



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Musculoskeletal MRI segmentation using multi-resolution simplex meshes with medial representations

Benjamin Gilles*, Nadia Magnenat-Thalmann

MIRALab – University of Geneva, Battelle, Building A, 7 Route de Drize, CH-1227 Carouge, Switzerland

ARTICLE INFO

Article history:

Received 22 July 2008

Received in revised form 21 January 2010

Accepted 28 January 2010

Available online 1 March 2010

Keywords:

Image segmentation and registration
Musculoskeletal system
Magnetic Resonance Imaging
Discrete deformable models

ABSTRACT

The automatic segmentation of the musculoskeletal system from medical images is a particularly challenging task, due to its morphological complexity, its large variability in the population and its potentially large deformations. In this paper we propose a novel approach for musculoskeletal segmentation and registration based on simplex meshes. Such discrete models have already proven to be efficient and versatile for medical image segmentation. We extend the current framework by introducing a multi-resolution approach and a reversible medial representation, in order to reduce the complexity of geometric and non-penetration constraints computation. Our framework allows both inter and intra-patient registration (involving both rigid and elastic matching). We also show that the introduced representations facilitate morphological analysis. As a case study, we demonstrate that muscles, bones, ligaments and cartilages of the hip and the thigh can be registered at an interactive frame rate, in a time-efficient way (<30 min), with a satisfactory accuracy (~1.5 mm), and with a minimal amount of manual tasks.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Musculoskeletal disorders are certainly the most notorious and common causes of severe long-term pain and physical disability. For their diagnosis, the surgical planning and their post-operative assessment, the patient-specific modeling of the musculoskeletal system from medical images is an important problem (Blemker et al., 2007). The musculoskeletal system exhibits a complex geometry, difficult to model realistically (multiple organs in contact), a complicated mechanical behavior (viscoelastic, anisotropic, hyperelastic and non-linear), and complex interactions (e.g. confined cartilages within articulations). While simplified representations such as stick-figures and action lines (Delp and Loan, 1995) have proven to be useful for many applications in biomechanics, they have a limited accuracy as stressed in Teran et al. (2005), Ng-Thow-Hing (2000), Blemker and Delp (2005): they are unable to represent large attachment areas and accurately simulate global constraints such as volume preservation and non-penetration. Improvement in terms of accuracy could be achieved using surfacic models (Blemker et al., 2007) or equivalent reduced representations such as medial axis (Blum, 1964). Current interactive modeling methods (Ng-Thow-Hing, 2000; Aubel and Thalmann, 2001; Teran et al., 2005; Blemker and Delp,

2005) remain time-consuming and are not suitable for clinical use. Indeed, orthopaedists, biomechanicians and kinesiologists would like to simulate, visualize and navigate through articulations with a minimum amount of manual tasks. The automatic segmentation of bones and cartilages from medical images has been achieved in previous studies (Lorigo et al., 1998; Fripp et al., 2007). However, the complete segmentation of muscular complexes is not yet available.

Diagnosis tools used in the daily medical practice, especially medical scanner, are becoming increasingly precise, available, standardised, as well as less and less invasive. Besides traditional 3D images available from various modalities such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI), kinematical data are now getting more accessible (e.g. cine MRI, real-time MRI, ultrasound). MRI is a flexible modality for imaging both soft and bony tissues non-invasively, and has been chosen in this study. The potentially large amount of data makes data difficult to exploit. In this context, 3D anatomical models (e.g. shapes, surfaces, volumes) and 4D kinematical models (e.g. joint angles, deformation maps) could provide more insight and help in fusing (registering) data from diverse modalities or scan sessions. Registration is a fundamental problem in image and geometry processing (Maintz and Viergever, 1998; Zitová and Flusser, 2003). The matching problem is often expressed as an energy minimization and is generally difficult to solve due to the presence of local solutions. To tackle this, researchers are focusing on finding, in specific contexts, adequate registration features, similarity measures, parameterizations of the deformation and optimization methods.

* Corresponding author. Present address: University of British Columbia, Department of Computer Science, 2366 Main Mall, Vancouver, Canada V6T 1Z4.

E-mail addresses: benjamin.gilles@miralab.unige.ch, bgilles@cs.ubc.ca (B. Gilles), thalmann@miralab.unige.ch (N. Magnenat-Thalmann).

The goal is to evolve the solution (i.e. transformation) in a spatially coherent way through the introduction of priors, and along an as-convex-as-possible matching energy. Prior shape information relies on assumptions about surface regularity (smoothness, curvature) and variability across the population, while topological constraints exploit prior knowledge about organ interrelationships (e.g. contacts, motion discontinuities). These constraints can be efficiently applied on deformable models (Terzopoulos and Fleischer, 1988; Montagnat and Delingette, 2001). Particularly, discrete deformable models offer both flexibility and efficiency (Delingette, 1994; Montagnat and Delingette, 2005; Park et al., 2001; Ghanei et al., 1998; Lötjönen and Mäkelä, 2001), and are thus well suited for large datasets and displacements such as the ones involved in musculoskeletal imaging. Since complexity is an important issue, scalable approaches (McInerney and Terzopoulos, 1995; Szeliski and Lavallée, 1996; Lötjönen et al., 1999; Park et al., 2001) are also particularly adequate in our context.

In this paper, we present a new framework for image segmentation and registration based on discrete deformable surfaces, that is well suited to the musculoskeletal system. Our main contribution is the introduction of scalable simplex surfaces equipped with reversible medial representations. We consider the different issues involved in this process: the initial construction and adaptation of generic models, the computation of relevant internal and external constraints and the model evolution. Although our methods could be adapted to any type of discrete models (and most parametric models), we focus on simplex surfaces that are efficient in terms of flexibility and computational cost (Delingette, 1994; Montagnat and Delingette, 2005). To perform multi-object registration, we also introduce non-penetration and attachment constraints that take advantage of the multi-resolution and medial representations. Finally, we present quantitative results for the hip joint and the thigh using different MRI protocols. We show that our method can achieve quasi-automatic musculoskeletal MRI segmentation with a satisfactory accuracy and computational cost.

2. Methods

2.1. Simplex meshes

In this section, we describe the geometry and topology of simplex surfaces, and propose new local and global mesh adaptation mechanisms.

2.1.1. Definition

A k -simplex mesh is a k -manifold discrete mesh where vertices have exactly $k + 1$ distinct neighbors. Depending on their connectivity k , simplex meshes can represent various objects such as curves ($k = 1$), surfaces ($k = 2$) or volumes ($k = 3$), with any kind of topology. They were first described in Delingette (1994) for constrained 3D shape reconstruction and segmentation, and extended in Montagnat and Delingette (2005) for 4D segmentation, through the introduction of spatiotemporal constraints. Simplex meshes are dual to standard solid meshes exhibiting constant cell connectivity (e.g. three edges per face in a triangle mesh). In this present work, we focus on simplex surfaces ($k = 2$) as we perform object boundary registration. Volumetric properties are however given to the surfaces through our proposed medial representation (see Section 2.2.4). The power of simplex meshes lies in their simple local geometric description due to their constant vertex connectivity. The global three degrees of freedom of vertices can be turned into three local simplex parameters. From them and vertex neighbors, it has been shown that vertex positions are uniquely defined (Delingette, 1994). Simplex parameters can be viewed as reduced pyramid coordinates (Sheffer and Kraevoy, 2004). These parameters are

translation, rotation and scale invariant, so they are particularly suited for estimating the elastic, the bending and the strain energy of the surface, thus deriving geometric constraints. Fig. 1 shows the local geometry of simplex surfaces around a vertex \mathbf{P} . The linear combination of neighbors positions \mathbf{P}_i with weights ε_i (i.e. the metric parameters) defines the orthogonal projection \mathbf{P}_\perp onto the neighborhood domain along the normal \mathbf{n} . The angular parameter ϕ is related to the elevation h of \mathbf{P} above \mathbf{P}_\perp . The discrete mean curvature is given by $\sin(\phi)/r$ (Delingette, 1994).

2.1.2. Local topology adaptation

The geometric quality (uniformity of vertice repartition) and topological quality (uniformity of edge number among faces) of simplex meshes are important to ensure a good matching of the model with the object we want to represent. A regular simplex surface (constant edge length) is made of hexagons and its dual mesh of equilateral triangles (this is the centroidal Voronoi tessellation). Delingette (1994) shows that there are six basic operators from which all possible mesh connectivity changes are realizable. Most of the time, one wants to keep the object topology (same genus and number of holes) while optimizing the mesh topology. Just two operators (TO_1 and TO_2) remain in this case. It is possible to combine them to form macro-operators such as the exchange operation that changes the repartition of edges across faces without changing the total number of vertices, edges and faces (Delingette, 1994). For the dual triangle mesh, it corresponds to the swapping operator applied to two adjacent triangles. When proceeding to an exchange on 2-simplex meshes, we want to decrease the deviation of the number of edges per face. This yields to the criterion in Table 1 where n_i is the number of edges of the face i . Local resolution adaptation can be performed by combining TO_1 and TO_2 . These operations replace a faces with b faces inside a given domain $\sum_a S_i$ where S_i is the surface of the face i . Given the target face surface \tilde{S} , the improvement criterion is $\sum_b (S_i - \tilde{S})^2 < \sum_a (S_i - \tilde{S})^2$. Supposing that elements are equally distributed in space (similar size), the criterion is reduced to:

$$\begin{aligned} \sum_a S_i < \tilde{S}\sqrt{ab}, & \quad \text{if } (b - a) > 0 \text{ (res. decrease)} \\ \sum_a S_i > \tilde{S}\sqrt{ab}, & \quad \text{if } (b - a) < 0 \text{ (res. increase)} \end{aligned} \quad (1)$$

We have retained two particular macro-operators (TO_d and TO_i in Table 1) because they maintain a quasi-regular mesh (hexagons). Criteria based on the surface are not sufficient to ensure a good geometric quality, as faces can stretch in one direction. So

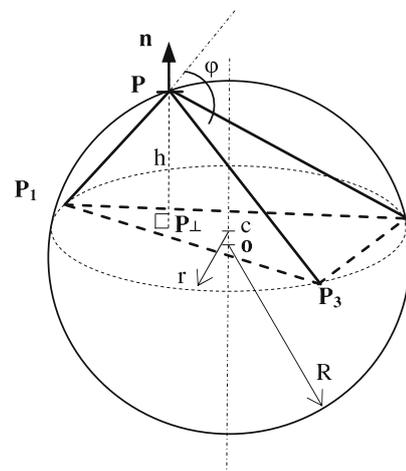
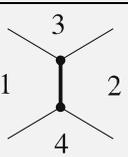
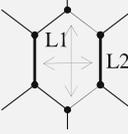
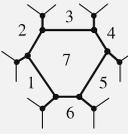


Fig. 1. Simplex surface local geometry.

