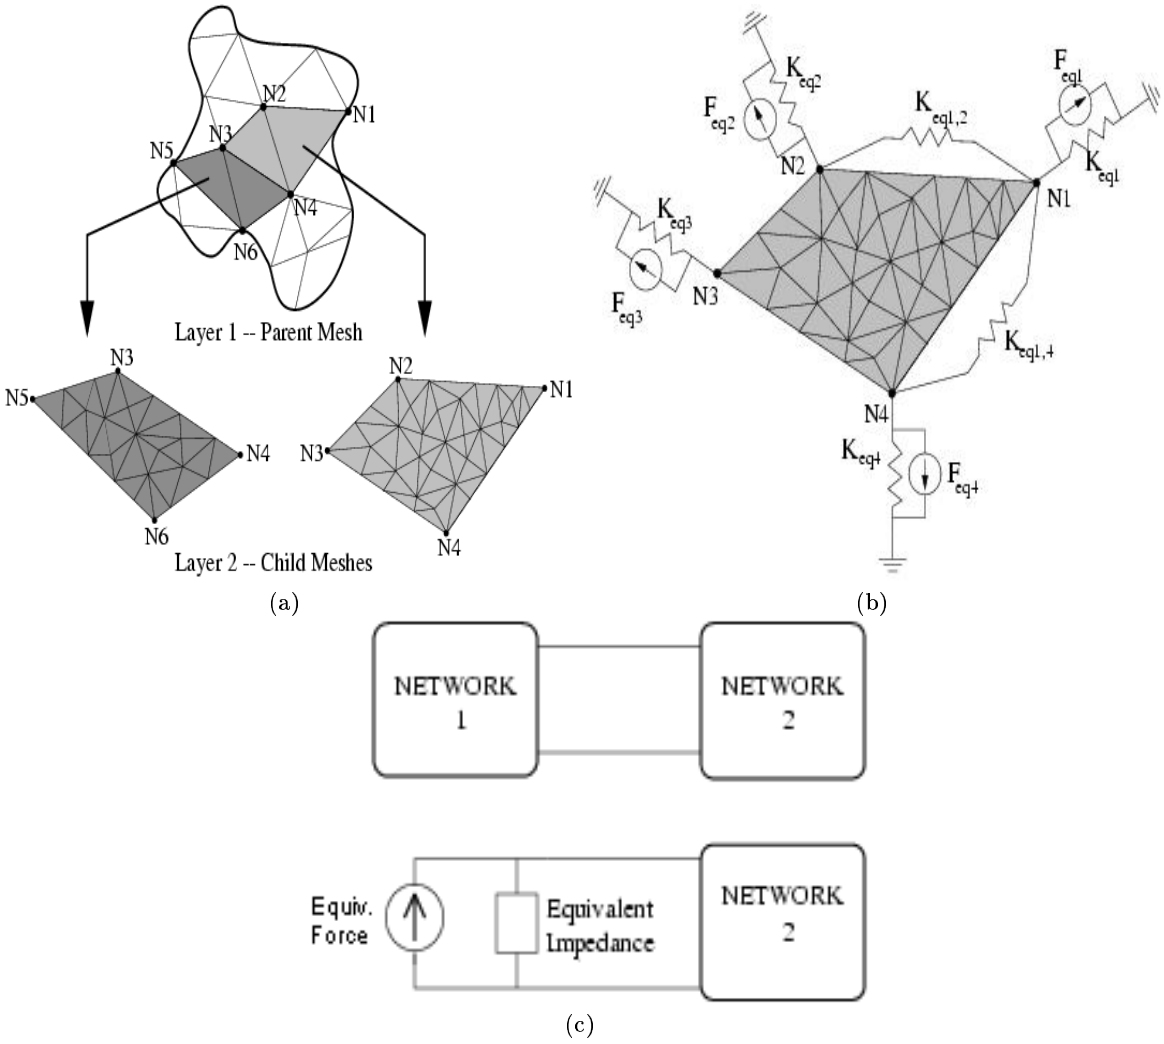

Fig. 2. Illustration of biomechanical modeling principle. (a) Partition of FE domain into parent and one or more child meshes. Subsequently, either subregion, parent or child, can be represented in the stiffness matrix of the other as an equivalent based on the Norton equivalent in circuit analysis. Each node shared by parent and child is considered (b), and, in computing the decoupled stiffness matrix of the child mesh, the parent mesh as seen by this node can be expressed as force and stiffness equivalents (reproduced with permission from [2]).

globally accurate results. However, this method suffers from misregistration precisely where the distortion is most pronounced, and where the greatest accuracy is required by our clinical application. Our registration technique, which addresses this shortcoming, is comprised of the two following stages:

1. a standard *global 7-parameter MI procedure* (rigid + scale), followed by
2. a *locally weighted linear MI procedure* based exclusively on information "near" the anatomy of interest, using a spherical mask easily specified by the user.

