Towards the virtual physiological human

N. Magnenat-Thalmann [1], B. Gilles[1], H. Delingette[2], A. Giachetti[3] and M. Agus[3]

[1] MIRALab - University of Geneva, Switzerland
[2] INRIA - Asclepios, France
[3] CRS4 - Visual computing group, Italy

Abstract

The objective of this tutorial is to train students and researchers in the various domains involving the modelling and simulation of the human body for medical purposes. Human body representations have been used for centuries to help in understanding and documenting the shape and function of its compounding parts. Nowadays, medical acquisition devices especially medical scanners are able to produce a large amount of information, such as high-resolution volumes, temporal sequences or functional images, more-and-more difficult to analyse and visualise. In other words, we measure much more than we understand. In this context, higher-level information such as anatomical and functional models is increasingly required to support diagnosis and treatment. Three levels of complexity can be distinguished (geometry, function and control) where modelling and simulation methods take place. In this tutorial, we will present the current research issues towards the patient-specific reconstruction of virtual models and their functional simulation. We will focus on the computer graphics aspects involved in medical modelling/ simulation: deformable model-based segmentation, mesh optimisation, data fusion, interactive mechanical simulation, cost-efficient visualisation, etc. Examples will be given in the musculoskeletal, cardiac and neurological domains. All speakers are partners of the EU project "3D Anatomical Human" led by MIRALab - University of Geneva.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations. I.3.5 [Computer Graphics]: Physically based modelling. I.3.3 [Computer Graphics]: Viewing algorithms. I.3.8 [Computer Graphics.]: ApplicationsComputer-aided diagnosis and treatment.

1. Anatomical and kinematical reconstruction from medical data

1.1. Introduction

Nowadays, medical imaging devices are able to produce a large amount of information, such as high-resolution volumes and temporal sequences, more-and-more difficult to analyse and visualise. In this context, higher-level information such as 3D or 4D models is increasingly required to support medical diagnosis. Such models can have predictive capabilities (for computer-aided treatment, prevention, rehabilitation and ergonomics). For simulation, the use of medical images is important, as they allow the production of patient-specific models (anatomical models), expressing subtle inter-subject variability (possible cause of pathology), and the validation of mechanical prediction through temporal studies (kinematical models).

1.2. Acquisition modalities

When modelling and simulating real objects, data acquisition is essential. It aims at extracting model parameters and validating simulation output. Data can be divided into geometric data (e.g. medical images), kinematical data (e.g. optical motion capture data), mechanical parameters (e.g. strain-stress curves) and physiological data (e.g. muscle activation). Due to measurement difficulties (for instance, the access to bone kinematics, internal muscle activation patterns) and musculoskeletal system complexity, it is common to make assumptions about this data and use prior knowledge. Relying on these, computer models need to be validated correctly. Obviously, geometric models (shape) deal with few assumptions, whereas functional models require a lot of prior information for their simulation. In this section, we focuss on musculoskeletal macroscopic acquisition,

which is more relevant for the computer graphics community.

1.2.1. Static acquisition

1.2.1.1. Cadaveric studies: Until XX century, cadaver analysis has been the only one way for assessing organ shapes and interrelationships (e.g. Gray's anatomy [Gra00]). The visible human [NLM99], a mixture of medical images and histological cuts of the same individual, is up to now the most complete dataset for anatomical description. It has supported many projects in the fields of medical image processing and biomechanics. The measurement of submillimeter anatomical features such as in synovial joints still requires specimens. However, medical imaging modalities, more and more accurate and standardised, tend to supplant these kind of studies.

1.2.1.2. Radiography and Computed Tomography (CT): Historically, the radiography has been the first modality for in-vivo anatomical imaging of internal structure (Roentgen, end of XIX century). It is based on the photo measurement of the intensity of X-rays traversing the body. Sensors are generally phosphor plates, that can be subsequently digitalised (Computed Radiography - CR), but new digital sensors are emerging (Digital Radiography - DR), though they are still expensive. The resulting 2D projections highlight anatomical structures, absorbing X-rays differently. Combining several projections (Computed Tomography - CT), volumes (typically 512x512x128 voxels) can be reconstructed with a resolution of about 0.5x0.5x1 mm in 30s. CT is a tomographic modality, meaning that output images, reconstructed from several projections, are cross-sections. The intensity in images is measured towards the Hounsfield scale, leading to a simple segmentation of tissues. However, different tissues can share the same value, and artefacts disturb this identification. It is known that CT is more dedicated to chest and bone investigations (about 70% of all imaging examinations), as bone-to-soft tissue contrast is high. Thresholding, few regional constraints and simple post-processing techniques (e.g. morphological operations) are generally sufficient to segment bones, or structures highlighted with suited contrast agents (e.g. angiography). Direct volume rendering is also popular in the clinical environment [HPP*03]. In orthopaedics, where most of pathologies rely on bones (e.g. fracture, cartilage ossification), CT is the privileged modality, although it is invasive (high dose X-ray absorption can damage cells and cause cancer).

1.2.1.3. Ultra-sound (US): Ultra-sounds propagate in homogeneous tissues but are partly reflected at tissue interfaces. The principle of US is to measure this reflection in time, the emission/ reception time interval being proportional to the depth of the interface (assuming few wave velocity variations). US scanners generally operate in brightness mode (B-mode), that is grey-scale in images represent an amplitude map of the reflected sound. 2D cross-sections and 3D volumes can be interpolated by casting rays in different directions (tomography). Either mechanical and free-hand probes are available [FDC01], allowing for some of them position tracking. General interpolation techniques may be used to resample unparallel slices. US is low-cost, fast and non-invasive. Unfortunately, image quality is low: speckle is present due to wave interference, geometric distortions occur due to wave velocity changes and ultra-sound can be totally absorbed by some tissues (e.g. air, bones) and at large depth, leading to occlusions. US is typically used for heart and foetus dynamic imaging, but is also adequate for superficial musculoskeletal imaging. Particularly, muscle fibers direction can be visualised in-vivo.

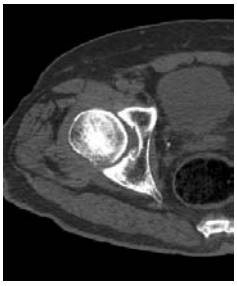

Figure 1: *CT imaging of the hip*

Figure 2: *Ultrasound assessment of the pennation angle of the gastrocnemius muscle in a relaxed state, from University of California, Davis*

1.2.1.4. Magnetic Resonance Imaging (MRI): Atomic nuclei have a magnetic moment that can be oriented by an external magnetic field. This moment rotates about the axis of the field, with a frequency depending on the type of nucleus and on the strength of the magnetic field. In this condition, atoms can briefly jump from one equilibrium state (spin $+1/2$) to another one (spin $-1/2$), under radio impulses emitted at this frequency (resonance). When returning to the initial state, atoms generate a damped radio frequency signal (relaxation). MRI aims at measuring this signal from the resonance of hydrogen atoms, contained in all tissues (water). Indeed, tissues react differently, depending on the concentration of water (proton density), but also on the way that protons revert back to their resting states after the initial RF pulse (relaxation time). Relaxation time is measured

in two directions: T1 (longitudinal) and T2 (transversal). In addition, fluid flow may decrease the signal. These four tissue parameters: proton density, T1 relaxation time, T2 relaxation time and fluid flow; determine the signal intensity on the MR images. By weighting these parameters, MRI is capable of producing various types of contrast, conferring to this modality a great versatility. The self-diffusion of water is restricted by membranes, and consequently relies on tissue geometry (e.g. fiber orientation). By applying different gradient directions, fluid flow/ fiber direction can be measured (six directions: diffusion tensor MRI; three directions: diffusion weighted MRI). Local, voxel-based directions can be combined with a fiber-tracking algorithm (tractography) to reconstruct complete paths [MZ02]. It is admitted that MRI does not produce any harmful ionizing radiation [SC04]. Typical resolution is about $1mm^3$. Strong magnets (>1 Tesla), which are horizontal tubes of about 60cm diameter, lead to a higher signal-to-noise ratio (SNR), and consequently to a higher resolution and shorter scan times than feeble magnets. In research, the spatial resolution can attain $1\mu m^3$, using powerful magnets (e.g. 4 Tesla, 7 Tesla).

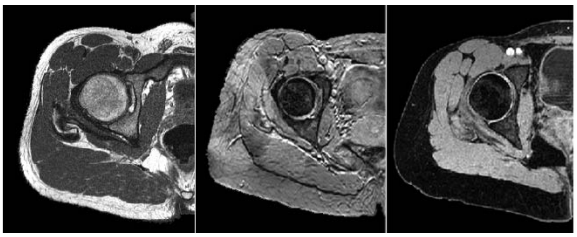

Figure 3: *Different MRI sequences for imaging the hip: a) Spin-Echo T1 b) Gradient-Echo T2* c) Gradient-Echo T1*

1.2.1.5. Body scanning: For computer graphics and animation purposes, accurate skin models of the complete body are generally digitalised using commercial optical scanners. The reflection of white light or eye-safe laser is measured by mobile cameras, providing a point cloud of the scene that need to be subsequently triangulated (resolution: $\simeq 1mm$, acquisition time: $\simeq 30s$). Another approach uses stereo photogrammetry from at least two camera views. This technique is faster ($\simeq 1.5ms$) and more accurate ($0.5mm$) than laser scanning when using efficient mesh reconstruction techniques, however this has not been applied to large scale (e.g. full body) yet.

1.2.2. Kinematics and dynamics acquisition

1.2.2.1. Ultra-sound (US): US is a fast imaging modality that can provide useful kinematic information, with a frequency up to 100 frames/s (e.g. foetal echography). 1D temporal analysis (x axis: time; y axis: depth), called M-mode, is useful for visualising highly sampled (about 1000 pulses/sec) superficial muscle contraction patterns or heart valves movements. Direct tissue/ liquid velocity visualisation is also possible with US, exploiting the Doppler effect. Colour Doppler aims at visualising flow direction and velocity, whereas color power Doppler (CPD) only deals with velocity amplitude.