Table 1
Summary of mesh topology adaptation operators. n_i, S_i and L_1/L_2 are respectively the number of edges, surface and principal lengths of a face i . \tilde{S} is the target face surface.

Operators/criteria	
	Exchange $\rightarrow n_1 + n_2 - n_3 - n_4 > 2$ $\leftarrow n_3 + n_4 - n_1 - n_2 > 2$
	TO_1/TO_2 $\rightarrow L_1 > L_2\sqrt{2}$ $\rightarrow S > \tilde{S}\sqrt{2}$ $\leftarrow \sum_2 S_i < \tilde{S}\sqrt{2}$
	TO_6/TO_7 $\rightarrow \sum_6 S_i > \sqrt{42}\tilde{S}$ $\leftarrow \sum_7 S_i < \sqrt{42}\tilde{S}$

we compute the two principal lengths L_1 and L_2 of the face and subdivide it along the 2nd direction if it leads to more homogeneous edge lengths.

2.1.3. Global mesh adaptation

It is necessary to change mesh resolution globally in order to adapt the complexity of the system according to the complexity of the problem. In image registration, the global-to-local approach (or coarse-to-fine approach) is widely used since it increases robustness at a lower computational cost (McInerney and Terzopoulos, 1995; Szeliski and Lavallée (1996), Lötjönen et al. (1999), Park et al. (2001)). With simplex meshes, it is possible to define different level of details through multi-resolution schemes that preserve the constant vertex connectivity. We propose a multi-resolution scheme that is based on the tessellation of the dual triangle mesh (see Fig. 2). This scheme preserves shape features as low-resolution vertices are contained in higher resolutions. When increasing the resolution, edge length is approximately divided by two and the number of vertices quadrupled (this is exact for regular hexagons).

2.2. Internal constraints

Because of noise, local solutions and aperture problems (apparent motion), it is desirable to incorporate prior geometric information through internal constraints. Here, we formulate for simplex surfaces three different types of constraints (smoothing, shape memory and volume conservation) that are particularly relevant in the context of musculoskeletal segmentation. We subsequently introduce radial constraints based on a medial representation. Due to the simple local geometric description of simplex meshes, constraints take computationally efficient formulations. Here, we express constraints related to a vertex \mathbf{P} through its desired position $\tilde{\mathbf{P}}$. In Section 2.5, they will be converted into forces.

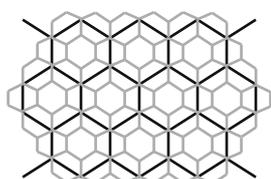


Fig. 2. Global resolution change. In black: coarse resolution. In grey: fine resolution.

2.2.1. Smoothing

We first exploit the fact that biological tissues have a smooth shape. First order elastic forces aim at minimizing the curvature. The popular Laplacian smoothing attracts vertices towards the centroid of its direct neighbors: $\tilde{\mathbf{P}} = \langle \mathbf{P}_i \rangle_1 (\langle \cdot \rangle)_\eta$ operator denotes local averaging in the neighborhood η of \mathbf{P}). Globally, this does not regularly spread vertices. Instead, we propose to use a weighted sum to average the size of the domains associated to vertices (rather than their spatial positions). Let be s_i the surface associated to a vertex \mathbf{P}_i . The new target position is defined as the barycenter of its neighbors, weighted by the size of their associated domains: $\tilde{\mathbf{P}} = \sum s_i \mathbf{P}_i / \sum s_i$. The normal component of $(\tilde{\mathbf{P}} - \mathbf{P})$ globally reduces the surface. To overcome this, several authors (Ghanei et al., 1998; Montagnat and Delingette, 2005) have proposed to replace it by a 2nd order term that averages the curvature (bending force). Here, we use a local average of the elevation $h = \|\mathbf{P}_\perp - \mathbf{P}\|$ (see Fig. 1). The final goal position is given by:

$$\tilde{\mathbf{P}} = \sum s_i \mathbf{P}_i / \sum s_i + \langle h_i \rangle_\eta \mathbf{n} \quad (2)$$

2.2.2. Shape memory

During segmentation, it is effective to use prior shape information through a reference template model. By predefining simplex parameters, we can enforce hard constraints on shapes (Delingette, 1994). The use of the simplex angle ϕ is restrictive because \mathbf{P} is uniquely defined only if the projection \mathbf{P}_\perp lies inside the circumscribed circle (see Fig. 1). Instead we directly use the elevation parameter h , and more exactly the parameter $h_n = h \cdot St^{-1/\beta}$ where St is the area of the triangle formed by the three neighbors and β a parameter that tunes the scale invariant aspect. With $\beta = 2$, the definition is similitude invariant; with $\beta = \infty$, the definition is only invariant through rigid transforms. In practice, we take $\beta = 4$ which has proven to be adequate to limit surface distortion. Given the three reference shape parameters $\tilde{\varepsilon}_1, \tilde{\varepsilon}_2$ and \tilde{h}_n , the final target position is given by:

$$\tilde{\mathbf{P}} = \tilde{\varepsilon}_1 \mathbf{P}_1 + \tilde{\varepsilon}_2 \mathbf{P}_2 + (1 - \tilde{\varepsilon}_1 - \tilde{\varepsilon}_2) \mathbf{P}_3 + \tilde{h}_n St^{1/4} \mathbf{n} \quad (3)$$

2.2.3. Global volume preservation

For intra-patient registration, it is useful to exploit the incompressible property of biological tissues. We apply global volume conservation through normal constraints. We perform fast volume computation by applying the divergence theorem to the dual triangle mesh. Let \tilde{V} , V and ζ be the target volume, current volume and current surface of a closed mesh. The assumption of a similar displacement for all vertices along the normal yields to:

$$\tilde{\mathbf{P}} = \mathbf{P} + \frac{\tilde{V} - V}{\zeta} \mathbf{n} \quad (4)$$

2.2.4. Medial representation and radial constraints

Muscles generally have a smooth and tubular shape, that can be represented by an underlying piece-wise action line (Delp and Loan, 1995; Blemker and Delp, 2005; Teran et al., 2005). However, muscles with large attachment areas and several origins/insertions require many action lines that are not well defined (Ng-Thow-Hing, 2000; Aubel and Thalmann, 2001). On the contrary, the medial axis, in the geometric sense, is unique and is able to reconstruct exactly any object (Blum, 1964). It is defined as the union of maximal balls inside an object (parameters are positions and radii) and is characterized by three main properties: homotopy equivalence, good localization and reversibility. We propose a new method to automatically compute medial surfaces of musculoskeletal tissues, which will better follow these properties than previous methods. It leads to enhanced mechanisms for constraining surfaces and char-

acterizing shapes. Several medial axis approximation methods have been proposed, based on Voronoi diagrams, on distance maps, or on thinning (see Attali et al. (2007) for a complete review). These approximations need to be pruned to remove undesirable branches and spikes (Amenta et al., 2001). To create a simplified medial axis, we take a different approach: we iteratively fit a predefined shape to the center of the object like in Terzopoulos et al. (1987) (where a parametric curve is pushed towards the center of a parametric surface by energy minimization). This guarantees homotopy equivalence and smoothness properties as shown in Fig. 3.

We make the hypothesis that, for muscles, ligaments and cartilages, medial axis will be homeomorphic to planes. We represent both the boundary and the medial sheet of these organs with simplex surfaces, and apply constraints of them. Medial vertices \mathbf{Q}_j are constrained to remain in the center of the object, while boundary vertices \mathbf{P}_i are constrained to match the reconstructed shape from the medial axis. This dual representation is flexible as it allows both the approximation of the true medial axis and the shape control of the object boundary through radial constraints. Our approach closely match the work on m-reps by Pizer et al. (2003). In their study, additional parameters are considered for each medial vertex (local frame and object angle) to explicitly derive geometric sections of the boundary mesh. In our method, the sampling and the geometry of the boundary mesh is not fixed by the medial surface; the two surfaces being related by spring-like forces.