2.2 Global Structure-preserving Voxel Classification

Voxel Classification is a mapping of feature vectors, typically comprised of tomographic modalities such as CT or MRI, to a discrete set of tissue classes. Up to now, virtual anatomical models for surgical simulation have been obtained by highly elaborate manual segmentation. Ultimately, the tissue classification on which the anatomical model is based must account for clinically relevant tissues, including critical tissues such as vasculature and nerves, while also fulfilling our objective to produce a series of such models with little supervision. These conflicting requirements lead to the necessary consideration of a priori anatomical information in the classification, as feature space alone does not provide enough information to discriminate between all classes relevant to the simulation. Recent techniques [16] [20] exploit essentially local spatial constraints on tissue classification. In constrast, we exploit the *global spatial structure of tissues*, whose recruitment in disambiguating two or more clinically relevant tissue types, overlapping in feature space, has been neglected up to now in the literature.

Our classification method, somewhat inspired from the semi-supervised Fuzzy C-Means technique [6], begins with a Minimum Distance (MD) classification from a small training set, and then tries either to consolidate, or to invalidate and recompute, membership on the basis of global spatial constraints. We integrate the Fast Marching (FM) [24] method with our minimally supervised, iterative classifier. Constraints can easily be placed on the FM front propagation to implement assumptions about global structure. These assumptions include (for more details see [4]):

1. *Contiguity with training points* [1]: given an initial classification, the contiguity of voxels $\mathbf{x}_k$ with the training points of class C_i can be enforced by an outward front from these points, propagating only on voxels of class C_i.
2. *Prior assumptions about spatial extent*: the FM method described in 1 can be restricted by a scale factor, preventing the front propagation to continue beyond this limit; e.g.: pathology, cranial nerve and vasculature voxels can be assumed to be close to their training points, if the latter are well chosen.
3. *Embedded structure*: a very useful cue for classification relates to how contiguous tissues fit within each other, particularly in relation to a tissue class

[1] For classes where this constraint is useful, training points should be chosen not only for their intensity, but for their spatial relation to tissues. We consider this added burden on the user to be manageable, since the training set remains small.

which can be identified reliably such as bone; e.g.: grey and white matter occur inside, while muscle and fat are found outside, the cranium.

4. *Similarity and proximity to confidently classified voxels*: if there is still ambiguity, feature similarity and spatial proximity to the boundary Γ_i of a set of contiguous blobs of confidently classified voxels $\tilde{\mathbf{x}}_{k,i}$ can be exploited.

The philosophy of this method is not to favour any single class over all others, but merely to discount from consideration any class which at a given position has essentially zero likelihood, based on prior information.

2.3 Simplex-based Tissue-guided Surface Meshing with Resolution Control

Once a descriptive tissue classification is computed, our next step is to establish the triangulated tissue boundaries relevant to the simulation by exploiting this classification, a discrete surface model, and the position of the training points of each class. The *n-simplex* mesh is a discrete active model [12], characterized by each vertex being linked to each of $n+1$ neighbours by an edge. A balloon force can act on this mesh to cause it to expand until some image-based force halts this expansion. A surface model in 3D is realized as a 2-simplex, characterized by each vertex having 3 neighbours, and this representation is the dual of a triangulation, with each simplex vertex coinciding with a center, and each simplex edge being bisected by an edge, of a triangle. Furthermore, this surface model also features other internal forces [12] which nudge each simplex face, and consequently each dual triangle, towards having edges of equal length, and towards $\mathcal{C}_0$, $\mathcal{C}_1$ or $\mathcal{C}_2$ continuity, for example.