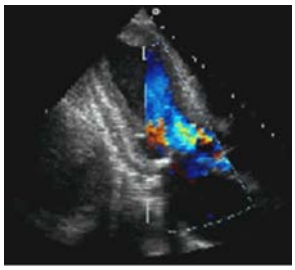

Figure 4: *B-mode ultrasound slice with superimposed colour Doppler showing blood flow inside the heart (from AnaesthesiaUK)*

1.2.2.2. Fluoroscopy and Computed Tomography (CT): Fluoroscopy produces real-time images of internal structures in a similar fashion to radiography. X-ray fluoroscopy is faster than CT, as it performs only one projection from a continuous input of X-rays, while CT requires multiple projections to reconstruct cross-sections. CT volumes are converted into 2D projections using digitally radiographs reconstruction (DRR) techniques. Fluoroscopy has been widely used for intra-operative bone tracking and biomechanical studies [YSAT01]. Implanted pins are used for accurate bone localisation, but non-invasive bone motion tracking can be obtained through 2D/ 3D registration between fluoroscopic images and CT data. New technical advances in CT scanner including large 2D detectors, fast rotating devices and fast systems for transmission and reconstruction, have recently permitted 4D CT acquisition (e.g. 256-slices, field of view along $z = 12cm$, resolution $= 1mm \times 1mm \times 0.5mm \times 1sec$ [MET*04]).

1.2.2.3. Magnetic Resonance Imaging (MRI): Several MRI techniques has been made available for dynamic tissue analysis, through the decrease of MRI acquisition time. Particularly, joint kinematic studies with MRI have been reported. The passive sequential (or stepped) acquisition technique (i.e. loaded stationary joint during acquisition) is well-established but has limited use. However, the problem of acquiring volumetric image data in real-time with MRI during active motion remains to be solved due to inherent trade-off in the MR imaging technique between Signal-to-Noise Ratio (SNR), spatial resolution and temporal resolution. Specific imaging sequence (trueFISP: $6frames/s$ [QLH*02]) and parallel imaging techniques exploiting spatiotemporal correlations with reference images (e.g. SENSE, Philips Medical Systems, $38frames/s$ [TBP03]) can greatly speed up the acquisition. Motion periodicity can be exploited by combining information across several cycles (*cine-MRI*), which

is particularly suited for cardiac imaging. Because tissues keep their magnetic orientation for a time (∼T1 time), it is possible to create temporary magnetic fiducial markers, or tags, and hence trace tissue deformations (*tagged MRI*) [ZPR*88]. Velocity imaging is possible using phase contrast sequences (*pcMRI*) [PHSE91]: two datasets are acquired, one being flow sensitive and the other flow compensated. The raw data (phase information) from the two data sets are then subtracted. pcMRI can be coupled with cine-MRI as in [DP94].

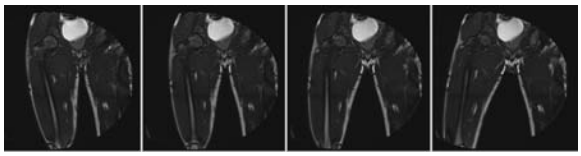

Figure 5: *Real-time dynamic MRI of the thigh [GPMTV04]*

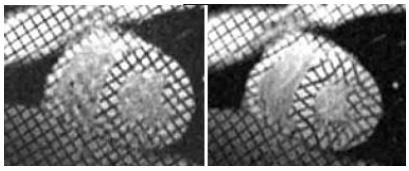

Figure 6: *Cine tagged-MRI of the heart*

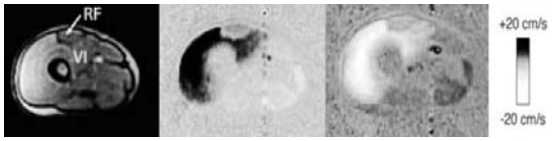

Figure 7: *Cine phase-Contrast MRI of the thigh during knee flexion/extension [ABGD02]*

1.2.2.4. Motion Capture (MoCap): While medical imaging provides precise but local kinematic measurements, motion capture (MoCap) techniques aim at acquiring global skin motion (larger scales and ranges of motion). Avoiding the burden inherent to medical scanners, MoCap provides in-situ kinematical data (e.g. gait), and is compatible to other modalities for force (e.g. force plates) and physiological (e.g. EMG) measurement. Motion capture of the complete body is achieved through the recording of landmarks attached to the skin. In a post-processing phase, landmark positions are reconstructed in 3D, and rigid motions of the different body segments can be established, approximation skeletal movements and joint centers [LO99]. Three main motion capture systems can be distinguished [Men00]: the ones based on electromagnetic trackers (*flock-of-bird*) ($\sim 140Hz$, $\sim 2mm$); the ones based on mobile electromechanical devices such as inclinometers, goniometers, accelerometers and gyroscopes ($\sim 90Hz$, $\sim 0.5deg/2mm$); and the most used ones based on reflective optical markers ($< 1000Hz$, ∼0.5mm). With optical systems, a post-processing (interactive) phase is necessary to manage occlusions. As the relative skin/ bone motion introduces errors in joint center and skeletal kinematics computation, implantable pins and error minimisation techniques have been used for biomechanical studies.

1.2.2.5. Video-based methods: Instead of MoCap, raw video images or laser scans can be used to assess body kinematics. Whereas MoCap provides accurate and high-rate motion of the skin at sample points, it can be useful to assess the skin with more complete spatial information. However, video-based methods have a lower accuracy and frame-rate, and require efficient post-processing methods (computer vision techniques). In [DHW*04], a dynamic body scanner system ($15Hz$) is used for spine dynamic analysis where internal anatomical structure are estimated for the shape of the skin. Commercial products using stereo photogrammetry can achieve around $60Hz$ for imaging body parts. In [ACP02], skin animation is performed from scattered scan data, trough interpolation techniques.

1.2.2.6. Force plates: For dynamic simulations of the musculoskeletal system, it is necessary to evaluate external forces acting on the body, such as ground reaction forces. MoCap, EMG and force plates are typically associated. Force plates can record ground reaction forces (6 DOF) [DP02]. Alternatively, one can use pressure measurement soles [WC03], leading to slightly different results (shoe/ground vs. foot/shoe interaction). Both systems can record at approximately $500Hz$, with a resolution of about $1cm$.

1.2.3. Physiology acquisition

1.2.3.1. Electromyography (EMG): The measurement of muscle stimulation patterns is essential to understand musculoskeletal dynamics and control, and consequently the causes of abnormal movements. Electromyography (EMG) consists in assessing electrical signals ($0-6mV$, $10-500Hz$) being generated by a muscle at the same time that it contracts [Cut93]. Surface EMG (sEMG), where electrodes are placed at the skin above muscles, is a non-invasive approach (contrary to needles) for muscle stimulation recording. It has been used in most studies regarding the forward dynamic simulation of the musculoskeletal system, the validation of inverse dynamic simulation, and the diagnosis of neuro-musculoskeletal disabilities [DL95]. Due to differences in electrode properties, configuration and location, which is problematic for the use and exchange of sEMG data, standard protocols have been recently established [HFDKR00]. Indeed, a good electrode/skin contact need to be insured to reduce noise, and a precise electrode placement to minimise crosstalk from neighbouring muscles. To tackle this, researchers have investigated

more sophisticated electrodes such as bipolar electrodes and more recently electrode arrays. EMG data need to be post-processed to extract relevant information and to allow intra and inter subject comparison [DeL97]. Because there is no absolute signal value, EMG are generally normalised according to the maximum voluntary contraction or MVC (highest peak during maximum effort). In addition, EMG analysis in the Fourier domain is particularly useful to separate the contribution from the different muscles, and analyse physical fatigue. Indeed, fatigue is characterised by an average frequency shift towards lower frequencies.

1.2.3.2. Mechanomyography (MMG): MMG is a non-invasive technique for measuring muscle contraction speed of muscle groups, and is the mechanical equivalent to EMG (electrical activity). Its principle is to measure skin displacement under neuromuscular stimulation. Measurements are commonly made through a variety of transducers (condenser microphones, piezoelectric contact sensors or accelerometers) leading to various output names: muscle sound, soundmyogram, acousticmyogram, phonomyogram, accelerometermyogram or vibromyogram.

1.2.4. Mechanical testing

1.2.4.1. Mechanical devices: Biological tissues mechanical parameters can be established from various modalities measuring the uniaxial/biaxial tissue strain response in relation to mechanical or ultrasound impulses. Indeed, the relationship between the strain (tension, compression, bending or torsion) and the stress (pressure or force) provides the elastic property of the material, while viscoelasticity can be analysed from stress and strain rates (stress-relaxation and creep testing) over multiple loading/unloading cycles. The biaxial stress-strain relation of a soft tissue can not be obtained by superimposing two individual uniaxial test results due to the large deformation, but data regression methods are usually performed to fit the constitutive model to data. Commonly, tissue specimens need preconditioning (several loading/unloading cycles) before obtaining reproducible results. Then, parameters are established from the loading and unloading curves. The use of dead specimens is problematic, as they may behave differently than living tissues: they need to be correctly hydrated (e.g. saline baths) and kept at a suitable temperature. In addition, tissue initial milieu is difficult to reproduce (e.g. pressurised synovial joints). In-vivo mechanical measurement of internal tissues poses extreme experimental difficulties, but is becoming possible through the development of miniaturised devices. Buckle transducers, telemetered pressure sensors and strain gauges have been investigated for in-vivo force, pressure and deformation measurement. Non-invasive approaches are also available for isotropic materials such as aspiration tools [NMK*03] or torsional resonator devices [VM04]. Highly anisotropic materials (e.g. ligaments) are, however, difficult to measure in situ. Implantable force transducers and video-based measurement techniques have been used, but the reference state (zero stress) remains problematic to determine [WG01].