The goal of radial constraints is to reach a state where boundary vertices \mathbf{P}_i lie on the surface of the reconstructed object, that is the union of maximal balls (see Fig. 4). In other words, we want that the distances between \mathbf{P}_i and their orthogonal projections on the medial axis $\mathbf{P}_{i\perp}$ correspond to the interpolated radii r_i . The error can be calculated as $err = \langle |r_i - \|\mathbf{P}_{i\perp} - \mathbf{P}_i\|| \rangle$ ($\langle \cdot \rangle$ denotes averaging over all points). Let R_j be the radii associated to medial axis vertices \mathbf{Q}_j and w_{ij} the barycentric coordinates of the projection such as: $\mathbf{P}_{i\perp} = \sum_j w_{ij} \cdot \mathbf{Q}_j$ and $r_i = \sum_j w_{ij} \cdot R_j$. The parameters w_{ij} and R_j define the links of the medial surface with the boundary mesh. In the step of building medial surfaces from boundary meshes, we fix the boundary vertices and update these parameters at each iteration: we compute w_{ij} by projecting \mathbf{P}_i orthogonally to the medial surface and estimate the radii R_j through the weighted mean: $R_j = \sum_i w_{ij} \|\mathbf{P}_{i\perp} - \mathbf{P}_i\| / \sum_i w_{ij}$. We experienced that this expression leads to nearly optimal r_j , minimizing err , at a reasonable computational cost. Now, given that all w_{ij} and R_j are known (they have been updated or memorized from a reference state), we derive particle constraints through the desired positions $\tilde{\mathbf{Q}}_j$ and $\tilde{\mathbf{P}}_i$ that minimize the error. We want to displace \mathbf{P}_i and $\mathbf{P}_{i\perp}$ in opposite directions, so we have:

$$\mathbf{P}_i - \tilde{\mathbf{P}}_i = \tilde{\mathbf{P}}_{i\perp} - \mathbf{P}_{i\perp} = \sum_j w_{ij} (\tilde{\mathbf{Q}}_j - \mathbf{Q}_j) \quad (5)$$

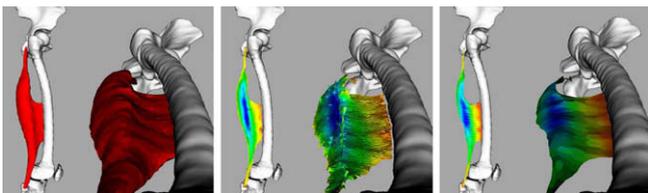


Fig. 3. From left to right: boundary surface of the biceps femoris muscle; its medial axis constructed with the powercrust method (Amenta et al., 2001) where some branches have been pruned; its medial axis constructed with our method. Colors represent the radius parameter of the medial axis. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

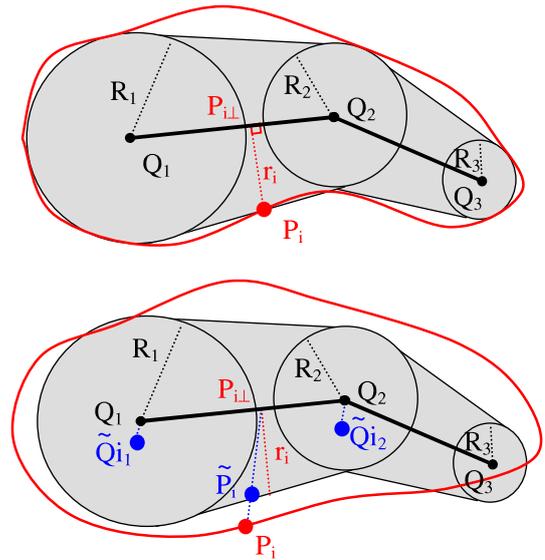


Fig. 4. A 2D illustration of the desired state (left) and deformed state (right). The boundary surface is represented in red, the medial surface in black and the reconstructed object from the medial surface in grey. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

To comply with the momentum conservation principle and thus avoid ghost translation and spinning, we suppose that the displacements of \mathbf{Q}_j are collinear with $(\mathbf{P}_i - \tilde{\mathbf{P}}_i)$ and proportional to w_{ij} . We obtain $(\tilde{\mathbf{Q}}_j - \mathbf{Q}_j) = w_{ij} (\mathbf{P}_i - \tilde{\mathbf{P}}_i) / \sum_j w_{ij}^2$. The desired position $\tilde{\mathbf{P}}_i$ can be defined so that $\tilde{\mathbf{P}}_{i\perp} - \mathbf{P}_i$ remains normal to the surface. However, this would produce undesirable distortions and self intersections when the model is bent. A better behavior is obtained by seeking reference radii only. This allows model vertices to slide onto the medial balls while smoothing forces will prevent from excessive sliding. The final goal positions of boundary vertices are given by:

$$\tilde{\mathbf{P}}_i = \mathbf{P}_{i\perp} - r_i \frac{\mathbf{P}_{i\perp} - \mathbf{P}_i}{\|\mathbf{P}_{i\perp} - \mathbf{P}_i\|} \quad (6)$$

Vertices \mathbf{Q}_j share multiple projections $\mathbf{P}_{i\perp}$, so we average their contribution. This leads to the goal positions of medial vertices:

$$\tilde{\mathbf{Q}}_j = \mathbf{Q}_j + \left\langle w_{ij} \frac{\mathbf{P}_i + r_i \frac{\mathbf{P}_i - \mathbf{P}_{i\perp}}{\|\mathbf{P}_{i\perp} - \mathbf{P}_i\|} - \mathbf{P}_{i\perp}}{\sum_j w_{ij}^2} \right\rangle_i \quad (7)$$

2.3. Inter-object constraints

In this section, we derive constraints related to the interaction of the different objects (contact and attachments).

2.3.1. Medial axis-based collision handling

In multi-organ segmentation, collision handling is important to guarantee that objects do not self-penetrate. Collision detection is often considered as the bottleneck for simulation, since it involves pair-wise penetration tests for *all* primitives (points, edges, faces, cells) of *all* models involved in the simulation. Significant speed up can be achieved based on distance fields, spatial subdivision, image-space techniques, bounding volume hierarchies or stochastic methods (Teschner et al., 2005). The *distance field* method with uniform grids is superior for rigid models because collision query is very fast (the main idea is to pre-process and store the distance to the surface within a uniform grid). We apply this method after bone segmentation, to accelerate soft tissues-to-bone collision

detection. For deformable objects, we chose a bounding volume hierarchy (BVH) method to quickly eliminate most of the non-colliding primitives. We use 18-Discrete Oriented Polytope (18-DOP) quadtree (Mezger et al., 2003) that exhibits good properties (fast tests and updates, compact bounding volumes). A k -DOP is a convex polyhedron defined by k half-spaces $H_i, i \in \{0, \dots, k-1\} : H_i = \{\mathbf{P} \in \mathbb{R}^3 | \mathbf{n}_i \cdot \mathbf{P} \leq b_i\}$ where normal vectors \mathbf{n}_i are chosen by opposite pairs to form $k/2$ intervals. They generalize the concept of bounding boxes (6-DOP), by adding other directions such as edge pairs (18-DOP) and corners (26-DOP).

Collision detection can benefit from our medial representation which is more compact. The idea is to store faces of medial axis into the BVH and inflate bounding volumes according to radii (see Fig. 5). The object inside is always contained in the bounding volume (contrary to the bounding volumes of the surface). So, collision detection, like for any volumetric model, is simpler because colliding primitives result in overlapping bounding volumes. Out of the detection process, a collision vector \mathbf{P}_c between a vertex \mathbf{P} and face composed of vertices \mathbf{P}_i is expressed as the linear combination: $\mathbf{P}_c = \sum_i w_i \mathbf{P}_i - \mathbf{P}$. When permanent non-sliding contacts are expected, we pre-detect collisions at the initialization (i.e. we store the indexes of the colliding vertices and the weights). During simulation, we only need to update \mathbf{P}_c according to the current positions, which is fast.

Collision response aims at constraining particles to reach a non-penetrating state. Typically, from a current state where collisions are detected, it consists in altering particle positions, velocities and/or accelerations to cancel collisions in the current state and to prevent from collisions in the next states. A good scheme has to respond smoothly to particle changes. This is why a proximity region is commonly considered to allow a gradual collision correction scheme and prevent particle jumping at object interfaces (Volino and Magnenat-Thalmann, 2000a). After testing the convergence time and stability of different alteration methods in our context, we have decided to use an impulse-based method (Mirtich and Canny, 1994): we alter collision velocities such as $\Delta \mathbf{V}_c = -\mathbf{V}_c - \mathbf{P}_c/dt$. In the proximity region, the normal component is gradually altered. Then, the different corrections related to each particle are combined through a simple weighted sum.

2.3.2. Attachments

Hard constraints on vertex positions can be enforced by setting their mass to infinite. This technique is used to attach soft-tissues (muscles, cartilages and ligaments) to bones. To select the vertices to attach and their positions, we have developed a method to define attachment areas. We use cardinal splines with a reduced number of control points to interactively place and adjust the attachments. For each attachment, two splines are used: one on the bone and one on the soft organ. By using simple projection and filling techniques, we select boundary and internal soft tissue

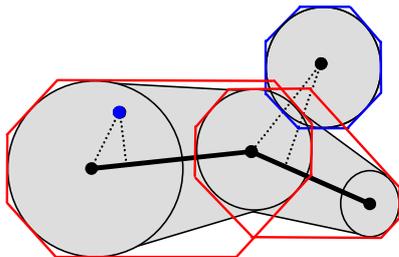


Fig. 5. In this 2D example, a vertex and a medial point (bounded by blue 8-DOP) are tested towards a medial surface (bounded by red 8-DOPs). Our medial axis-based collision detection method detects 4 collisions (dashed lines). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

vertices (see right picture of Fig. 6). Boundary vertices are easily warped to bones through spline curvilinear coordinates. For internal vertices, we use a relaxation approach: we apply smoothing constraints until convergence, and finally project internal vertices to the bone surface. To transfer attachment locations from one subject to another, we warp spline control points through their local coordinate on the bone surface (defined as barycentric coordinates). As shown in Kaptein and der Helm (2004), most attachments rely on bone geometrical features. Therefore, our approach can effectively estimate subject-specific attachments from a generic model. Attachment localization from MRI was difficult. Instead, we relied on the literature on anatomy to place the generic splines.

2.4. Image-based constraints

Here, we express different external constraints that can be applied to align models with the data. They will be compared in the context of musculoskeletal MRI in Section 3.2.2. These constraints are expressed in terms of optimal shifts $\tilde{\mathbf{j}}$ in the surface normal direction. Using normal direction reduces aperture problems (Lötjönen and Mäkelä, 2001; Montagnat and Delingette, 2005). Along this direction, we perform regular sampling according to a step size s (set to the minimum voxel dimension). We have $\tilde{\mathbf{P}} = \mathbf{P} + \tilde{\mathbf{j}}\mathbf{n}$.