We apply the 2-simplex surface model to our anatomical meshing problem because of the explicit control on mesh characteristics which can be achieved, many of which are already implemented [12][13]. Each simplex is initialized with a small spherical mesh centered on a tissue class training point, and expanded until halted by image information, in the form of a set of tissue classes assumed "outside" the volume of interest. A few adaptations to the simplex model are proposed here for automatic FE mesh generation. First, we can endow the surface model with absolute mesh size thresholds, expressed as a target simplex area scale A_s and a percent tolerance ε_{A_s}: a simplex face falling below the area minimum $A_s - \varepsilon_{A_s}/100$ results in the fusion of two contiguous faces into one by eliminating a shared edge [12], while a face above the maximum area $A_s + \varepsilon_{A_s}/100$ is subdivided into two by a new edge. This *constant* factor can be replaced with a *radially varying function* $A_s(\mathbf{x})$, determined by the distance $R(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}_c\|$ between the centroid $\mathbf{x}$ of each simplex face and a user-defined central point $\mathbf{x}_c$, e.g.: inside the pituitary gland:

$$A_s(\mathbf{x}) = \begin{cases} A_{s,min} & \text{if } R(\mathbf{x}) <= R_{child} \\ A_{s,min} + (A_{s,max} - A_{s,min}) \left\{ 1 - \exp\left[\frac{-(R(\mathbf{x}) - R_{child})}{R_{scale}} \right] \right\} & \text{otherwise,} \end{cases} \tag{2}$$

where $A_{s,min}$ and $A_{s,max}$ specify smallest and largest target areas, and R_{child} and R_{scale} determine the behaviour of the function bridging the two values: a small, constant scale for a radius less than R_{child} and an exponential function tending towards $A_{s,max}$ as the distance of the simplex face to the central point increases. This scale function thereby produces small mesh faces near the pituitary gland and larger faces away from it. The triangulated surface coinciding with the anatomical boundary, obtained by duality with the final simplex mesh, is the objective of this stage. More than one children are possible, at the cost of selecting other "central" points of interest $\mathbf{x}_{c,i}$, in which case we define $R(\mathbf{x}) = \min_i \|\mathbf{x} - \mathbf{x}_{c,i}\|$.

Finally, a number of other adaptations are underway, including:

- A *conformality-preserving force*, which would cause two contiguous boundaries to share vertices wherever desirable, and
- *Topological adaptivity* applied to the 3D simplex model (previously demonstrated for 2D simplex models [13]), which would allow two or more spheres emerging from neighbouring training points to fuse together upon contact, or conversely would allow a simplex boundary to break apart over large gaps. In particular, a promising approach is that of Lachaud [17], which sets an inner bound on edge length and detects topological changes on the basis of an inter-vertex distance falling under this bound.

2.4 Almost-regular Volumetric Meshing with Radial Resolution Control

The last stage in our procedure partitions each volume bounded by a triangulated mesh, coinciding with a relevant tissue class, into tetrahedral elements consistent with the FE method. The volumetric meshing stage is essentially a published technique [15] which automatically produces an optimal tetrahedralization from a given polygonal boundary, such as a triangulated surface. In this case, optimality is defined as near-equal length of the tetrahedral edges, along with a sharing of each inner vertex by a nearly consistent number of edges and tetrahedra. This method features the optimal positioning of inner vertices, expressed as a minimization of a penalty functional, followed by a Delaunay tetrahedralization. The resulting near-regularity is important for FE stability and efficiency [23].

We modify this technique by integrating into the penalty functional the now-familiar radially varying scale function, which is specified as a target edge length $L_t(\mathbf{x})$ for each tetrahedron. Based on the relationship between the number of simplex and triangle vertices $V_t \approx V_s/2$ [12], a target simplex mesh size of A_s works out to a triangular area of $A_t \approx A_s/2$, and to the following triangular and tetrahedral target edge length (assuming edges of near-equal lengths):

$$L_t(\mathbf{x}) \approx \sqrt{[2A_s(\mathbf{x})/\sqrt{3}]}\,. \tag{3}$$

The separation of the resulting tetrahedral mesh into child and parent is as follows: contiguous tetrahedra whose edge lengths approach $L_{t,min} = \sqrt{[2A_{s,min}/\sqrt{3}]}$ comprise the child mesh, while the other elements constitute the parent.