1.2.4.2. Magnetic Resonance Elastography (MRE): MRE is a new technique to measure material internal stiffness in vivo and non-invasively [MRL*96]. It is based on the measurement of tissue response under shear waves, propagated from a transducer attached to the skin surface. The stiffer is the tissue, the more rapidly waves are propagated (longer wavelength). Wave displacement velocity is measured with Phase-Contrast MRI (pcMRI), whereas in ultrasound electrography, the measurement is based on Doppler techniques. MRE has been applied to measure muscle actuation, exploiting the fact that the more contracted is a muscle, the stiffer it is. In [PBHS05] muscle anisotropy (3 parameters) has been measured with MRE, based on a transversely isotropic linear material. MRE is a very promising approach allowing tension distribution within a muscle to be imaged. In addition, contrary to EMG, MRE can provide information about passive properties of muscle.

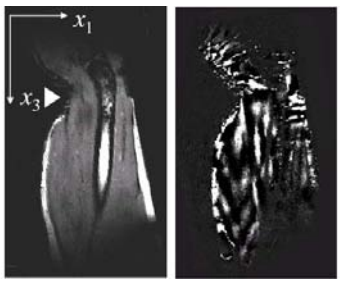

Figure 8: *MRE measurement based on a harmonic shear excitation in x1-direction, from [PBHS05]*

1.3. Image segmentation and registration

1.3.1. Introduction

Anatomical and kinematical modelling deal with image segmentation, that is the identification, in 2D, 3D or 4D images, of the region or contour of interest. Segmentation has been an intensive research field for many years, and is generally associated to registration (the search of spatial correspondences across datasets). Indeed, registration can indirectly perform segmentation as soon as a generic/ reference dataset has been segmented (the generic model is iteratively deformed/ registered to images). On the contrary, direct segmentation can be achieved with few assumptions about the problem (organ shape, image intensities, deformation, etc.). Direct approaches such as thresholding, region growing or edge detection are generally ad-hoc techniques that fail to handle noisy images with a large amount of textural information such as medical images, while indirect approaches use prior information (reference segmented datasets, generic models, anatomical and statistical information, texture database) to perform the segmentation.

Given ordinary clinical images (low-resolution images with anisotropic voxels, noise and partial volumes), the output of segmentation (binary volume), even manually obtained, is a more-or-less correct representation of organ shapes. A post-processing step is necessary to reconstruct a model, from the binary map, that can be simulated and efficiently visualised. The Marching Cubes algorithm, originally proposed by Lorensen et al. [LC87], is considered to be a standard approach to the problem of extracting isosurfaces from a volumetric dataset. However it produces aliased meshes, with a poor geometrical quality (irregular triangles) and an excessive number of triangles. Constrained reconstruction techniques (such as [Del99]) where surface smoothness and mesh topology can be controlled are more efficient. In this context, deformable models are nowadays a widely spread approach for segmentation, since they allow geometrically constrained image segmentation (no post-processing).

Figure 9: *marching cubes reconstruction of the knee (left) and constrained reconstruction using deformable models (right)*

1.3.2. Direct segmentation

Direct segmentation of anatomical structures in medical images is particularly challenging because medical images are noisy (e.g. speckle in ultrasound) and generally complex from a textural point of view. Moreover, the ratio between image resolution and the size of structures to segment can be high. Low resolution images produce partial volume effects, making frontiers between organs unclear. Direct segmentation makes very little assumption about organ shape, organ size or image intensity, so is very generic. For complex problems, direct segmentation is noise-sensitive, not robust and quite inaccurate. However it can help in speeding up manual segmentation and enhancing/ extracting image features for further algorithms (e.g. registration).

Generally speaking, direct segmentation is an iterative process involving a "detection" step where regions are identified in images and a *classification* step where regions are combined to create new regions. A pre-processing step may initially filter images to normalise them, reduce artefacts and noise and enhance specific features. After segmentation, the resulting atlas can be converted into surfacic or volumetric model through reconstruction techniques.

1.3.2.1. Pre-processing: Well-known filtering algorithm may be applied to images in a pre-processing step. Noise can be attenuated through low pass filtering. Anisotropic filtering [PM90] is particularly suited, as it does not alter edge sharpness. When segmenting several datasets from different patients or sessions, intensity normalisation may be applied to get more robust results (because most algorithms rely on image intensity). This is frequent to fit image contrast and brightness to a reference image. Normalisation is generally performed by an (piece-wise) affine transform of the intensities. Between cross-sections, intensity bias can be present in medical volumes and attenuated through intensity normalisation across slices.

1.3.2.2. Region detection: Regions in images can be defined by their internal voxels or by their boundaries. Similarly, segmentation methods aim at discriminating regions, looking at image intensity levels and their derivatives. Another way is to assess the probability of voxels to belong to a region, considering local intensity values.

- **Edge-based approaches:** Organ interfaces are characterised by a more or less steep intensity level variation in images and are quantified through gradient amplitude and orientation. Various convolution filters have been proposed for gradient measurement. Among them, 1st order derivative filters such as the Sobel or Prewitt filter, and 2nd order derivative filters such as the Laplacian filter are particularly used.
- **Region-based approaches:** When regions are more or less homogeneous, voxels and their neighbours look similar, except at boundaries. Voxel intensities, standard deviations or gradients are computed and compared to the ones of neighbouring voxels (classification step) in an ascending or descending fashion (e.g. split and merge algorithm). When regions are not homogeneous, but present repetitive patterns, texture analysis may be performed. It is based on local parameters (e.g. frequency or statistics), extracted within a specific window. Alternatively, cross-correlation with prior texture primitives may be performed. An example of the use of statistical parameters is the cooccurence method.
- **Statistical approaches:** Bayesian approaches try to maximise the knowledge of the scene (classification) given the image, through its probability density function $p(S|I)$. According to the Bayes rule: $p(S|I) = p(I|S)p(S)/p(I)$, this is an optimisation process involving the likelihood (probability of the image given the scene) and the prior (probability of the scene). It leads to $\tilde{S} = argmax(log(p(S)) + log(p(I|S)))$ (Maximum A Posteriori or MAP solution). For expressing the likelihood, additive measurement noise (usually Gaussian) is considered as well as a signal model (such as the Gibbs distribution, considering that the signal is a Markov field). It is also common to incorporate a model of spatial interaction to account for signal continuity in the image. It is a very general framework that can be used also in deformable segmentation and registration [SD92], as we will see later.

1.3.2.3. Classification: Detected regions, or voxels are classified into classes (e.g. tissue type) using homogeneity criterions. When using one parameter, thresholding is performed at a specific level. The criterion is generally image intensity level or probability density. The threshold may be automatically computed through histogram analysis (e.g. Gaussian modelling of grey level distribution). Alternatively, classification can be performed through multiple parameters (multidimensional classification), such as the region mean level, the standard deviation, the contour mean curvature or the object principal axis. Unsupervised classification tries to delineate automatically the different classes in the parameter space using proximity criterions, whereas supervised classification makes use of prior knowledge about parameter values of the different classes. Two strategies are possible regarding the detection/ classification loop. The ascending strategy (split algorithm) starts from under-segmented regions and make them grow iteratively by detecting similar neighbouring regions. In the descending strategy (merge algorithm) the image is split into many regions (over-segmentation) which are subsequently merged through homogeneity/ proximity evaluation and decision making (e.g. quadtree partition). The two may be mixed (split and merge algorithm) [HP78].

1.3.3. Image Registration

1.3.3.1. Introduction Image registration, that is the computation of spatial correspondences between images, is one of the main fields in image processing. Indeed, registration is required in various domains such as image fusion (multimodal registration), segmentation (template registration), motion tracking (temporal registration) and navigation (viewpoint registration). Image registration aims at finding the displacement field that maximises the similarity between images, while preserving a certain regularity of the displacement. Hence, it can be seen as an energy optimisation problem, where the external energy represents the attraction potential field from images, and the internal energy, the regularity of the displacement. Alternatively, image registration can be viewed as a force (energy derivative) equilibrium process where external and internal forces self-balance. The diversity of problems, approaches and comparisons has made any attempt to exhaustively classify registration methods difficult. The various existing reviews [Bro92] [MV98b] [FHM00] have focused on different more-or-less independent criterions such as:

- **The nature of features:** Features, also called primitives or landmarks, are the information, either geometric or iconic, extracted from images to be registered.
- **The similarity measure:** this is the criterion upon which the registration is performed.
- **The problem regularity:** the type of transformation that we are looking for (e.g. rigid, elastic, etc.)
- **The resolution method:** how the optimal transformation, maximising the similarity between features, is found.

1.3.3.2. Registration features: We commonly distinguish geometric features (shape) and icons (image), but it is possible to combine them.

Iconic features: Icons rely on the (pre-processed) photometric information of a voxel (e.g. optical flow [HS81]) or a region (e.g. template, intensity profile [MD05]). With iconic registration, there is no pre-segmentation or geometric feature extraction. However, images may be pre-processed to remove noise and bias, and normalise intensities. Iconic registration has shown to be very efficient in multimodal registration (higher relevance of intensity distribution correlation over extracted features matching).

Geometric features and models: There are two main approaches when using geometric features for image registration. On one hand, they can be segmented separately in the two datasets and then, registered geometrically (using geometrical distance as the similarity measure) [AFP00]. On the other hand, they can be extracted from the source dataset, and subsequently used to segment the target one. With this approach, geometric features (the model) are rigidly or elastically matched, in an iterative process. With deformable models, segmentation, registration and reconstruction are performed at the same time, leading to a possible combination between iconic (image-based) and geometric registration. Moreover, they allow spatially coherent regularisation mechanisms.