2.4.1. Gradient maximization

To directly segment an image/volume \mathbf{T} , the main criterion that discriminates object boundaries is the gradient $\nabla \mathbf{T}$. We can either maximize the gradient magnitude (G measure) or the alignment of model normals with the gradient vectors (GV measure). Given a search distance d , the goal positions are given by:

$$G : \tilde{\mathbf{j}} = \underset{(-d/s < \tilde{\mathbf{j}} < d/s)}{\operatorname{argmin}} (-\|\mathbf{VT}(\mathbf{P} + \tilde{\mathbf{j}}\mathbf{n})\|) \quad (8)$$

$$GV : \tilde{\mathbf{j}} = \underset{(-d/s < \tilde{\mathbf{j}} < d/s)}{\operatorname{argmin}} (\pm \mathbf{VT}(\mathbf{P} + \tilde{\mathbf{j}}\mathbf{n}) \cdot \mathbf{n})$$

In the GV expression, the sign is obtained ad-hoc from the structure to segment. If not known, one can use the absolute value of the dot-product.

2.4.2. Image registration

In deformable model-based image registration, the goal is to align a source (reference) image \mathbf{S} to the target image \mathbf{T} via model deformations. The model, initially aligned to the segmented reference dataset (vertex positions \mathbf{P}^0), is iteratively deformed until it matches the target image, that is where the similarity in vertex neighborhoods η is maximal:

$$\tilde{\mathbf{j}} = \underset{(-d/s < \tilde{\mathbf{j}} < d/s)}{\operatorname{argmin}} \Delta(\mathbf{S}(\eta(\mathbf{P}^0)), \mathbf{T}(\eta(\mathbf{P} + \tilde{\mathbf{j}}\mathbf{n}))) \quad (9)$$

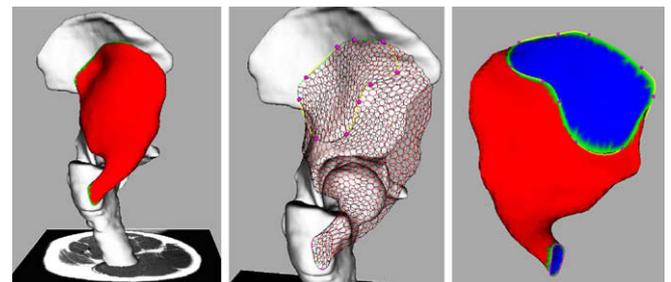


Fig. 6. An example of spline-based attachment between a muscle (gluteus minimus) and hip bones. From left to right: a 3D visualization of the models; a closer view showing the spline attached to the pelvis; the selected attachment vertices of the muscle (in green: the boundary vertices; in blue: the internal vertices). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

where Δ is an image distance measure (Holden et al., 2000; Maintz and Viergever, 1998; Zitová and Flusser, 2003) such as the sum of absolute differences (AD) or the normalized cross-correlation (NCC). As in Montagnat and Delingette (2005), we apply a 1D registration of intensity profiles: vertex neighborhoods η are composed of regularly sample points along the surface normal direction \mathbf{n} , which is the most relevant direction (i.e. direction of expected intensity change). The sampling period is set to s . 1D intensity profiles have a length sNo in the direction of the normal (i.e., outside the model) and a length sNi in the opposite direction (i.e., inside the model). It is possible to filter intensity profiles by convolving them with 1D discrete masks in order to enhance edges (e.g. Prewitt mask $[-1, 0, 1]$ or $[-1, 2, 1]$). Different distance measures and masks will be compared in Section 3.2.2.

2.5. Dynamic evolution

In this section, we describe how we deform models according to the constraints derived in the previous sections, and how we exploit the multi-resolution scheme to globally and locally control shapes.

2.5.1. Numerical integration

Deformation is obtained through the minimization of the internal and external energy. We combine target positions of the different constraints through the weighted sum: $E = 1/2 \sum \alpha_i (\mathbf{P}^i - \mathbf{P})^2$ with $\sum \alpha_i = 1$ (here \mathbf{P}^i denotes the target position of \mathbf{P} for the constraint i and α_i is the stiffness of this constraint). Because of the non-convex aspect of the energy (noise, multiple local minima, etc.), there is no analytical solution for the minimization and one must use traditional multivariate optimization strategies (Maintz and Viergever, 1998; Zitová and Flusser, 2003). By analogy with a mechanical system, one strategy is to consider discrete models as punctual masses (lumped mass particles) evolving under the Newtonian law of motion. The Newton equation leads to a first-order differential equation system relating forces to particle state (velocity and position). Time discretization leads to either explicit or implicit methods (Press et al., 1992). Finite differences lead to explicit resolution methods (e.g. explicit Euler, Runge-Kutta) that are known to go forward blindly so that the stability is only ensured for small time steps. A Taylor expansion of the force is more suitable but requires the evaluation of force derivatives with regards to particles position and velocity and the inversion of a sparse matrix. This is known as the implicit method. It has shown good performances for stable real-time applications (Baraff and Witkin, 1998; Volino and Magnenat-Thalmann, 2000b), despite its possible inaccuracy (ghost numerical damping due to non-linear non-stiff behavior). In this work, we have chosen an implicit Euler integration and a conjugate gradient resolution such as in Baraff and Witkin (1998), Volino and Magnenat-Thalmann (2000b). We have observed its faster convergence and higher stability in our context, compared to other schemes (explicit Euler, Runge-Kutta, Implicit-Midpoint, BDF, Newton-Raphson).

For simplicity, we consider that target positions are independent from the current positions. So, the force \mathbf{F} acting on \mathbf{P} is expressed as: $\mathbf{F} = -\nabla E = \sum \mathbf{F}^i = \sum \alpha_i (\mathbf{P}^i - \mathbf{P})$. Taking into account the dependence improves the accuracy but leads to complex expressions (particularly for the Jacobians) that significantly increase computational cost. Our tests showed that, in our context, it does not improve convergence time and positions. This assumption allows us removing all non-diagonal terms of the stiffness matrix (which is equivalent to the Jacobi preconditioner): $\partial \mathbf{F}_i / \partial \mathbf{P}_j = \mathbf{0}$, $i \neq j$. For diagonal terms, we approximate the Jacobian of the constraint i by:

$$\frac{\partial \mathbf{F}^i}{\partial \mathbf{P}} = -\alpha_i \mathbf{c} \mathbf{I} - \alpha_i (1 - c) \mathbf{F}^i \mathbf{F}^{i\top} / \|\mathbf{F}^i\|^2 \quad (10)$$

where c is a constant related to force isotropy (see Volino and Magnenat-Thalmann (2000b) for more details). For smoothing and shape memory forces, we consider an isotropic variation: $c = 1$ (an infinitesimal displacement does not change the target position). The other constraints (volume preservation, radial constraints, image-based forces) are anisotropic (an infinitesimal displacement does not change the direction of the force). For them, we set $c = 0.1$ (a perfectly anisotropic variation leads to instability, and a small isotropic contribution, that act as a damping factor, need to be added in practice). To reduce oscillations that can occur in the particle system, we add a damping force inversely proportional to the particle velocity \mathbf{V} in the direction of the force: $\mathbf{F}_d = -\gamma (\mathbf{V}^T \mathbf{F}) \mathbf{F} / \|\mathbf{F}\|^2$ with $\gamma = 0.5$.

2.5.2. Multi-resolution scheme

The use of different levels of details (LODs) during the deformation process is valuable for registration since it reduces system complexity and sensitivity to local solutions. Here, we propose to use the different scales generated with our global topology adaptation scheme to derive multi-resolution constraints. A straightforward approach is to use LODs successively and independently to allow a coarse-to-fine optimization, as done in most multi-scale registration methods (McInerney and Terzopoulos, 1995; Szeliski and Lavallée, 1996; Lötjönen et al., 1999; Park et al., 2001). On the contrary, our approach combines all LODs at the same time. Our goal is to keep global constraints when fine LODs are used and thus prevent from global drifting. The idea is to quickly propagate constraints from lower resolutions to a current simulation level. This is simply done by linear interpolation using the same technique than for positions (see Fig. 2). Given a vector \mathbf{U}_i^r attached to the vertex i at resolution r and the mean vector \mathbf{U}_i^r related to the face l , the vector interpolation scheme is given by:

$$\text{Resolution increase : } \mathbf{U}_i^{r+1} = (\mathbf{U}_i^r + \overline{\mathbf{U}}_i^r) / 2 \text{ or } \mathbf{U}_i^{r+1} = \mathbf{U}_i^r$$

$$\text{Resolution decrease : } \mathbf{U}_i^r = \mathbf{U}_i^{r+1}$$

\mathbf{U} can be either a force, a displacement, a velocity correction term or a Jacobian. So, all constraints can be finally summed into a unique LOD. A significant benefit of our multi-scale method is related to collision handling: collision detection can be performed at coarser LODs, while the response is being passed to the current resolution. Exact contact computation is most of the time not relevant in our context, as fat separates organs.

For image constraints, multi-resolution forces are not appropriate because they do not depend on the local resolution of the mesh. Because image constraints are noisy, we need to regularize them globally. Such as in Montagnat and Delingette (2005), we compute for each model the linear transform $\overline{\mathbf{M}}$ that optimally approximates the target displacements $(\overline{\mathbf{P}}_i - \mathbf{P}_i)$ in the least-squares sense: $\overline{\mathbf{M}} = \text{argmin} \sum \|\overline{\mathbf{P}}_i - \mathbf{M} \mathbf{P}_i\|^2$. Closed form formulas exist for common linear transformations such as the affine, similarity and rigid transforms (Horn et al., 1988). To tune the global/local effect of image-based constraints, we replace target displacements by the weighted sum $\lambda \overline{\mathbf{P}}_i + (1 - \lambda) \mathbf{M} \mathbf{P}_i$ where $\lambda \in [0, 1]$ is the control parameter.