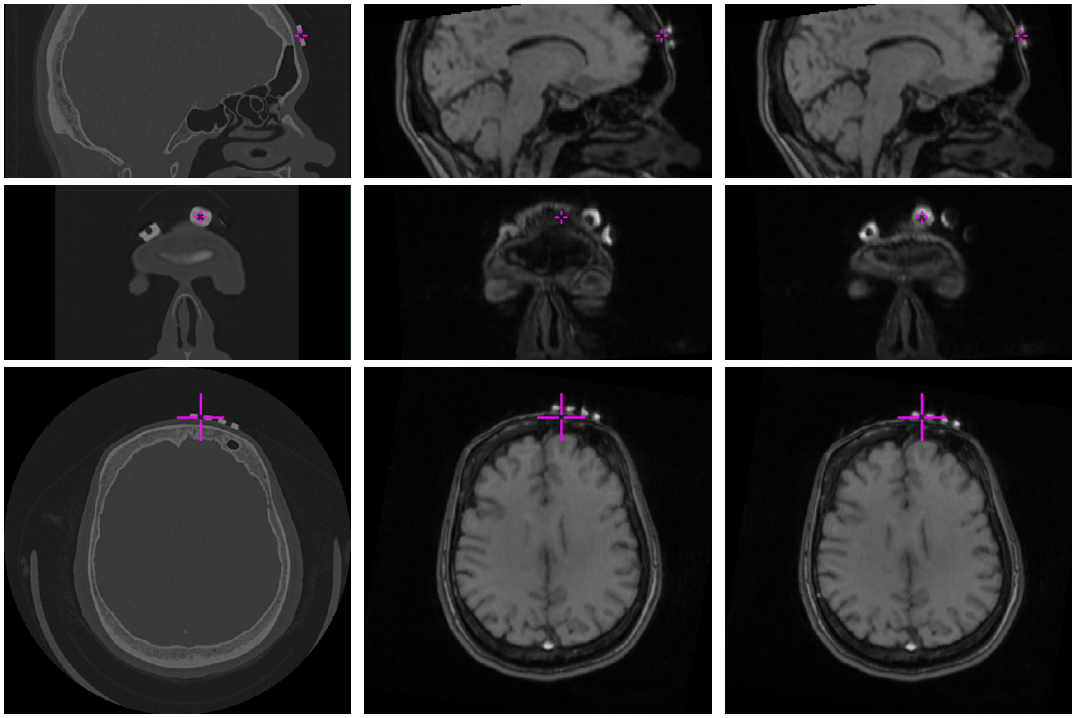

(a) (b) (c)

Fig. 3. Illustration of local weighting: (a) CT scan with manually identified fiducial; (b) MR scan transformed to CT space based on global MI registration, with CT fiducial overlaid; (c) transformed MR data as in (b), using the locally weighted method.

3 Results

Ongoing validation for each stage is based on the application of the procedure to real patient data obtained in collaboration with the Tokyo Women's Hospital. Validation of the registration exploits fiducial registration error (FRE) statistics, based on manually identified fiducials in both spaces, and is documented in [4]. A study evaluating the relative merits of different linear transformations is currently under way and will be published shortly. We use either a small set of homologous anatomical landmark pairs or a principal axis transform to provide the global MI stage with an initial transformation, with the former proving to be slightly less accurate than the latter under ideal circumstances, but more robust to artefacts and non-overlapping information, such as the presence of a headrest in one scanner but not in the other. The size of the spherical mask is typically 100 mm. An illustration of typical results is shown in figure 3.

Classification validation, while currently qualitative, provides a stark justification for incorporating global spatial constraints, given the number of classes which can be discriminated, in comparison to what is achievable otherwise, as

illustrated in figure 4. Future validation will make use of a digital anthropomorphic phantom as well as CT and MR simulators.

The validation of both the surface and volumetric meshing are also based on qualitative studies with patient data, as shown in figure 5, and in the future will exploit synthetic anthropomorphic data.

4 Conclusion and Future Directions

This paper presented results of ongoing research on a semi-automatic procedure for computing anatomical models for patient-specific surgical simulation. The procedure features distortion-tolerant MI registration of MR and CT, classification which exploits the global structure of tissues, simplex-based tissue-guided surface meshing, and automatic almost-regular volumetric meshing, with explicit resolution control on both meshing techniques. We believe that the near-regularity of the resulting mesh has positive bearing on the stability and efficiency of the FE method [23]. It should be emphasized that this meshing strategy will produce a mesh of radial density but heterogeneous material properties, resulting from the consideration of triangulated boundaries of different tissues. In the event that a particular tissue is essential to the simulation but is impractical to tessellate at the scale target computed for it, it may become imperative for our target to be relaxed somewhat. Moreover, some tubular structures may benefit from a modeling as curvilinear, rather than volumetric, elements, for the sake of computational performance. Ultimately, forthcoming experimentation with the multirate FE architecture of [2] will validate our meshing strategy, or suggest modifications to it.