Geometric features of lowest dimension are **points**. We differentiate artificial (or extrinsic) and anatomical (or intrinsic) fiducial points [FHM00]. Artificial markers are fixed invasively (e.g. bone pins) or non-invasively (e.g. skin markers) to organs for inter-patient image registration, intra-patient image registration or motion tracking in image sequences. Their discrete locations can be extracted from marker regions using image processing methods or directly registered through iconic registration (block matching). Intrinsic landmarks rely on salient anatomical features. They can be extracted from images, but robust point extraction (automatic or manual) is often a difficult task [RFS99]. Point registration can lead to rigid or elastic registration: from known spatial correspondences at discrete locations, it is possible to estimate the global rigid movement or the surrounding elastic displacements using error minimisation and adequate interpolation methods such as radial basis functions (e.g. thin plate splines). Interpolation methods are usually not realistic at locations distant from landmarks. So that accurate point-based registration is generally limited to rigid registration (tracking applications). For rigid registration, a **coordinate system** can be used instead of a set of landmarks. This can be obtained through gravity center, principal axis and moment calculation from grey-scale or binary images, but also from organ shape (intrinsic anatomical axes). To overcome unrealistic point-based elastic registration, landmarks orientation and anisotropic error computation have been incorporated in the classical Thin-Plate-Spline (TPS)

interpolation method [RFS99]. This provides better results, especially at organ interfaces where the registration transformation is not continuous.

Higher dimensional geometric features, such as curves, surfaces and volumes, have various mathematical representations. The type of representation is a key-point with deformable models, as it influences geometric descriptors, the evolution process and the visualisation methods associated to them. McInerney et al. [MT96], Montagnat et al. [MD05], Singh et al. [SGT98] present detailed reviews and classifications of deformable models. *Curves*, such as crest lines, are perceptive features that may be used for registration; however, *surfaces* remain the most used geometric primitives, since boundary is commonly the most noticeable object feature. In images, homogeneous objects, are, in fact, only characterised by their shape. We broadly present the various representations used in image registration, and to a larger extent in tissue modelling and simulation.

There are three distinct families of deformable models: **continuous models** are definite through the mapping of parameter ranges in space; **discrete models** are defined at discrete locations in space and connectivity relationships; and **implicit models** are indirectly defined by function isovalues. Continuous models offer the ability to compute differential quantities (e.g. normals, curvatures) easily and are well controllable through their limited number of parameters (intrinsic geometric regularisation) but their shape is limited by the parameterisation. On the contrary, discrete models are more flexible but difficult to constrain geometrically. Implicit models can benefit from their volumetric nature, and merging capabilities, but their abstract nature makes spatial interactions with them, difficult.

The **continuous** representation has been the first proposed one, for deformable models. In 1988, Kass et al. [KWT88] propose the well-known *snake* representation (also called active contour, or explicit contour). A snake is an explicit mapping of a real parameter to 2D coordinates. This has been generalised to 3D case by Terzopoulos et al. [TF88b] and adapted by McInerney et al. to handle topology change [MT99] (T-snake). Parametric models are continuous models where the mapping is made through particular mathematical functions, defined by few parameters (degrees of freedom of the model). Continuous models are popular in medical image analysis since their regularities are simply evaluated; examining mapping function derivatives with regards to material coordinates. Moreover, parametric models are intrinsically regularised thanks to their limited number of parameters (DOFs). B-splines, (bicubic) Hermite surfaces, and superquadrics are the most used parametric model/ functions, since they can describe a wide range of shapes, while offering physically significant regularisation/ deformation. Shape limitation can be overcome by adding local shape perturbations [TM91], or by applying FFD subsequently to registration [BCA96]. Continuous models are particularly suited for statistical analysis, because their mapping functions can be efficiently decomposed into a linear combination of variation modes. Different basis (e.g. Fourier, harmonic, principal components) have been exploited in image segmentation [SD92] [CT01] for constraining model evolution into relevant deformations, that can be learnt from examples. Explicit continuous models need to be discretised for evolving (Section 1.3.4), while, with parametric models, control points (part of the parameters) may be directly used. The main problem with this kind of models remains the inverse problem (defining shape functions and parameters from a set of locations), which is required for reconstructing/ initialising models and interacting with them. This is carried out through energy functional minimisation, as a geometric registration process.

Figure 10: *topology adaptive continuous model (T-snake) based segmentation from [MT99]*

A **discrete** mesh is represented by vertices and connectivity relationships. Connected points form faces and cells. In discrete contours (respectively triangle meshes, tetrahedral meshes and rectangular grids), each cell contain exactly 2 (resp. 3, 4 and 8) vertices. They are characterised by a constant cell connectivity. Dually, simplex meshes [Del99] are characterised by a constant point connectivity (in p-simplex meshes, each vertex is connected to exactly $p+1$ vertices). Discrete surfaces are very popular in the field of computer graphics and image processing, as they are flexible (correlation with anatomical structures) and memory cheap. Among the numerous registration methods, [PMTK01] use triangle meshes, while [FWG*99] use tetrahedral meshes and [MD05] use 2-simplex meshes. Abstract volumetric deformation grid have shown to be an efficient support for image deformation (FFD) and hence registration. In [LRMK99] Lötjönen et al. present a coarse-to-fine approach with multiresolution grids. Szeliski et al. [SL96] propose an octree grid, to ensure better fitting with anatomical structure. However, FFD-based registration methods, often fail to handle displacement discontinuities at boundaries.

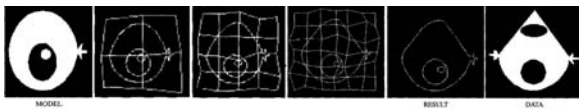

Figure 11: *Multi-resolution FFD-based registration from [LRMK99]*

Implicit Models are defined by the zero of a potential field F taking spatial coordinates as parameters. Com-

plex shapes can be obtained by merging simple implicit primitives, defined by polynomial function, such as superquadrics or hyperquadrics. Ellipsoids are part of superquadrics. Superquadrics also have a parametric formulation and have been widely used in image segmentation such as in [BCA96]. Level-set [OS88] is the most used implicit model representation in medical image analysis. It has become popular since they allow a great flexibility in terms of shape and topology, through higher dimensional parameters. However, as emphasised by Montagnat et al. [MD05], level-sets are computationally heavy and generally difficult to constrain spatially.

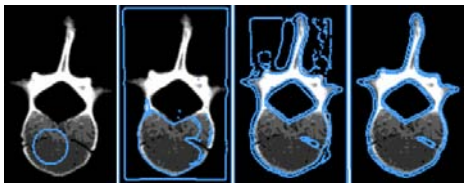

Figure 12: *Level-set based segmentation from [MD05]*

1.3.3.3. Similarity measure: The similarity (or external energy) between images (or between geometric features) is a measure of the matching quality. It is an important parameter upon which the searched transformation is computed. Global external forces and moments can be defined through the examination of the global energy variation according to transformation parameters (**standard approach**). In opposite, local external forces, subsequently converted into the closest global transformation, correspond to the **pair and smooth approach** (e.g. Iterative Closest Point (ICP) algorithm [BM92]).

In **geometric registration**, the similarity is derived from the distance between features: for instance, the Euclidian distance or the p-order Minkowski distance between two points, the Hausdorff distance or Mahanalobis distance between polygons, etc. A review of geometric registration techniques in medical image processing can be found in [AFP00].

A common methodology in segmentation is to derive the external energy from **image gradient** values [KWT88] around models. In this case, models converge to locally maximal gradients, but may be attracted by undesired noisy edges. This can be partly corrected by regularising the gradient vector field. Nevertheless, a more robust approach is to use prior information about image intensity and gradient through **intensity (or iconic) registration**. In iconic registration, the similarity relies on the correlation between the intensities of aligned voxels or regions (in the source reference dataset and the target dataset). Depending on the assumptions about this correlation [RMA00], many similarity measures have been proposed [HHD*00] [MV98b]. Intensity conservation along displacement, which can be valid in motion analysis, is assumed with the sum of absolute differences (SAD), the sum of square differences (SSD), or the cross-correlation (CC) measures. Between these measures, CC is the least noise-sensitive but the heaviest in terms of computation. In mono-modal registration, the hypothesis of affine relationship between intensities from the two registered datasets improves robustness with regards to global intensity change across acquisition protocols and sessions. In this case, the cross-correlation normalised by the standard deviation (NCC) measure [HHD*00] [RMA00] eliminates the affine terms. In multimodal registration, assuming that intensities values can be mapped between datasets through a certain function (functional relationship), measures have been proposed, such as the correlation ratio (CR) or the variance of intensity ratio (VIR) measure [WCM92]. Finally, for multimodal registration, information theory has lead to popular and flexible measures, based on the joint entropy. Minimising the joint entropy (or joint histogram spreading) is equivalent to optimise intensity value prediction from one image to another, without a-priori knowledge of the functional relationship. The (normalised) mutual information (NMI, MI) results from the normalisation of the joint entropy [CMD*95] [Vio95]. MI is robust with regards to change in the size of the overlapping region. In [RMA00], Roche et al. unify the MI and CR measures within a maximum likelihood (ML) framework.

Prior to similarity computation, intensity filtering can be performed to decrease noise (e.g. anisotropic smoothing) or enhance particular features by convolution with discrete masks. Particularly, some authors have used image gradient magnitudes [MD05] or vectors [GPMTV04] instead of basic intensities.