3. Results

In this Section, we apply our algorithm on one part of the musculoskeletal system: bones, muscles, ligaments and cartilages of the hip of the thigh. We generate regular generic models through our mesh adaption techniques and construct the associated medial surfaces. We individualize these models from patient-specific MRI

images and validate the methods through comparisons with manual segmentation. We finally present performance tests.

3.1. Generic model reconstruction

3.1.1. Mesh topology optimization

Thanks to our fast collision handling method, users can constrain models according to manually-defined positions in space. By putting internal/external/boundary constraint points, we have reconstructed generic models to be used as reference models for quasi-automatic registration. For the different organs to be reconstructed, a simplex sphere or a cylinder is used as a starting primitive. Meshes are progressively optimized to improve their geometric and topological quality according to the shape we want to approximate (Section 2.1.2). Here, we present the example of the femur due its particularly irregular shape. The straightforward matching of a simplex sphere (without any mesh optimization) leads to undesirable folding and overstretched faces as shown in Fig. 7, left picture. To test our method, we fix the target edge length to $\bar{l} = 15$ mm, to 8 mm and then back to 15 mm (Fig. 8). As we want to seek hexagons, the corresponding target face surface is $\bar{S} = 3\sqrt{3}\bar{l}^2/2$. In our tests, the average edge length quickly matched the target values with a low standard deviations (typically 1.5 mm). A common measure of the regularity of triangle meshes is to plot the histogram of radius ratio (that is $2r/R$ where r is the radius of the inscribed circle and R the radius of the circumscribed circle of a triangle). As shown Fig. 8, our resulting meshes possess quasi-regular elements (the histograms of the dual triangle meshes are close to 1). The introduced operators TO_i/TO_d are clearly valuable compared to the use of TO_1/TO_2 only, since they reduce computation time and closer match target edge lengths.

3.1.2. Medial axis construction

We apply the methodology of Section 2.3.2 to attach muscles, ligaments and cartilages to bones. From the attachments, we automatically generate medial axis as following: for each model, we initialize a simplex plane as the optimal plane approximating attachment centroids (see Fig. 9B and C). We project vertices of the boundary model to the plane and prune exterior faces (see Fig. 9D). After this cropping stage, we simulate the medial surface with radial and smoothing forces, while projecting model vertices and updating radii at each timestep as described in Section 2.2.4. The medial surface converged to the medial axis in a couple of seconds. This process has proven to be robust for all the models we have tested (21 muscles, 3 cartilages, 3 ligaments). We have mea-

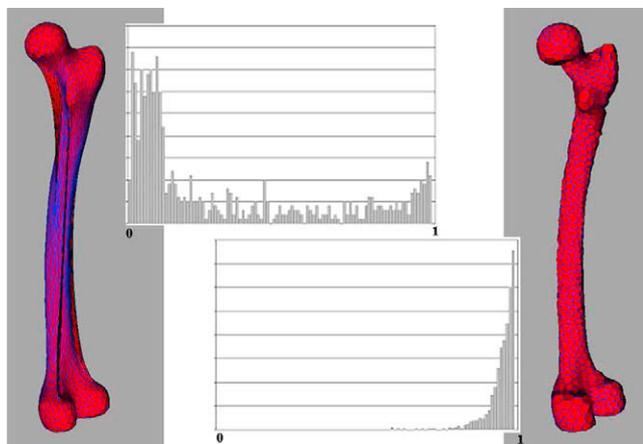


Fig. 7. The unoptimized femur model (left) presents a poor topological and geometrical quality as shown by its radius ratio histogram, while the optimized model (right) possess quasi-regular faces.

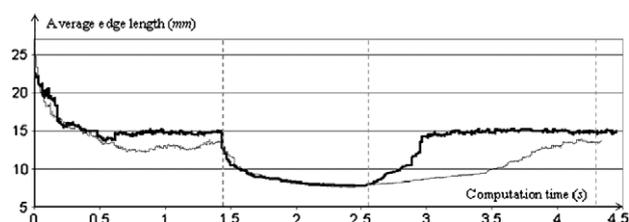


Fig. 8. Example of mesh topology adaptation: the target edge length is successively fixed to 15, 8 and 15 mm. The addition of TO_i/TO_d (in black) to TO_1/TO_2 (in grey) significantly speeds up convergence.

sured the error err between the reconstructed surface from the medial axis and the boundary mesh. As shown in Fig. 10, there is a relation between the error and the sampling of the medial surface. We measure the compression factor as the ratio between the number of model vertices times 3 (three spatial coordinates) and the number of medial vertices times 4 (three spatial coordinates plus the radius). As a trade-off between the compression factor and the error, we have decided to use an initial edge length of 8 mm for planes, leading to an average compression factor of 14 and an average error of 0.61 ± 0.65 mm. It allows, on one side, a precise shape representation (for shape analysis, radial regularization and medial axis-based collision handling) and, on the other side, a significant reduction of the required information (about 14-fold reduction).

3.2. Patient-specific model reconstruction from MRI

3.2.1. MRI acquisition

MRI protocols definition has been done in close collaboration with physicians from radiology and orthopaedic departments (HUGE – Geneva). The goal was to obtain images carrying sufficient information with clinically achievable protocols (fast). Acquisition has been performed on a 1.5 T Intera MRI system (Philips Medical Systems, Best NL). We have investigated a high-resolution and a low-resolution protocol to perform the full acquisition of the hip and the thigh. The first protocol is based on an Axial 2D T1 Turbo Spin Echo sequence (Table 2, sequence A). Due to restricted acquisition time, it was not possible to achieve a precise (isotropic) acquisition over the full region, so the slice thickness was adjusted to 2 mm for the hips, 4 mm for the knees or 10 mm for the thigh. The total acquisition lasted about 40 min (~ 250 slices). To improve image quality (SNR), a surface coil around the hip was used. This protocol was applied on 13 young healthy volunteers (7 male, 6 female). Besides this high-resolution protocol that provided a set of high-quality models, we performed a clinical study on 30 young female dancers where several poses were analyzed: the neutral position and the right and left splits. The former protocol was not applicable, so a faster (but low-resolution) protocol was proposed (Table 2, sequences B., C. and D.). It consisted of two axial 3D series (acquisition time: 3 min). A high-resolution sequence centered on joints with isometric voxels was run to improve cartilage/ligament and bone models near articulations (acquisition time: 4 min per hip). Moreover, because the purpose was to study coxo-femoral conflicts and labrum deformations, radial acquisitions were run to provide highly detailed images around the acetabulum (acquisition time: 6 min per hip).

3.2.2. Optimizing image-based constraints

To improve contour discrimination in images, we have optimized and compared parameters related to the constraints defined in Section 2.4. Our gold standards are manually registered models from two subjects. They correspond to the optimal shifts \hat{j}_i s of vertices \mathbf{P}_i in the normal directions, and their accuracy is about 1 mm

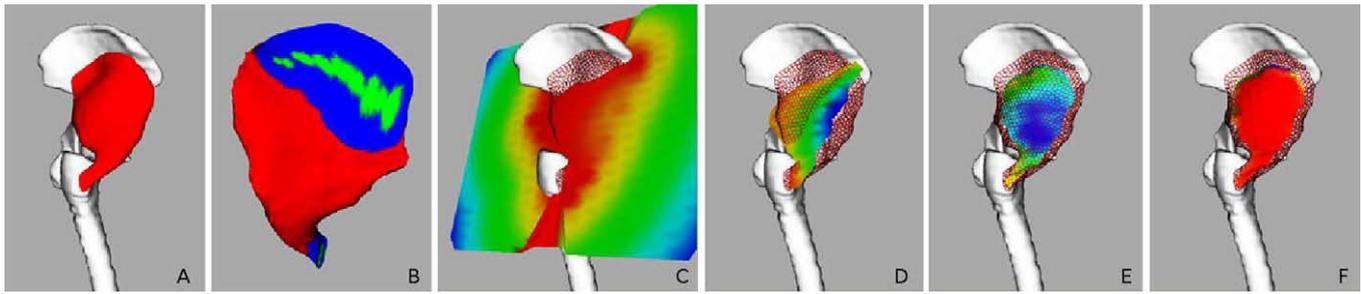


Fig. 9. Medial axis generation. From left to right: A. Original gluteus minimus model, attached to hip bones. B. Attachment centroid faces (in green). C. Plane initialization. D. Plane after cropping. E. Medial axis after fitting. F. Reconstruction error. In C, D and E, colors show the radius of the medial axis. In F, colors show the reconstruction error (blue: max, red: min). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

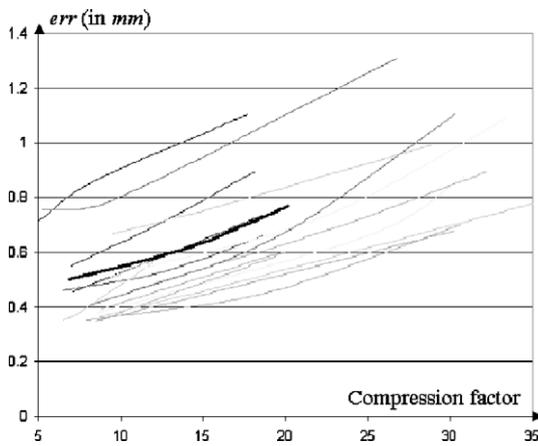


Fig. 10. Final medial axis error in function of the compression factor for different models (average in bold).

(voxel size). Fig. 11 shows intensity and similarity profiles for the femur bone.

In a first step, we compute the optimal size of intensity profiles by minimizing the distance between the optimal shifts $\hat{j}_i s$ found with the considered profiles (see Section 2.4), and the gold standard shifts $\hat{j}_i s$. Optimal inside and outside lengths (sNi and sNo) were found fairly independent from the similarity measure. Being more variable, the external part of intensity profiles was less relevant. The final lengths were $Ni = 25$ and $No = 5$ for bones and $Ni = 20$ and $No = 10$ for muscles, with a step size of $s = 0.5$ mm.