In our tissue classification, there is currently no special treatment of critical tissues, but given their importance to the simulation, we recognize that our framework must be refined to consider them as a separate case. In research currently underway, user-provided "training points" from blood vessels and cranial nerves are being used to anchor a minimal path (MP) through them [11][14], and a tissue classification which proceeds outwards from these minimal paths. In contrast with existing MP techniques, which exploit only image *gradient magnitude* information and may be undermined by low-contrast areas, our MP computation makes use of a robust image feature for detecting tubular structures. This feature, which takes *gradient direction* into consideration, is the *image gradient flux* [26], which distinguishes between *sources* and *sinks* of a gradient vector field, coinciding with points where the *outward flux* of this field is *positive* and *negative* respectively. For visible vessels and nerves, tubular structures of higher intensity than their surroundings, points along their central axis coincide with strong negative outward flux, so that a minimal path along this central axis can be defined based on a flux-weighted potential. The subsequent critical tissue classification can make use of the sign of image gradient flux, as well as the proximity of a voxel to a particular minimal path.

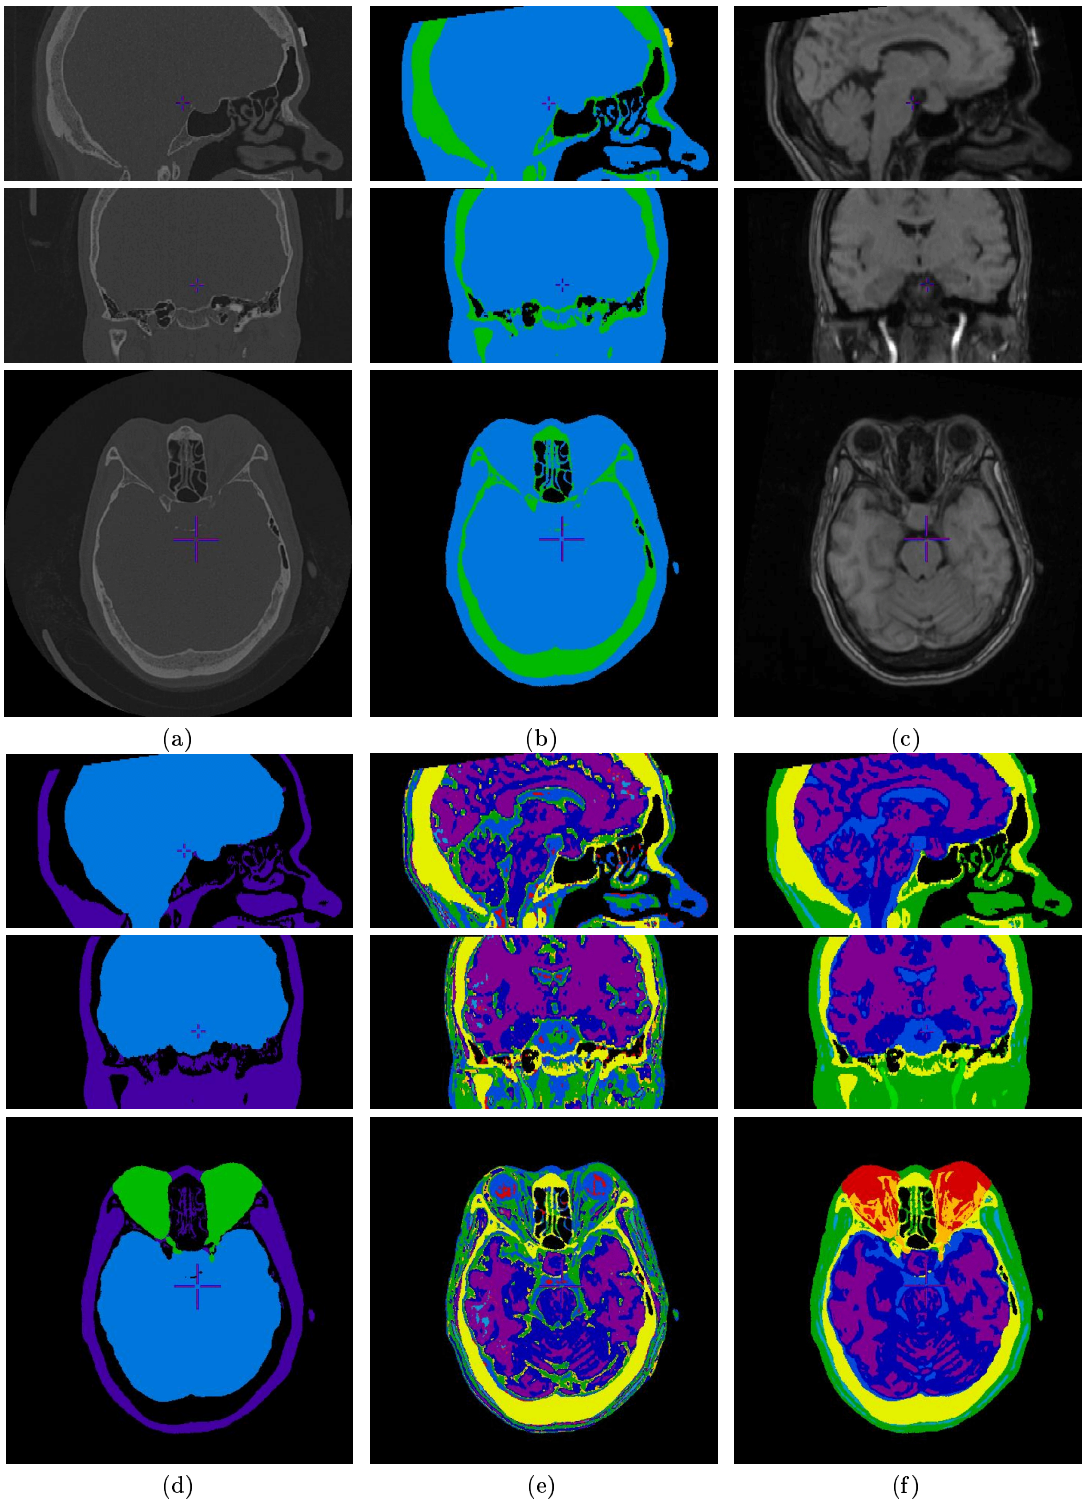

Fig. 4. Illustration of improvement due to global spatial constraints: (a) CT data; (b) CT-based classification of hard and soft tissues, via constraints 1, 2 and 3; (c) MR data; (d) embedded intracranial and intraorbital regions (constraint 3); classification of 9 soft tissues (e) without, and (f) with, constraints 1, 2, 3 and 4. In (e) there are many corticospinal fluid (royal blue) and muscle (green) false positives, as well as many vasculature (bright green) and ocular tissue (orange and red) false negatives.