1.3.3.4. Regularisation: Due to noise, local solutions and the aperture problem (apparent motion), the registration process need to be constrained through degrees of freedom minimisation (**parameterisation**) and **internal regularisation**. Making hypothesis about the regularity of the problem and the form of the solution (initial distance between the two datasets) is a key-point in a registration process, for at least two reasons: 1) it determines the solution search space (dimension, order of magnitude) and consequently rules most of the computational charge (parameterisation of the resolution); 2) it decreases the uneven effect of external forces and energy through intrinsic constraints. Assumptions are generally application-dependent: for instance, rigid transformations are involved in intra-patient multimodal fusion, bone tracking or viewpoint registration problems.

Parameterisation: Parameterisation of the resolution aims at limiting solution search to relevant degrees of freedom. Indeed most of the difference between datasets can be recovered from a limited number of parameters. The researched transformations can be global or local, whether which subset of images is concerned. Coarse-to-fine approaches increase the number of DOF and the localness of

the transformation during the registration process, avoiding fall into local minima of the energy. There are basically two ways or parameterising a registration process: the standard approach maximises the similarity directly in the space of admitted transformations (global parameterisation), while the Pair and Smooth (P&S) approach locally maximises the similarity (pairing) according to a local parameterisation, and subsequently smooth the result (through standard smoothing of the displacement or explicit resolution).

In the standard approach, global parameterisation can be performed using **homogeneous transforms**. Indeed, standard 3D geometric transformations can be formulated with 4×4 homogeneous matrices derived from successive basic transforms (e.g. rotation around one axis, shear in one direction, etc.).

Geometrically continuous elastic deformations can be derived (interpolated) from the **structured deformations** of a regular lattice (e.g. FFD [SP86]). The interpolant function is a triple tensor product of 1D weighting functions. The use of cubic splines is particularly popular since they allow various continuity constraints: Hermite, Bezier and Catmull-Rom enforce C1 continuity, while B-Splines are C2 continuous. Hence, B-splines are widely used in image registration [RSH*99]. The use of a regular lattice can be limiting for handling displacement discontinuities at organ boundaries. Indeed, a careful matching between model geometry and anatomy provides more realistic deformations. As extensions of B-splines, octree B-splines [SL96] and triangular B-splines [WHQ05] allow the use of more flexible lattice in terms of mesh topology.

Besides, **unstructured deformations** from irregular control points can be performed through radial basis function - RBF. The interpolated displacement is defined as the weighted combination of non-linear functions (or radial basis) of the distance from control points plus an affine or polynomial term. Weights are pre-computed according to control point displacements (boundary conditions). The TPS radial basis [RSS*96] is popular for interpolating displacement since it minimises the bending energy; moreover it can handle different local continuity constraints [RFS99] and is well adapted when the number of control points (landmarks) is small (however, it turns to heavy computation when it grows). The use of RBF is difficult in a variational framework, because, weights need to be recomputed at each control point displacement. In general, with RBF, scattered spatial correspondences are manually selected, or semi-automatically segmented; and not iteratively found in an optimisation framework. However, in [RAD03] Rohde et al. use a compactly supported radial basis allowing weights to be optimised locally and independently for each group of 8 points.

In P&S registration approaches, feature local pairing is constrained through **local parameterisation**. Because image registration highly relies on object boundaries registration, it is particularly suitable to constrain deformation in the normal direction of these boundaries and within a certain range (similarity is maximised in this search space) [Del99] [MD05] [LM01]. In [HS81] [Thi95], pairing is done in the direction of image gradient, with a certain magnitude (optical flow). With the DFD [HS81], the displacement is computed as $\mathbf{u} = (T - S)\nabla\mathbf{T}$. While, with the demon algorithm [Thi95] we obtain $\mathbf{u} = (T - SoM)\nabla\mathbf{T}/(\|\nabla\mathbf{T}\|^2 + \alpha^2(T - SoM)^2)$ (SoM is the current transformed source image within the iterative process), leading to a normalisation of the displacement (its amplitude remains below the constant $1/(2\alpha)$). Ghanei et al. [GSZW98] use both surface normal and image gradient directions.

The last class of parameterisation techniques is the one based on **statistics**. That is new shapes are obtained by a linear combination of N sample shapes. Consequently, deformations are limited in a space of N dimensions. Relevant samples are generally obtained by PCA. The example-based approach is quite popular in image segmentation and registration (brain) since it takes into account shape and appearance (image intensity) variability within the population [CT01] (Active Shape Model - ASM and Active Appearance Model - AAM) [WS00] [DSMK01].

During the registration process, a widespread approach is to progressively increase the number of DOF and the localness of the transformation in order to improve robustness and computational speed; this is the **coarse-to-fine approach**. In [MD05] and [RSH*99], geometric transforms, from rigid to affine, are successively applied. In [LM01], deformation spheres become smaller and smaller. Multi-resolution lattices [SL96] [RAD03] and multi-resolution surfaces [PMTK01] [GMMT06] are also used.

Internal forces/ energy: Intrinsic regularisation aims at ensuring a certain continuity of shapes and displacements by penalising other configurations during optimisation. Because the image registration problem is ill-conditioned (no stability) due to noisy external terms, extra regularisation terms (stabilisers) are needed within the resolution process. Either forces or regularisation energy terms can be applied depending on the chosen evolution process (see Section 1.3.4). Basically, they are of three types: smoothing, physically-based and example-based.

Smoothing energy or forces enforce shape continuity through curvature, elastic energy and/or bending energy minimisation. It relies on the assumption that biological tissues are relatively smooth. Most of regularisation terms, used in the literature, are based on 2nd order Tikhonov differential stabilisers [TPBF87], since they are a generic expression of shape deformation through differential terms up to the second order. 1^{st} order terms are related to the elasticity of the model (minimisation of the curvature, Laplacian smoothing), while 2^{nd} order terms deal with bending (averaging of the curvature). Forces derived from the minimisation of the energy can be obtained from the energy using

the Euler-Lagrange theorem. When using explicit or parametric models, the differential terms can be easily computed due to their continuous nature. This is more complicated for discrete models where geometric approximations need to be made. Standard Laplacian smoothing forces (attraction of vertices towards the barycenter of their neighbors) minimise the elastic energy but produce model shrinking. This well known problem can be tackled by compensating them with surfacic balloon forces [Coh91] or removing their normal contributions [MD05]. Their tangential contributions prevent from surface stretching. Bending forces rely on the averaging, in the normal direction, of the discrete curvature. There are several ways for estimating the curvature: Park et al. [PMTK01] average the normal contribution of Laplacian forces; Montagnat et al. [MD05] smooth the simplex angle; Ghanei et al. [GSZW98] smooth the Gaussian curvature. Similarly to shape, deformation can be regularised [TPBF87] [LM01], by enforcing continuous variation of the deformation with Tikhonov stabilisers. Deformable models, fitting to organ boundaries, implicitly handle deformation discontinuities at boundaries. However, this is not the case with FFD-based registration, where boundaries are embedded into the deformable volumetric lattice, so that Tikhonov stabilisers would not suit (isotropic regularisation). Different anisotropic stabilisers, reviewed in [DKA95], have been developed for handling discontinuities. They basically provide anisotropic smoothing perpendicularly to image gradient. Radial regularisation of surfaces is useful for constraining models possessing a medial representation. Indeed, this provides to surfaces a volumetric aspect and higher level control parameters through simple mechanisms such as radial springs [PFJ*03] [GMMT06].

Smoothing forces, minimising the bending and/or the elastic energy can be considered as weak prior information [MD05] because biological tissue generally satisfy this condition. On the contrary, **example-based regularisation** provides strong prior information as it may be specific to a particular problem, organ, individual, instant, etc. Montagnat et al. [MD05] add shape and trajectory constraints through forces attracting shapes towards a reference configuration. In contrast, example-based parameterisation combines sample deformations to compute model deformation.

The minimisation of the elastic strain energy ($E_{reg} = \int W dv$) has been used in several image registration studies [Chr94] [WS00] [BN96] [VGW06], since it provides **physically-based regularisation** (one-to-one mapping, no negative volume). Equivalently, one can use body forces (divergence of the stress), in a force equilibrium evolution process 1.3.4, where an equilibrium with external forces (image-based) is iteratively found. In the linear elasticity framework, body forces are expressed as (linear elasticity operator):

$$div(\boldsymbol{\sigma}) = \mu\Delta\mathbf{u} + (\lambda+\mu)\nabla(div(\mathbf{u})) = -\mathbf{f_{ext}}$$

where λ and μ are the Lamé coefficients and $\mathbf{u}$ the displacement. Since infinitesimal displacement is not a valid assumption, Christensen et al. [Chr94] introduced the fluid registration, where the internal stress is reinitialised at each iteration. To speed up the rather slow algorithm from Christensen et al. (several hours for 2D registration) Bro-Nielsen et al. [BN96] implement a fast filtering method from the impulse response of the linear elasticity operator. As a natural extension, authors have investigated hyperelastic registration to allow large displacement in the elastic framework (no stress re-initialisation). Nowadays, sophisticated constitutive models are used for registration such as in [VGW06], where Veress et al. incorporate an anisotropic hyperviscoelastic constitutive model for cardiac registration. In addition to the large computational weight of such methods, there are some restrictions in using mechanical parameters, making physically-based regularisation not necessarily more robust than geometric regularisation: first, external forces from images do not have any physical meaning, so that they are difficult to blend with body forces. There are also topological problems due to the linearization of the elasticity operator. In addition, inter-patient registration based on mechanics does not have much sense (there is no matter deformation). Finally, the mechanical model is difficult to fit to the data because of assumptions (small displacement, linear isotropic behaviour, etc.) and unknown mechanical parameters. The problem can be over constrained such as the over penalisation of large rotations when using linear elasticity.

In multiple object segmentation, it is valuable to add extra constraints, preventing from interpenetrations. Few authors have actually incorporated neighbouring constraints and collision handling within an image registration process. In [MT99] McInerney et al. check self-penetration of their T-snake, using spatial hashing, to allow topology change. Non self-intersecting forces are applied in [PMTK01] on discrete models. Recent advances in collision handling from the computer graphics community [TKH*05] and topological constraints (organ attachment) have being applied in [GMMT06] on simplex surfaces.