No similarity measure has proven to perform the best in the general case (Holden et al., 2000). Here, we compare, in our context, gradient maximization techniques with standard similarity measures widely used in mono-modal registration: the sum of absolute differences (AD) and the normalized cross-correlation (NCC). Our similarity profile length (search space) is fixed to $d = 1$ cm. We applied the five evaluation criteria proposed in Skerl et al. (2006). The normalized cross-correlation (NCC) was always the best similarity measure. The use of gradient images is not worth for this type of mono-modal registration. However the discrete mask $[-1, 2, 1]$ (gradient enhancement) slightly improves

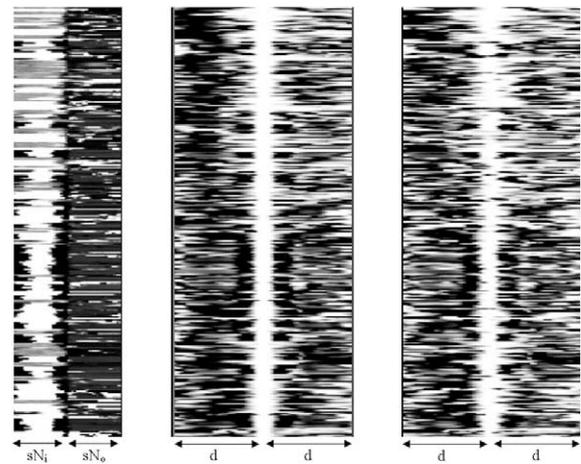


Fig. 11. Left: Sample intensity profiles (femur). Center: The corresponding self-similarity profile (NCC measure). Right: An inter-subject similarity profile (NCC measure).

the results without significant extra computational cost. Gradient maximization (G and GV measures) leads to good results in the close neighborhood of the solution and is not sensitive to errors in the generic models. Therefore, for registration, we combine a measure based on intensity profile (NCC) with a large search space, and a measure based on the gradient with a reduced search space.

3.2.3. Model registration

Our registration scheme works sequentially. Bones are first aligned. Then cartilages, ligaments and muscles are initialized from the bones. These soft-tissues are deformed to match individual anatomy.

3.2.3.1. Bone initialization. First, generic bone models are coarsely initialized using manually placed points corresponding to anatomical landmarks (7 per pelvis and 8 per femur) and thin-plate-spline (TPS) interpolation.

3.2.3.2. Bone modeling. Bone models are automatically deformed from the coarse level to the fine level using multi-resolution shape

Table 2
Summary of the MRI protocols.

Sequence	TR (ms)	TE (ms)	FOV (cm)	Matrix	FA (°)	Res. (mm)
A. Hi-res. axial 2D T1 TSE	578	18	40	512 × 512 × 250	90	0.78 × 0.78 × (2–10)
B. Low-res. axial 3D T1 (VIBE)	4.15	1.69	35	256 × 256 × 150	10	1.37 × 1.37 × 5
C. High-res. sagittal 3D T2' (TrueFISP)	10.57	4.63	20	384 × 384 × 112	28	0.52 × 0.52 × 0.6
D. Radial 2D intermediate weighted TSE	2180	13	16	384 × 384 × 18	180	0.41 × 0.41 × 10°

memory forces, non-penetration constraints (multi-resolution collision handling) and image-based external forces (NCC measure and GV gradient maximization). We have found that setting the total contribution of external forces to 30% of the total forces led to good results. During the process, the affine regularization is applied to external forces with a decreasing contribution ($\lambda = 0-0.7$). During the process, the search distance for computing intensity profile-based forces is progressively reduced from 1 cm (0.5 cm outside and 0.5 cm inside) to 0. The full process, including both the femur and the pelvis lasts about 1 min. A comparison with manually segmented bones from 13 high-resolution datasets shows an average distance of 1.7 ± 1.6 mm. The use of multi-resolution forces is effective as it prevents from large displacements when using fine resolutions (the systematic difference in the error is about 0.5 mm). Fig. 12 illustrates bone matching on images from the high-resolution and low resolution protocols.

3.2.3.3. Cartilage, ligament and muscle initialization. Standard bone coordinate systems (Wu et al., 2002) and soft-tissue attachment locations are automatically wrapped from the generic models. Then, soft-tissue models are initialized based on the joint angles using a skinning algorithm adapted from Kalra et al. (1998).

3.2.3.4. Cartilage and ligament modeling. Cartilages and ligaments are thin periarticular tissues. Despite of the highest resolution achievable with conventional MRI (~ 0.5 mm), the use of image information is not sufficient to allow an accurate reconstruction and geometric constraints start to be more reliable. For the hip joint, we know from anatomy literature that femoral and acetabular cartilages fill the inter-articular space, the interface being very smooth (Shepherd and Seedhom, 1999). The labrum links the femoral head and the acetabular rim to which it is attached. Ligaments have a quasi-constant thickness. Using our methods, it is possible to enforce such constraints: stable and accurate contact management enforces inter-organ constraints while the radial representation provides a way to control the thickness. To ensure quasi-constant thickness, we smooth medial axis radii locally. To ensure a perfect contact between cartilages and between the labrum and femoral cartilage, we add forces based on vertex projection (with $\alpha = 0.2$). Other forces are divided into smoothing and radial forces. We uniformly fix ligament thickness to 3 mm. Fig. 13 illustrates the reconstruction of the two hip cartilages, the labrum and the ilio/pubo/ischiofemoral ligaments (10,584 vertices in total). The medial representations are made of 1216 vertices. The fitting is performed sequentially: cartilages are deformed first, then the labrum, then ligaments. At each step, the previous models are constrained with infinite mass. The complete process is performed in ~ 3 min with a minimum frame rate of 10 frames/s. Thus, interaction (constraint point placement) is possible during the segmenta-

tion process. Because cartilage and ligaments are difficult to identify in MRI, our geometrical-based algorithm is difficult to validate. We have found a good agreement between 3D models and anatomical structures, but a quantitative validation would be required from more suitable images (i.e. arthro-MRI using contrast agent injected inside the capsule). The advantages of our algorithm are the low computational time and the flexibility. It can be easily tuned with different default thickness, different attachment areas and with user constraint points.

3.2.3.5. Muscle modeling. We have modeled the 21 skeletal muscles of the hip and thigh (71,328 model vertices and 4272 medial axis vertices in total). To perform muscle segmentation, we can use a different amount of prior information. The first level consists in using segmented bones to perform attachment wrapping and skinning as previously described (steps A and B of Fig. 14 and Table 3). In a second level, generic muscle shapes and muscle relative positions can be used to approximate patient-specific models based only on internal forces (the prior information is the medial axis shape parameters). Permanent contact are enforced with a proximity distance of 10 mm (step C). At the third level, we introduce image forces (combination of gradient and intensity profile registration constraints) (step D). During the whole segmentation process, we progressively increase the number of degrees of freedom while reducing the magnitude of admissible deformation (see Fig. 14): we successively increase model level of details, we reduce image force search space and the global regularization parameter, and we update, at the end of the segmentation, generic shape parameters through medial axis radii updates. Until now, our muscle segmentation method is fully automatic. But a fourth level of prior information can be added from the user. Indeed, it is possible to interactively place constraint points on the images, to get a faster matching and a more accurate segmentation (the frame rate is always >1 frames/s, which allows interaction). A medical student in our group has used this interactive segmentation to accurately segment 7 high-resolution datasets. It has provided gold standard datasets for validating our automatic segmentation method. Typically, about 2000 constraint points are necessary (about 30 min of manual work). But a little number of points can prevent from reaching most of the local minima. In Table 3 (method *constraint points*), we evaluate the placement of about 100 points during the fitting process. This Table presents the average error of our method over the seven different subjects (3 males and 4 females, different in terms of size and origin). Tests on two low-resolution MRI show a comparable accuracy than with high resolution MRI (see Fig. 15). However, results are biased because gold standard models are less accurate due to a more difficult manual delineation of muscles (partial volume effects). The overall computation time for the automatic method is around 8 min on a standard PC, for which 3/4 of the time is spent on external forces computation.

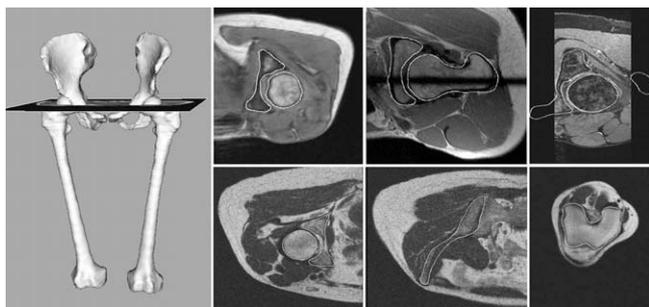


Fig. 12. Left: Individualized 3D bone models; Right: Sample slices with superimposed models generated by our quasi-automatic segmentation method (Top: low-resolution protocol, Bottom: high-resolution protocol).

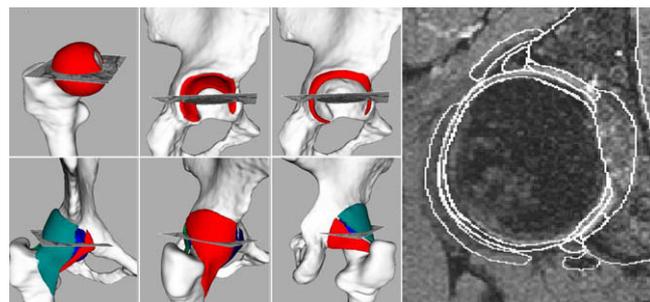


Fig. 13. Automatically segmented periarticular tissues.