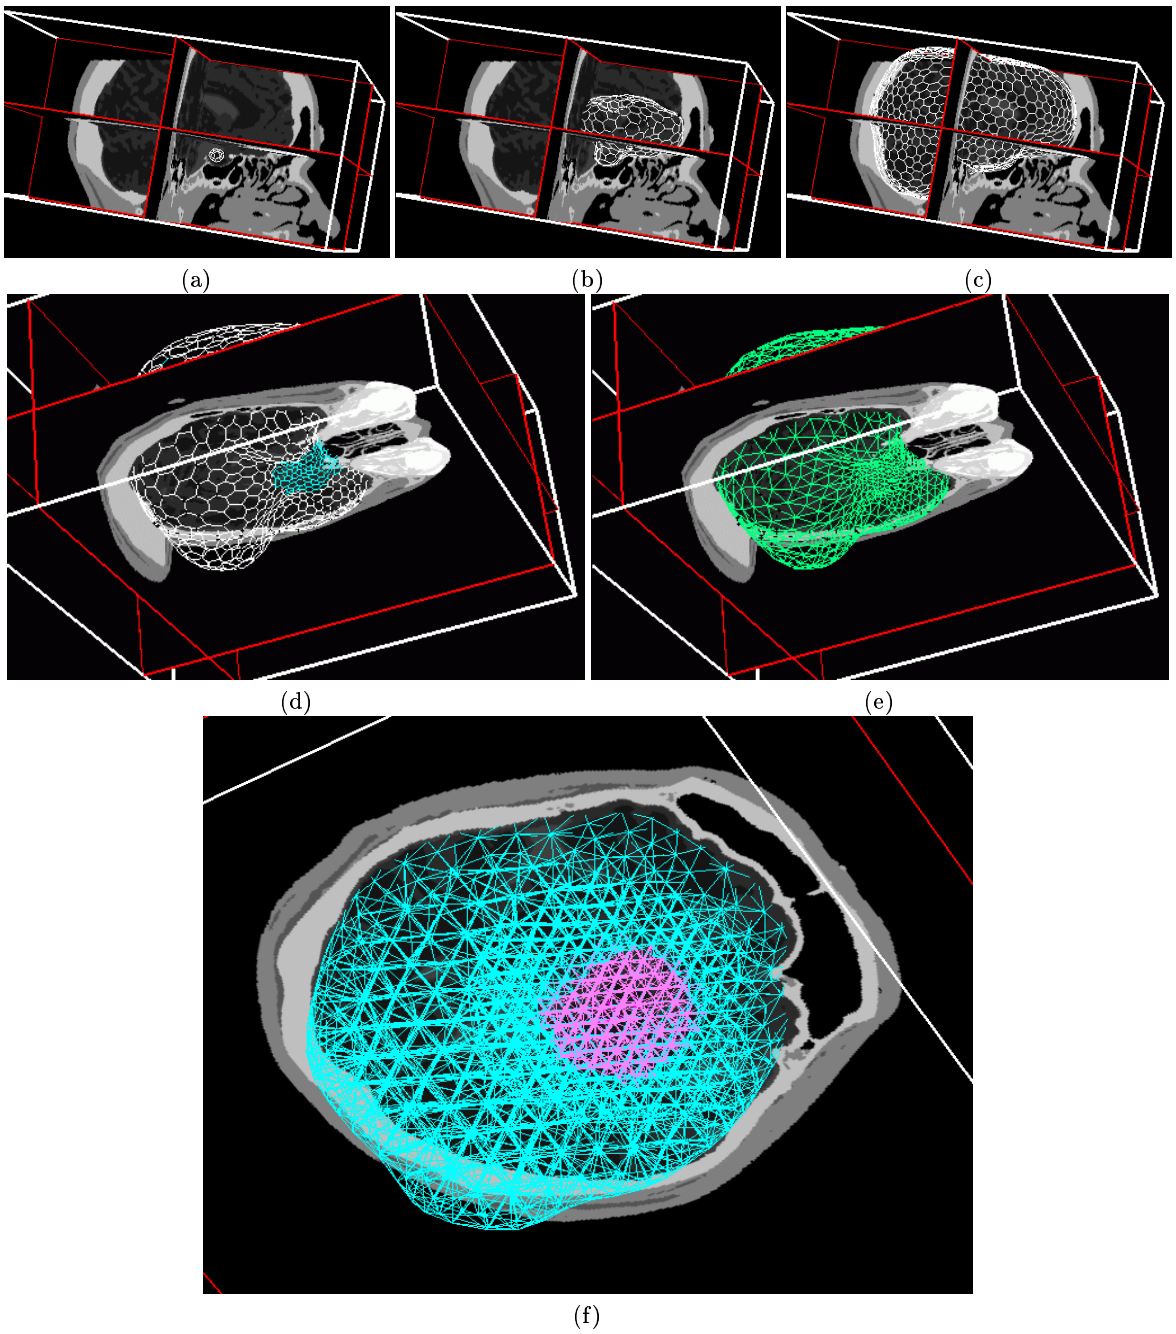

Fig. 5. Illustration of surface and volumetric meshing: (a)-(d) evolution of simplex mesh, at (a) 10, (b) 300 and (c) 1300 iterations at constant target mesh size; (d) radial target mesh size, with faces with minimal target area shown in turquoise; (e) final dual triangulated surface; (f) volumetric mesh, bounded by the surface in (e) (with child mesh shown in pink), featuring visibly near-regular structure.

5 Acknowledgements

The authors thank Dr. H. Delingette (INRIA) for generously providing simplex software, as well as Dr. H. Iseki and Mr. M. Sugiura (Tokyo Women's Hospital).