1.3.4. Model evolution (numerical resolution):

The solution of the registration problem is found by minimising the energy of the model, composed of an internal/ regularisation term and an external/ image-based potential energy: $\tilde{\mathbf{P}} = argmin_{\mathbf{P}} E(\mathbf{P})$ ($\mathbf{P}$: parameters). At minimum, competing forces equilibrate and we have $\mathbf{F} = -\nabla\mathbf{E} = \mathbf{0}$. In most cases, because of the non-convex aspect of the energy (noise, multiple local minima, etc.), there is no analytical solution. After discretisation, the energy is minimised iteratively using traditional multivariate optimisation strategies. Alternatively, when conferring a mechanical aspect to energy/ forces, the iterative process can be performed through a dynamic evolution. Another class of evolution method is the Eulerian evolution, where the space is transformed rather than the model itself. The solution (model shape) is hence

implicitly defined by the space state. Eulerian evolution has been investigated for the front propagation of implicit models (level-sets [OS88]) and will not be discussed here.

1.3.4.1. Explicit resolution: In some simple cases, where the system has a small number of DOFs, and the external energy a simple form, it is possible to analytically find the global solution. In [AHB87], the best rigid, similarity or affine transform that minimise the distance (error in the least square sense) between paired point-sets is found, by applying the quaternion method. This is particularly useful in body motion capture, where the rigid motion of the different body segments need to be estimated from noisy marker motion. When not available, pairing can be performed by minimising locally the distance (e.g. closest point). The pairing is then updated at each iteration, where the optimal transform is found, until the convergence (P&S approach). This is the principle of the popular Iterative Closest Point algorithm [BM92].

1.3.4.2. Energy minimisation: The minimisation, aimed at seeking the equilibrium, is driven by forces (energy spatial derivatives). Applying the Euler-Lagrange equation, the energy minimisation problem is turned into a force equilibrium equation (stationary equation) that can be use to derive the evolution of the model. Space discretisation is performed by the finite difference method [TPBF87], the finite element/volume method [MT95] or any parameterisation scheme (e.g. B-Splines) that lumps system degrees of freedom to control nodes. After discretisation, we obtain $-\nabla \mathbf{E}(\mathbf{P}) = \mathbf{F_{int}}(\mathbf{P}) + \mathbf{F_{ext}}(\mathbf{P}) = \mathbf{KP} + \mathbf{F_{ext}}(\mathbf{P})$ ($\mathbf{K}$: the stiffness matrix) which is zero at equilibrium (local energy minimum). Instead of minimising the energy (maximum likelihood scheme), some authors maximise the probability of model shape given the image (maximum a posteriori scheme), through a Bayesian framework [SD92] [WS00]. This can be viewed as an extra layer to which optimisation is performed. The Bayesian approach is interesting for incorporating statistical shape information (example-based approach) and image noise.

Local optimisation: to find the minimum of the energy, *exhaustive research* methods, or quasi-exhaustive methods (e.g. multigrid) may suit when the number of degrees of freedom is small (e.g. translation). In most cases, the complete inspection of the search space is too costly, and the search must be oriented. A popular oriented search method is the *Downhill Simplex method* [NM65]. The associated figure (simplex) is successively transformed through reflections, expansions and contractions, until a tolerance (simplex volume) is reached. Another way is to descend along to the energy gradient [Thi95] (the model evolves in the steepest energy direction with a certain speed). However, the convergence might require a large number of iterations (oscillations often occur around the solution). To speed up the convergence, one may adapt the time-step and the direction of the evolution. In the *conjugate gradient* algorithm, the direction of the evolution is forced to be orthogonal to the previous ones. The *Powell's method* provides conjugate directions without computing the derivatives. In the *Newton algorithm*, the time-step is computed, assuming that the energy is quadratic (2nd order Taylor development of the energy). The *Levenberg-Marquardt method* smoothly combines the *Newton* and the *gradient descent* algorithm (the Newton method is more suited near a local minimum, whereas the gradient descent works better far away from the solution). Newton method involves the expensive evaluation of the Hessian matrix of the energy. *Quasi-Newton methods* have been proposed to update an Hessian matrix estimate at each iteration, considering energy and energy gradient change. Considering the first order development of the energy that we want to be null in the next time step and assuming a displacement along the gradient we get the *Newton-Raphson method* where only first order energy derivatives (forces) need to be computed. In [MHHR06], Müller et al. use this method within a mechanical system to derive its dynamic evolution.

Global methods: Above methods are local, meaning they work well when the energy has few local minima (convex energy). If not, *stochastic optimisation* methods such as the simulated annealing or evolutionary algorithms can be used to find a global solution, despite their significant computational load. The simulated annealing (or stochastic relaxation) randomly changes parameter values at each iteration, and allows energy increase depending on the temperature (the higher the temperature, the higher rise is allowed). The temperature decreases until the instant freezing, where allowed perturbations only make a reduction of the energy. *Evolutionary algorithms* (e.g. genetic algorithms) are based on the Darwinian principle where a population of agents (parameters) evolves according to a fitness function (energy) and gene crossing principles. In [MH06], a genetic algorithm is used along with deformable models for medical image segmentation. *Dynamic programming* can be used to find a global solution. It is based on the decomposition of the energy into independent terms. Solving these sub-problems, the global solution can be recovered by finding the shortest path in the energy graph. Amini et al. [AWJ90] have applied dynamic programming to 2D snakes. But, this has not been extended in 3D, the critical point being the ability to decompose the energy.

1.3.4.3. Dynamic evolution: In a dynamic system, forces drive the model, that is the position and velocity of control points. This is particularly suited for discrete models that are commonly considered as lumped mass particles, moving according to forces. Dynamic systems are judged against their stability and accuracy, resulting from the numerical integration of the differential equation. The Newtonian law of motion leads to a first-order differential equation system relating the force vector $\mathbf{F}$ to the particle velocity vector $\mathbf{V}$ and position vector $\mathbf{P}$:

$$\frac{\partial \mathbf{P}}{\partial t} = \mathbf{V} \qquad \mathbf{M}\frac{\partial \mathbf{V}}{\partial t} = \mathbf{F}(\mathbf{P},\mathbf{V})$$

The force vector $\mathbf{F}(\mathbf{P},\mathbf{V})$ depends on particle position (internal and external forces/ energy derivative) and velocity (damping representing the viscosity of the milieu/ energy dissipation). When discretising the system with finite differences, we obtain:

$$\begin{aligned} \mathbf{u} = \mathbf{P_{t+dt}} - \mathbf{P_t} &= \mathbf{V}dt \\ \mathbf{V_{t+dt}} - \mathbf{V_t} &= \mathbf{M^{-1}}\mathbf{F}(\mathbf{P},\mathbf{V})dt \end{aligned}$$

The Lagrangian evolution, that assumes massless particles and unitary isotropic viscosity [TM91] ($\mathbf{F}(\mathbf{P},\mathbf{V}) = -\nabla\mathbf{E}(\mathbf{P_t}) - \mathbf{IV} = \mathbf{0}$) leads to $\mathbf{P_{t+dt}} - \mathbf{P_t} = -\nabla\mathbf{E}(\mathbf{P_t})dt$, which is equivalent to the gradient descent algorithm.

In the **explicit (or forward) methods**, velocities and forces are taken at time t ($\mathbf{V} = \mathbf{V_t}$ and $\mathbf{F}(\mathbf{P},\mathbf{V}) = \mathbf{F}(\mathbf{P_t},\mathbf{V_t})$), so that the new state vector (P_{t+dt}, V_{t+dt}) is explicitly defined in the system 1. However it goes forward blindly (no notice of force derivatives), so that the stability is ensured for very small time steps dt. A higher order forward method is the well known Runge-Kutta method [PTVF92]. Its principle it to use force multiple evaluation to better extrapolate the new state vector through higher order Taylor expansion. A simple way for improving stability, with no extra computational cost, is to take the implicit formulation of the first equation of 1: $\mathbf{P_{t+dt}} - \mathbf{P_t} = \mathbf{V_{t+dt}}dt$, where $\mathbf{V_{t+dt}}$ can be obtained with the forward step described above.

In the **implicit methods**, the idea is to estimate state vector temporal derivatives at $t+dt$ using its spatial derivatives [BW98] [VMT05]. Implicit methods are called backward since a forward step from new state vector at $t+dt$ brings it back to the initial state vector at t. Hence, results are always consistent, allowing larger time-steps. Since [BW98], implicit methods have become popular in computer graphics. However they are not necessarily more accurate than explicit methods, due to numerical damping. A comparison of explicit/ implicit methods in terms of stability and accuracy is done in [VMT05] [HE01]. The implicit resolution is equivalent to the resolution of a large sparse linear equation system $\mathbf{V_{t+dt}} - \mathbf{V_t} = \mathbf{H^{-1}}\mathbf{Y}$. The more stiff (large number of eigenvalues) is the system, the more complicated is the resolution (more solver iterations). In [DSB99], only the linear part is resolved, through pre-computed inverse matrix. In [EEH00], explicit resolution is applied to the non-stiff nonlinear part, while implicit resolution is applied to the stiff linear part (IMEX method), which improves accuracy (less numerical damping due to implicit resolution) and computational speed. To increase accuracy, authors have used history states to better approximate the new state. Particularly, Backward Differentiation Formulas (BDF) have been recently applied to deformable models [HE01]. They are computationally inexpensive, since they only require one system resolution at each step, contrary to Runge-Kutta methods. BDF are less robust in the cases of nonlinear systems, since they are sensible to sudden state change (e.g. collision response) due to their dependence with previous states. In practice, BDF are limited to their second order (first order is equivalent to the implicit Euler).