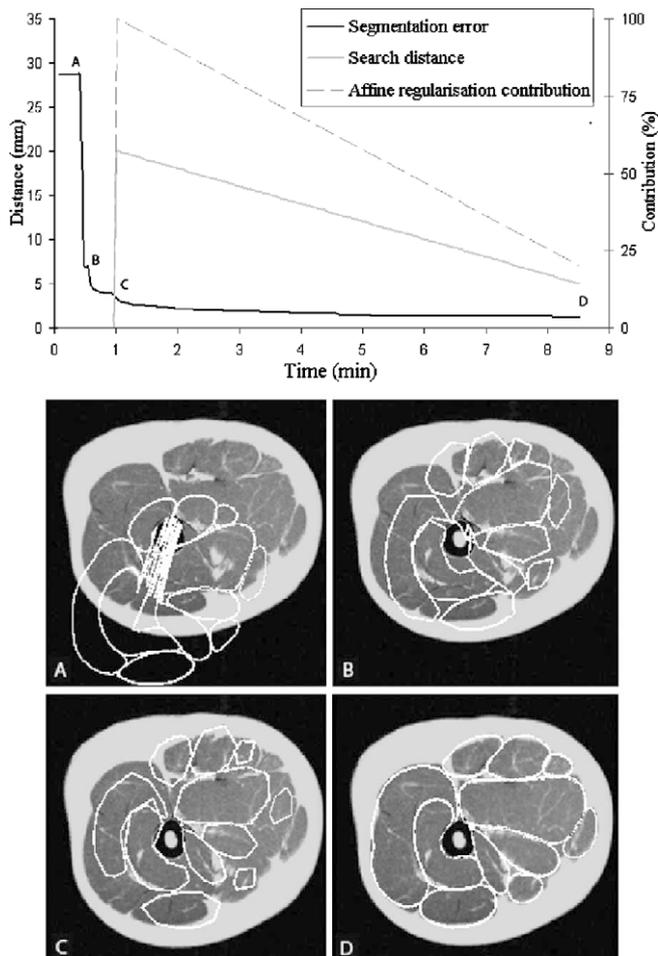


Fig. 14. Error variation during automatic muscle segmentation and segmentation results at different steps of the process.

3.2.4. Shape analysis

Musculoskeletal shape characterization is important for anthropometric comparison between individuals, and deformation analysis (temporal and longitudinal studies). For this purpose, our method registers anatomical features through shape and topological constraints. High-level descriptors such as the medial axis convey more information than local descriptors (curvature). Soft-tissue thickness can be simply analyzed through medial axis radii comparison, as shown in Fig. 16. Using approximated geodesic distances to attachments (Dijkstra algorithm), we compute normalized coordinates X and Y along medial surfaces, from which a thickness profile can be extracted (maximum radius in Y direction). For some muscles showing thickness steep changes, tendons lengths (which is an important biomechanical parameter) can be automatically extracted (see Fig. 16).

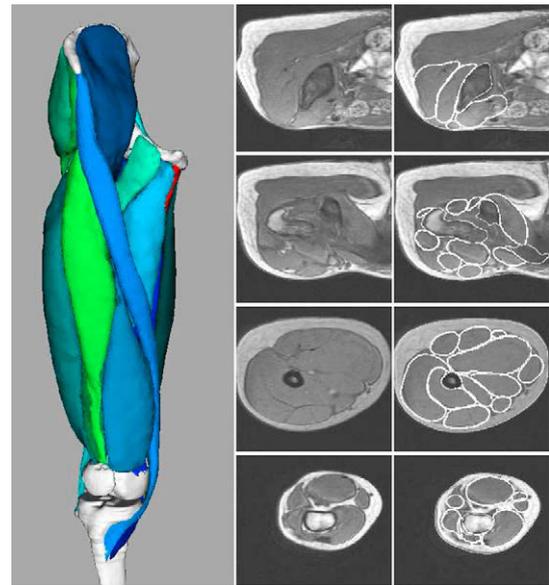


Fig. 15. Individualized muscle models and segmentation on low-resolution MRI sample slices.

Using the same methodology than for inter-patient registration, we performed intra-patient registration on some kinematic data. In a first step, we reconstructed models from scans in the neutral posture (volunteers in supine position). These models were subsequently registered in images from another posture. Bones were rigidly registered using a rigid regularization on external forces, while additional prior knowledge (volume preservation constraints) has been applied for soft-tissues. To compare the two postures, we have computed the compression/elongation of the organs based on the medial axis. This is done by calculating, for each medial vertex, the sum of the geodesic distances to the attachments (lengths L_r and L , respectively for the neutral and non-neutral posture). The compression/elongation percentage is then given by $100 \times (L - L_r) / L_r$. Visualizations of these ratios are shown in Fig. 17. For a small abduction, gluteal and adductor muscles are the ones that deform the most (up to 10% in length change) which is in agreement with functional anatomy knowledge. For a split posture, posterior muscles and the iliofemoral ligament are deformed up to 30%.

3.3. Performance tests

We simulate all muscle models of the hip and the thigh and perturb the system by simulating a global movement (20° of flexion). We check how vertices recover their initial positions under the action of medial axis shape forces and collision handling. Medial axis are driven by shape parameters, so stiffness are turned to 0.9 and 0.1 for shape memory and radial constraints respectively. Model forces are set to 0.6, 0.3 and 0.1 for radial, smoothness and volume

Table 3
Muscle segmentation results.

Prior knowledge	Method	Error \pm std. dev. (in mm)	Comp time (in s)
Individual bones + generic muscle models	A: Attachment wrapping	24.69 \pm 22.59	0
Joint transforms	B: Skinning	5.75 \pm 5.68	1
Generic shape parameters	Internal forces	5.48 \pm 5.82	15
Generic inter-object proximities	C: MA-MA collision handling	3.79 \pm 3.68	20
Generic intensity profiles	D: External forces	1.58 \pm 1.92	500
User inputs	Constraint points	1.37 \pm 1.56	500
	Gold standard	0 \pm 0	2000

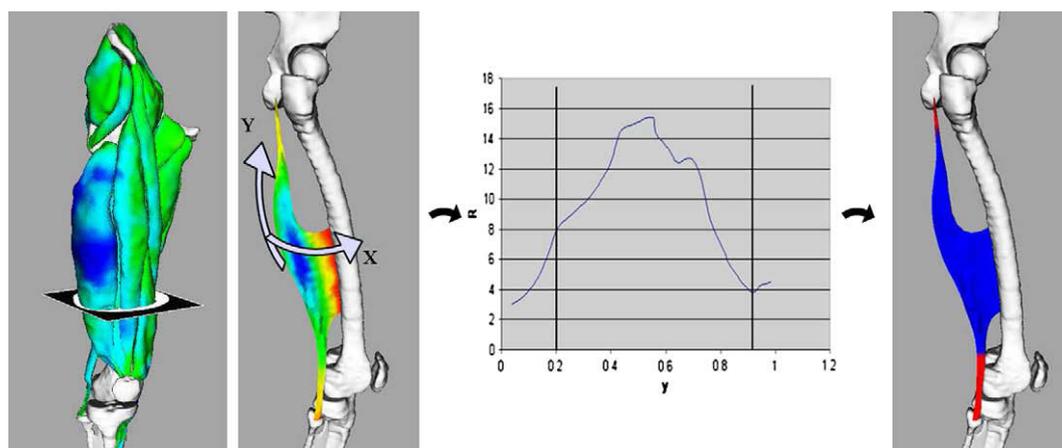


Fig. 16. Medial axis-based morphological analysis. Left: muscle thickness comparison between two subjects (Red: smaller; Green: equal; Blue: larger thickness). Middle: thickness profile analysis along the geodesic direction Y . Right: tendon extraction (in red) from the thickness profile. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

conservation forces respectively. Convergence is reached after about 2500 iterations corresponding to 160 s of computation. The different resolutions are applied sequentially (computation times per iteration are respectively 0.25 s, 0.06 s and 0.02 s). The computational time repartition is 12% for collision handling, 32% for force computation, 39% for the numerical integration and 17% for the visualization. The final error (distance to the initial model) is 0.8 ± 1.0 mm which roughly corresponds to the medial axis approximation error. We test the performance of our collision detection scheme by monitoring convergence times (time spent on collision handling), accuracy (average distances with regards to the most accurate collision handling scheme) and stability (remnant average vertex displacement) using different resolutions and collision detection features (standard *model-to-model* collision detection, the *model-to-medial axis* scheme where each model point is tested towards medial axis and the *medial axis-to-medial axis* method). We compare the *full* detection with the *permanent contact* scheme (where initial relative positions are enforced as presented in Section 2.3.1). The hypothesis of permanent contact speeds up collision detection by 3. It is valuable to make it for muscle registration as we do not expect large sliding between them due to fascia. The full detection is useful as an initial step for gen-

eric model reconstruction and when sliding is expected (e.g. bone, cartilage and ligaments registration). The multi-resolution collision detection improves convergence speed by a factor of 3–10 without prohibitive inaccuracies. Medial axis-based collision detection clearly improved performances: model/axis collision detection approximately divides computational costs by 2, while axis/axis detection divides them by 40.