References

1. O.R. Astley & V. Hayward, Multirate Haptic Simulation Achieved by Coupling Finite Element Meshes Through Norton Equivalents, *IEEE Int. Conf. Rob. & Auto.*, 1998.
2. O.R. Astley, *A Software Architecture for Surgical Simulation Using Haptics*, Ph.D. thesis, McGill University, 1999.
3. O.R. Astley & V. Hayward, Design Constraints for Haptic Surgery Simulation, *Proc. IEEE Int. Conf. Rob. & Auto*, pp. 2446-2451, 2000.
4. M.A. Audette & K. Chinzei, Global Structure-preserving Voxel Classification for Patient-specific Surgical Simulation, *Proc. IEEE EMBS-BMES Conf.*, 2002.
5. M.A. Audette et al., A Review of Biomechanical Modeling of the Brain for Intrasurgical Displacement Estimation and Medical Simulation, submitted for publication, *Annals of Biomed. Eng*, 2003.
6. A.M. Bensaid et al., Partially Supervised Clustering for Image Segmentation, *Pattern Recognition*, Vol. 29, No. 5, pp. 859-871, 1996.
7. G.C. Burdea, *Force and Touch Feedback for Virtual Reality*, John Wiley & Sons, 1996.
8. P. Cappabianca et al., *Atlas of Endoscopic Anatomy for Endonasal Intracranial Surgery*, Springer, 2001.
9. V.B. Chial, S. Greenish & A.M. Okamura, On the Display of Haptic Recordings for Cutting Biological Tissues, *Haptics 2002 - IEEE Virtual Reality Conference*, 2002.
10. I. Ciric et al., Complications of Transsphenoidal Surgery: Results of a National Survey, Review of the Literature, and Personal Experience, *Neurosurg.*, Vol. 40, No. 2., pp. 225-236, Feb. 1997.
11. L.D. Cohen & R. Kimmel, Global Minimum for Active Contour Models: A Minimal Path Approach, *Int. J. Comp. Vis.*, Vol. 24, No. 1, pp. 57-78, 1997.
12. H. Delingette, General Object Reconstruction Based on Simplex Meshes, *Int. J. Comp. Vis.*, Vol. 32, No. 2, pp. 111-146, 1999.
13. H. Delingette and J. Montagnat, Shape and Topology Constraints on Parametric Active Contours *Comp. Vis. and Image Under.*, Vol. 83, pp. 140-171, 2001.
14. T. Deschamps & L.D. Cohen, Fast Extraction of Minimal Paths in 3D Images and Applications to Virtual Endoscopy, *Med. Imag. Anal.*, Vol. 5, pp. 281-299, 2001.
15. A. Fuchs, Almost Regular Triangulations of Trimmed NURBS-Solids, *Eng. w. Comput.*, Vol. 17, pp. 55-65, 2001.
16. K. Held et al., Markov Random Field Segmentation of Brain MR Images, *IEEE Trans. Med. Imag.*, Vol. 16, No. 6, pp. 878-886, 1997.
17. J.-O. Lachaud & A. Montanvert, Deformable Meshes with Automated Topology Changes for Coarse-to-fine Three-dimensional Surface Extration, *Med. Imag. Anal.*, Vol. 3, No. 2, pp. 187-207, 1998.
18. M. Mahvash & V. Hayward, Haptics Rendering of Cutting: A Fracture Mechanics Approach, *Haptics-e - The Electronic Journal of Haptics Research*, Vol. 2, No. 3, Nov. 20, 2001.

19. Microstar Laboratories, *www.mstarlabs.com.*
20. N.A. Mohamed et al., Modified Fuzzy C-Mean in Medical Image Segmentation, *Proc. IEEE ICASSP*, pp. 3429-3432, 1999.
21. NVidia Corp., *www.nvidia.com.*
22. J.T. Oden, Finite Elements of Nonlinear Continua, McGraw-Hill, 1972.
23. V.N. Parthasarathy et al., A Comparison of Tetrahedron Quality Measures, *Fin. Elem. in Anal. & Des.*, Vol. 15, pp. 255-261, 1993.
24. J.A. Sethian, *Level Set Methods and Fast Marching Methods: Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*, 2nd ed., Cambridge University Press, 1999.
25. T.S. Sumanaweera et al., MR Susceptibility Misregistration Correction, *IEEE Trans. Med. Imag.*, Vol. 12, No. 2, June 1993.
26. A. Vasilevskiy & K. Siddiqi, Flux Maximizing Geometric Flows, *IEEE Trans. Patt. Anal. & Mach. Intel.*, Vol. 24, No. 2, pp. 1565-1578, 2002.
27. P. Viola & W.M. Wells, Alignment by Maximization of Mutual Information, *Proc. 5th Int. Conf. Computer Vision*, pp. 15-23, 1995.
28. O.C. Zienkiewicz, *The Finite Element Method*, McGraw-Hill, 1977.