The resolution of the system always involves the inversion of a large sparse matrix. Terzopoulos et al. [TPBF87] resolve the linear system with a Choleski decomposition and a relaxation method (Gauss-Seidel). Recent applications use the Conjugate Gradient (CG) algorithm, which is particularly adapted to sparse systems [PTVF92]. The CG algorithm works on symmetric positive-definite systems which is not the case with implicit schemes, due to mass-modified particles. Post-multiplication [VMT00] or pre-multiplication [BW98] by the inverse mass matrix can be performed to recover global symmetry. In [HE01], the CG algorithm is integrated in the Newton minimisation method, allowing error control and adaptive time stepping. The Conjugate Gradient can be accelerated using matrix preconditioning [She94]. System stiffness is reduced through matrix multiplication: $\mathbf{V_{t+dt}} - \mathbf{V_t} = (\mathbf{U^{-1}H})^{-1}\mathbf{U^{-1}Y}$ where $\mathbf{U}$ is the symmetric positive-definite preconditioning matrix. Indeed, if $\mathbf{U^{-1}H}$ contains less non-zero elements than $\mathbf{H}$, its inversion is cheaper. Popular preconditioners are the Jacobi preconditioner ($\mathbf{U_J}$ = diagonal of $\mathbf{H}$) and SSOR (Symmetric Successive Overrelaxation Method) conditioner ($\mathbf{U_{SSOR}} = (\mathbf{U_J} + \mathbf{L})\mathbf{U_J^{-1}}(\mathbf{U_J} + \mathbf{L})^T$ where $\mathbf{L}$ is the lower triangular part of $\mathbf{H}$). Following the idea of Terzopoulos et al. [TW88], forces can be applied on rigidly registered models (non-rotated references), leading to more sparse system, especially when using linear Cauchy strain [MG04]. This can be viewed as a particular preconditioning method.

1.3.5. Conclusion

Most proposed registration techniques remain generic and fairly independent from the targeted applications. The lack of specificity makes methods difficult to parameterise, difficult to evaluate and generally holds back their implementation in the clinical environment. However, it is reducing due to more-and-more established physical-based simulation methods, available data for example-based approaches and cross-validation studies. Unifying tissue simulation and segmentation could lead to a better specificity for daily clinical use. Moreover, physically-based constraints, taking into account not only pixel intensity but high-level parameters such as the shape, the function and the mechanics of tissues, improve the accuracy and the robustness of registration methods by reducing problem degrees of freedom. Transformations can be limited to admissible deformations through examples [DSMK01] or/ and prediction methods [VGW06]. This is particularly relevant for intra-patient registration, while, for inter-patient registration or longitudinal studies, it

is required to study how high-level parameters change across individuals or through time. For inter-patient registration, we expect that external forces would play a more significant role than for intra-patient registration (image-driven versus physics-driven). In this context, deformable surfaces offer a moderate complexity and a good flexibility to tune intrinsic, geometrical and data constraints. Besides, physically-based volumetric models are more suited for patient-specific tissue simulation and intra-patient registration.

1.4. Example: musculoskeletal analysis from MRI

1.4.1. Introduction

For the diagnosis, the surgical planning and the post-operative assessment of musculoskeletal disorders, the automatic segmentation of the patient musculoskeletal system is important for orthopaedists, biomechanicians and kinesiologists that would like to simulate, visualize and navigate through articulations with a minimum amount of manual tasks. As stressed by [TSB*05], [NTH00] and [BD05], usual simplified models (stick-figures, muscle action lines) are not able to take into account large attachment areas, as well as global constraints such as volume preservation and non-penetration; although they are important biomechanical parameters. Therefore, the relationship between musculoskeletal dynamics and organ shapes need to be better studied through image segmentation. Magnetic Resonance Imaging (MRI) is a flexible modality for imaging both soft and bony tissues non-invasively. However, due to the large amount of textural information, noise, low-resolution, organ imbrication, and large spatial variability; automatic, fast and robust musculoskeletal segmentation is a difficult task. As a consequence, existing methods for musculoskeletal modelling are interactive [NTH00] [AT01] [TSB*05] [BD05] and therefore time-consuming. A common approach to constrain a segmentation process is to use prior information: shape constraints rely on assumptions about surface regularity (smoothness, curvature) and variability across the population, while topological constraints exploit prior knowledge about organ interrelationships. Contrary to traditional pixel-based segmentation (e.g. level-sets) and registration (e.g. FFD) methods, these constraints can be efficiently applied on deformable models [TWK87a]. Physical-based simulation approaches (volumetric), including finite element or finite volume are appropriate for enforcing mechanical constraints, but their computational cost would exclude time-efficient and user-controlled segmentation. In this context, we propose to use scalable discrete deformable surfaces (simplex meshes), that can benefit from efficient geometric methods, popular in computer graphics applications.

Simplex meshes were first described by H. Delingette [Del99] for constrained 3D shape reconstruction and segmentation, and extended in 4D by J. Montagnat [MD05], with application to the heart (single model and resolution). A k-simplex mesh is defined by a set of vertices and a connectivity function (each vertex is connected to exactly k+1 neighbors). In this paper, we use 2-simplex meshes (dual to triangle meshes). Mesh geometric quality (uniformity of vertices repartition) and topological quality (uniformity of edge number among faces) are improved using simple topological operators. The most interesting property is its simple geometric description: three parameters (two barycentric coordinates plus the curvature) uniquely define vertex positions from their three neighbors. Based, on these parameters, smoothing and shape constraint forces are computed to regularize the segmentation. External forces are obtained through 1D registration of intensity profiles (at vertex positions and in normal direction) with generic profiles from a reference segmentation. The external force field is regularized using a local smoothing and global regularization based on the closest affine transformation.

Discrete models are commonly considered as punctual masses 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). After forces evaluation and time discretisation, particle state can be explicitly resolved (forward Euler, Runge-Kutta) with tight time step restrictions for ensuring stability. We prefer the more stable implicit (or backward) scheme [BW98] [VMT00], that however requires the resolution of a large sparse linear system. To simplify force derivatives evaluation, we consider that forces have a independent anisotropic action on each particle. Non-penetration constraints are based on collision handling techniques [TKZ*04]. In this paper, we apply a hierarchical collision detection scheme based on 18-discrete oriented polytope(DOP) quadtrees [TKZ*04]. We perform collision correction and response on particle positions, speeds and accelerations, such as in [VMT00].

From these background studies, we propose improvements in terms of computational speed and robustness by extending the simplex mesh framework with a multi-resolution scheme, topological constraints and medial surfaces. We show that the medial axis is suited for muscle shape analysis.

1.4.2. Methods

1.4.2.1. Multi-resolution Scheme The use of levels of details (LODs) aims at reducing system complexity and sensitiveness to local solutions. As shape constraints spatial influence depends on the resolution level, the idea is to quickly propagate forces from lower resolutions to a current simulation LOD in order to get multiscale regularization. In addition, collision detection (often considered as the bottleneck for simulation) is more efficiently performed on coarse LODs, collision response being passed to fine LODs. Indeed, exact contact computation is most of the time not relevant as fat separates organs. The tessellation of dual triangle meshes leads to a systematic and computationally efficient

LOD generation scheme for simplex meshes (linear combination of vertex positions). Shape features are preserved as low level vertices are contained in higher levels. The number of vertices is quadrupled when increasing the resolution. During the simulation, forces from lower resolutions are linearly combined like vertex positions and added to current resolution forces. Even if this is not exact (vertices having moved relatively to their neighbors, since resolution increase), the estimation is still relevant, assuming that shape constraints have enforced mesh local regularity.

1.4.2.2. Topological constraints The human musculoskeletal anatomy exhibits various organs interrelationships: muscles are attached to bones, they can merge into common tendon units, fascia binds muscles and enforces frictionless contact between them. We have developed a spline-based method for generating attachment areas, so that the placement and adjustment of areas have a reduced number of degrees of freedom. Spline control points are projected onto bone surfaces, while soft-tissues vertices are attached to the spline through curvilinear coordinates. These vertices are constrained using mass modification [BW98]: $\mathbf{M^{-1}} = \mathbf{0}$. For individualization, we wrap splines from a generic model using spline control point barycentric coordinates on bone surface. As shown in [KdH04], this approach is valid for most attachments as they rely on bone geometrical features. For merging deformable models, we sum forces and masses of the attached vertices. This is interesting for modelling common tendons (e.g. quadriceps femoris tendon), but also specific parts with high curvature, where smoothing forces are not appropriate (e.g. attachment between the adductor magnus and its inferior tendon). Fascia is modelled through collision handling: once generic collisions and proximities have been detected, they are subsequently used as a reference (springs) to smoothly enforce the relative position between models. In other words, we deform generic contacts.

1.4.2.3. Radial constraints Muscles generally have a smooth and tubular shape. Hence, they can be efficiently represented by an underlying piece-wise action line [TSB*05] [AT01], where isotonic contractions are modelled through action line shrinking/stretching, and isometric contractions through radial constraints applied to a wrapped surface. This has been applied in biomechanics for calculating joint moment arms [BD05]. However, muscles with large attachment areas and/or several origins/insertions require many action lines. We propose a continuous representation using medial surfaces, leading to a continuous radial regularization of the overlying surface and enhanced mechanisms for detecting collisions and characterizing shapes. Anchored medial surfaces (M-reps) have been successfully used in [TWK87a] [PFJ*03] for constraining deformable model-based segmentation. The medial axis transform (computation of maximal ball centers and radii inside an object [Blu67]) is an invertible transform that allows an efficient reduction of parameters dimension. Several methods have been presented for approximating the medial axis (MA), based on Voronoi diagrams, on distance maps, or on thinning.