4. Conclusion

Within a discrete deformable model framework, we have presented scalable methods (mesh topology adaptation, multi-resolution scheme and medial axis) for controlling shapes and deformations. Our work represents one the first attempt to partly model the musculoskeletal system from medical images in a cost-efficient way. Most existing methods are interactive, do not perform registration, and/or make use of non-clinical images such as the visible human dataset. Although more validation and extensions to handle pathologies are required, we believe that our method has a good potential for clinical use: we have tested our method with standard clinical images, the accuracy achieved is in the range of image resolution and the overall segmentation time is less than an hour. Moreover, our method can be ran in a quasi-automatic setting, meaning that only a few points are needed for the initialization (≈ 2 min of manual work). If needed, the user can interact during the segmentation to speed up the registration and/or prevent from convergence to local minima. We have demonstrated that our methodology allows representing the variety of shapes, topology and deformations of the hip joint complex. We expect that it can be generalized to other joints and organs. In most cases, our segmentation method has proven to be valuable to estimate organ shapes from images. However, it is not always robust in presence of false or fuzzy edges, and when anatomical variability from the generic model is too high. We believe that we could get a higher accuracy through a deeper study of intensity profile forces (combination of several metrics, weighting according to profile relevance, etc.). Statistical methods based on shape, deformation and intensity profile examples could also improve the results significantly, but it is currently not clear how to explore the potentially large example space (variable subjects, postures and actuation levels). We did not individualize attachments since they mostly rely on bone geometrical features. Visually, it was difficult to identify attachment locations on images, and we think that an individualization technique would require a different imaging protocol. Our

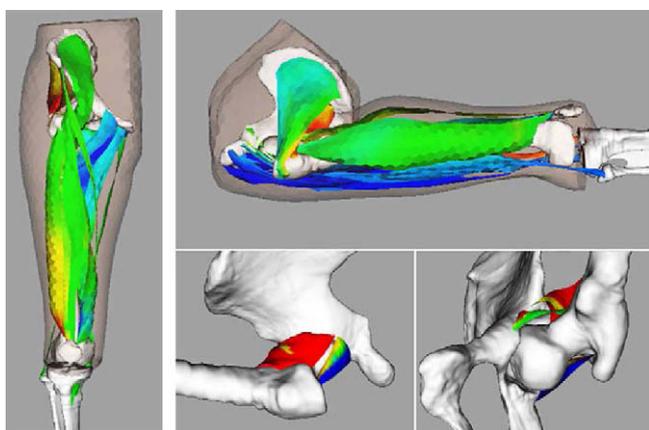


Fig. 17. Left: muscle medial surfaces during low amplitude abduction. Right: muscles and ligaments medial surfaces during split. Colors represent the elongation percentage (clamped at -10% (red) and $+10\%$ (blue) except for the top right figure where the clamping values are $\pm 30\%$). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

technique has a good potential for studying morphological and functional pathologies through static, dynamic and longitudinal studies. However, extensive tests on patient data need to be performed, and pathology models (e.g. lesions, cartilage calcification) need to be developed to detect them. Another perspective is the validation and the parameterization of functional biomechanical muscle models. This would require the addition of physically-based constraints in our framework (and therefore the incorporation of fiber geometry), along with actuation measurements.

Acknowledgments

This work is supported by CO-ME (Computer Aided and Image Guided Medical Interventions) project funded by Swiss National Research Foundation. We would like to thank Dr. Kolo-Christophe, Dr. N'Guyen and Dr. Sadri from the Geneva University Hospital for their collaboration.

References

- Amenta, N., Choi, S., Kolluri, R., 2001. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications* 19 (2–3), 127–153.
- Attali, D., Boissonnat, J., Edelsbrunner, H., 2007. Stability and computation of medial axes: a state of the art report. In: *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*.
- Aubel, A., Thalmann, D., 2001. Interactive modeling of the human musculature. In: *Proceedings of Computer Animation 2001*.
- Baraff, D., Witkin, A., 1998. Large steps in cloth simulation. In: *Proceedings of SIGGRAPH98, Computer Graphics*, vol. 32, pp. 106–117.
- Blemker, S., Delp, S., 2005. Three-dimensional representation of complex muscle architectures and geometries. *Annals of Biomedical Engineering* 33 (5), 661–673.
- Blemker, S., Asakawa, D., Gold, G., Delp, S., 2007. Image-based musculoskeletal modeling: applications, advances, and future opportunities. *Journal of Magnetic Resonance Imaging* 25 (2), 441–451.
- Blum, H., 1964. A transformation for extracting new descriptors of shape. In: *Symposium on Models for the Perception of Speech and Visual Form*. M.I.T. Press, pp. 139–146.
- Delingette, H., 1994. Simplex meshes: a general representation for 3d shape reconstruction. In: *Proceedings of CVPR'94*, pp. 856–857.
- Delp, S., Loan, J., 1995. A graphics-based software system to develop and analyze models of musculoskeletal structures. *Computers in Biology and Medicine* 25 (1), 21–34.
- Fripp, J., Crozier, S., Warfield, S., Ourselin, S., 2007. Automatic segmentation of articular cartilage in magnetic resonance images of the knee. In: *Proceedings of MICCAI'07*, vol. 2, pp. 186–194.
- Ghanei, A., Soltanian-Zadeh, H., Windham, J., 1998. A 3d deformable surface model for segmentation of objects from volumetric data in medical images. *Computers in Biology and Medicine* 28, 239–253.
- Holden, M., Hill, D., Denton, E., Jarosz, J., Cox, T., Rohlfing, T., Goodey, J., Hawkes, D., 2000. Voxel similarity measures for 3-d serial MR brain image registration. *IEEE Transactions on Medical Imaging* 19 (2), 94–102.
- Horn, B., Hilden, H., Negahdaripour, S., 1988. Closed-form solution of absolute orientation using orthonormal matrices. *Journal of the Optical Society of America* 5, 1127–1135.
- Kalra, P., Magnenat-Thalmann, N., Moccozet, L., Sannier, G., Aubel, A., Thalmann, D., 1998. Real-time animation of realistic virtual humans. *Computer Graphics and Applications* 18 (5), 42–56.
- Kaptein, B., der Helm, F.V., 2004. Estimating muscle attachment contours by transforming geometrical bone models. *Journal of Biomechanics* 37, 263–273.
- Lorigo, L.M., Faugeras, O., Grimson, W., Keriven, R., Kikinis, R., 1998. Segmentation of bone in clinical knee MRI using texture-based geodesic active contours. In: *Proceedings of MICCAI'98*, vol. 1496, pp. 1195–1204.
- Lötjönen, J., Mäkelä, T., 2001. Segmentation of mr images using deformable models: application to cardiac images. *International Journal of Bioelectromagnetism* 3 (2), 37–45.
- Lötjönen, J., Reissman, P., Magnin, I., Katila, T., 1999. Model extraction from magnetic resonance volume data using the deformable pyramid. *Medical Image Analysis* 3 (4), 387–406.
- Maintz, J., Viergever, M., 1998. A survey of medical registration. *Medical Image Analysis* 2 (1), 1–36.
- McInerney, T., Terzopoulos, D., 1995. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4d image analysis. *Computerized Medical Imaging and Graphics* 19 (1), 69–83.
- Mezger, J., Kimmerle, S., Etmuss, O., 2003. Hierarchical techniques in collision detection for cloth animation. *WSCG* 11 (2), 322–329.
- Mirtich, B., Canny, J., 1994. Impulse-based dynamic simulation. *The Algorithmic Foundations of Robotics*.
- Montagnat, J., Delingette, H., 2001. A review of deformable surfaces: topology, geometry and deformation. *Image and Vision Computing* 19 (14), 1023–1040.
- Montagnat, J., Delingette, H., 2005. 4D deformable models with temporal constraints: application to 4D cardiac image segmentation. *Medical Image Analysis* 9 (1), 87–100.
- Ng-Thow-Hing, V., 2000. Anatomically-based models for physical and geometric reconstruction of humans and other animals. Ph.D. Thesis, Department of Computer Science, University of Toronto.
- Park, J., McInerney, T., Terzopoulos, D., Kim, M., 2001. A non-self-intersecting adaptive deformable surface for complex boundary extraction from volumetric images. *Computers and Graphics* 25 (3), 421–440.
- Pizer, S., Fletcher, P., Joshi, S., Thall, A., Chen, J., Fridman, Y., Fritsch, D., Gash, A., Glotzer, J., Jiroutek, M., Lu, C., Müller, K., Tracton, G., Yushkevich, P., Chaney, E., 2003. Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision* 55 (2), 85–106.
- Press, W., Teukolsky, S., Vetterling, W., Flannery, B., 1992. *Numerical Recipes in C, second ed.* Cambridge University Press.
- Sheffer, A., Kraevoy, V., 2004. Pyramid coordinates for morphing and deformation. In: *Proceedings of Second International Symposium on 3DPVT*, pp. 68–75.
- Shepherd, D., Seedhom, B., 1999. Thickness of human articular cartilage in joints of the lower limb. *Annals of the Rheumatic Diseases* 58.
- Skerl, D., Likar, B., Pernus, F., 2006. A protocol for evaluation of similarity measures for rigid registration. *IEEE Transactions on Medical Imaging* 25 (6), 779–791.
- Szeliski, R., Lavallée, S., 1996. Matching 3-d anatomical surfaces with non-rigid deformations using octree-splines. *International Journal of Computer Vision* 18, 171–186.
- Teran, J., Sifakis, E., Blemker, S., Hing, V.N.T., Lau, C., Fedkiw, R., 2005. Creating and simulating skeletal muscle from the visible human data set. *IEEE TVCG* 11, 317–328.
- Terzopoulos, D., Fleischer, K., 1988. Deformable models. *The Visual Computer* 4 (6), 306–331.
- Terzopoulos, D., Witkin, A., Kass, M., 1987. Symmetry-seeking models and 3d object reconstruction. *International Journal of Computer Vision* 1, 211–221.
- Teschner, M., Kimmerle, S., Heidelberger, B., Zachmann, G., Raghupathi, L., Fuhrmann, A., Cani, M.-P., Faure, F., Magnenat-Thalmann, N., Strasser, W., Volino, P., 2005. Collision detection for deformable objects. *Computer Graphics Forum* 24 (1), 61–81.
- Volino, P., Magnenat-Thalmann, N., 2000a. Accurate collision response on polygonal meshes. In: *Proceedings of Computer Animation*, pp. 154–163.
- Volino, P., Magnenat-Thalmann, N., 2000b. Implementing fast cloth simulation with collision response. In: *Proceedings of the International Conference on Computer Graphics*, pp. 257–268.
- Wu, G., Siegler, S., Allard, P., Kirtley, C., Leardini, A., Rosenbaum, D., Whittle, M., D'Lima, D., Cristofolini, L., Witte, H., Schmid, O., Stokes, I., 2002. Isb recommendation on definitions of joint coordinate system of various joints for the reporting of human joint motion-part i: ankle, hip, and spine. *Journal of Biomechanics* 35, 543–548.
- Zitová, B., Flusser, J., 2003. Image registration methods: a survey. *Image and Vision Computing* 21 (11), 977–1000.