We propose an iterative method based on forces applied to constrained deformable surfaces. Hook's spring interaction forces $\mathbf{f_j^*}$ are applied to medial surface vertices for MA approximation (j indexes medial vertices). Reciprocally, forces $\mathbf{f_i}$ are applied to overlying surface vertices for MA-based shape regularization (i indexes model vertices). Each model vertex $\mathbf{P_i}$ is associated to a medial surface point $\mathbf{P_{i\perp}}$ with a certain radius r_i. Radii are linearly interpolated between medial surface vertices $\mathbf{Q_j}$ of radius R_j to allow continuous force definition (leading to a smooth reconstructed surface from the MA). Given the weights w_{ij}, we have $\mathbf{P_{i\perp}} = \sum_j w_{ij}.\mathbf{Q_j}$ and $r_i = \sum_j w_{ij}.R_j$. From a reference state (Figure 13a), weights are obtained by projecting orthogonally $\mathbf{P_i}$ onto the MA, and radii by the weighted mean $R_j = \sum_i w_{ij}.P_iP_{i\perp}/\sum_i w_{ij}$. Weights w_{ij} and radii R_j are subsequently memorized to compute radial forces for deformed states (Figure 13b). Model forces $\mathbf{f_i}$ are simply defined by $\mathbf{f_i} = k.(1 - r_i/P_iP_{i\perp})\mathbf{P_iP_{i\perp}}$ where k is the spring stiffness (or force weight). By applying the momentum conservation law [VMT00], we calculate the force contribution $\mathbf{f_{ij}}$ for one spring i at $\mathbf{Q_j}$ such as $\mathbf{f_{ij}} = -w_{ij}.\mathbf{f_i}/\sum_j w_{ij}^2$. We average all spring forces associated to a MA vertex: $\mathbf{f_j^*} = mean_{\,i}(\mathbf{f_{ij}})$. To obtain the reference medial surface for a given model, we first construct and attach a simplex plane, based on muscle attachments(Figure 15). After cropping, we simulate the medial surface with radial and smoothing forces, while projecting model vertices and updating radii at each timestep. The medial surface quickly converges to the medial axis. We measure the error between the reconstructed surface from the medial axis and the desired initial model shape by calculating the mean of $|r_i - P_iP_{i\perp}|$. The average error for all tested models is 0.6mm. This value remains stable during the segmentation when applying radial forces to both models and axis.

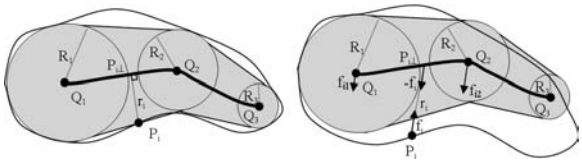

Figure 13: *Surface model and its associated medial axis in: A) A reference state; B) A deformed state; In grey, the reconstructed surface from the medial axis*

In addition, we use medial axis for collision detection: 18-DOP quadtrees hierarchies, inflated according to radii, are generated/updated for medial axis cells. Muscles vertices at a certain resolution are subsequently tested towards DOPs and projected to medial axis for the inclusion test (comparison with interpolated radii r_i). Independently of the model resolution, medial axis based-collision detection is around 3

times faster than surface-based collision detection. On top of MA based-collision detection, multi-resolution collision handling is used to improve computation time (by a factor of 7 per LOD). Given a current resolution, we use its first coarser LOD for collision handling in order to minimize errors.

1.4.3. Results

1.4.3.1. Automatic segmentation of the musculoskeletal system MRI protocol definition has been done in close collaboration with physicians from radiology and orthopaedics. The goal was to obtain images carrying sufficient information with clinically achievable protocols (fast). The final protocol was T1-weighted spin echo with TR/TE=578/18ms, Matrix/ FOV=512x512/40cm and slice thickness=2mm to 10mm between the foot the the iliac crests (five series in total, $\sim$150 slices, $\sim$30 minutes).

Based on interactive segmentation, we have reconstructed a generic model of the hip and the thigh, composed of a skin, 20 muscles, 4 bones and the corresponding attachment splines. After topological optimization, the different LODs (3 for soft-tissues, 4 for bones) and medial surfaces were generated. The final surface model (highest resolution) is composed of 71328 vertices for muscles and 85100 vertices for bones. Alternatively, muscles can be represented by their medial surfaces (1 resolution, 3821 vertices) with an average error of 0.6mm. The compression factor (parameter dimension reduction) is equal to (3x71328)/(4x3821)=14 (3 dimensions for model points, 4 for axis points, including radii).

Before automatic individualization, the generic model is coarsely initialized using ten manually placed landmarks corresponding to anatomical landmarks, and using thin-plate-spline (TPS) interpolation. After this step, bones are deformed (Euler implicit integration) from the coarse level to the fine level using multi-resolution internal forces (shape and smoothing constraints), non-penetration constraints (multi-resolution collision handling) and intensity profile-based external forces. Muscles are subsequently individualized (Figure 14) as follows: A) attachment splines are initialized on bone surfaces from their generic barycentric coordinates; B) a skeleton-driven deformation algorithm [KMTM*98] (skinning) is applied to generic muscles and medial surfaces according to joint angles; C) soft-tissues/medial surfaces are deformed using internal forces (radial/shape and smoothing constraints), proximity constraints (deforming contacts), and skin surface matching (gradient-based external forces); D) soft-tissues are deformed from the coarse level to the fine level using internal forces, proximity/non-penetration constraints and intensity profile-based external forces. The best intensity profile size and resolution have been experimentally defined from interactively segmented models by minimizing excursions. Being more variable, the external part of the profile is less relevant (thus shorter). The intensity profile search depth decreases from 20mm to 5mm during segmentation (Figure 14). We found that normalized cross-correlation is the most robust metric for intensity profile similarity measure. The affine regularization contribution decreases from 100% to 20%. For fine resolutions, internal forces from coarse levels are propagated (global/local constraints) as presented in **2.1**. Muscle shape constraints are derived from their medial surfaces with generic radii and weights. However, radii are updated for fine levels to give more freedom to surfaces. Radial forces are applied simultaneously to surfaces and to medial axis with an equal contribution. For coarse levels, generic proximity constraints are applied (deforming contacts). These constraints are released for the finest level, where multi-resolution medial axis-based collision detection is applied to allow surface sliding.

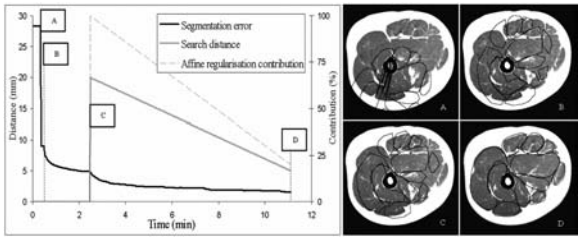

Figure 14: *Automatic muscle segmentation process and result on a sample slice: A) Generic model with wrapped attachments; B) Skeleton-driven model initialization; C) Segmentation using internal forces and skin matching (min. resolution); D) Final segmentation using external forces (max. resolution)*

During the segmentation, it is possible to interactively place constraint points on the images, to get a faster matching and a more accurate segmentation. Using our collision detection scheme, deformable models are forced to include or exclude these constraint points. In addition to the generic model, this method has been applied by a medical student to segment accurately 4 datasets from 4 different healthy subjects (2 females and 2 males). We compared automatically segmented models to these reference models (Figure 14). The average distance (std. dev.) was 1.25mm (1mm) for bones and 1.7mm (1.8mm) for muscles. The overall computation time for the automatic method is around 15min on a standard PC, for which 3/4 of the time is spent on external forces computation. Computation time for each timestep, including visualization, is around 0.5sec.

1.4.3.2. 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).

Muscle thickness can be simply analyzed through medial axis radii comparison, as shown in Figure 15. Using geodesic distances to attachments, 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 (Figure 15).

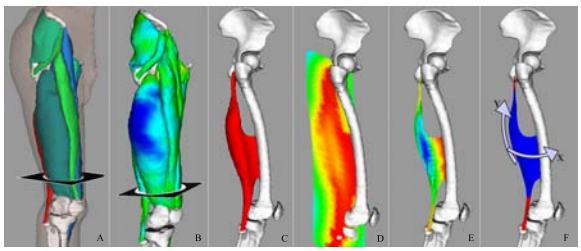

Figure 15: *A) Reconstructed generic model; B) Muscle thickness comparison with an individualized model (blue:≥5mm; red:≤5mm); C) Biceps femoris generic model; D) Initialized medial axis (blue: max radius, red: min radius); E) Medial axis after cropping and fitting; F) Tendon selection (in red) after thickness profile analysis along Y*

1.4.4. Discussion and future work

Prior low-level (e.g. curvature) and high-level (e.g. medial axis) shape information and topological relationships (e.g. proximities, attachments) are relevant for musculoskeletal modelling, and complexity can be efficiently decreased using a multi-resolution approach for force and contact computation. By constraining the problem, our goal is to get a fast and accurate segmentation from a minimum amount of information: we want to extend our previous work on bone motion extraction from real-time dynamic MRI [GPMTV04] (6 low resolution slices) by extracting soft tissue deformation. We believe that we can get a higher accuracy through a deeper study of intensity profile forces (combination of several metrics, weighting according to profile relevance, etc.) and by adding statistical constraints applied to high-level descriptors (e.g. medial axis radii) according to joint angles. The next step will be to relate these descriptors with dynamic (e.g. moment arms) and physiological parameters (e.g. muscle activation from EMG). This will provide useful information for validating functional biomechanical models.

2. Physically-based simulation of biological tissues

2.1. Constraints for Soft Tissue Modeling

There is no such thing as a good soft tissue model. The value of model lies in its adequacy with its applicative context. Thus, we first overview the constraints that arise from the use of biological tissue models inside a medical simulator.

2.1.1. Biomechanical Model

The obvious constraint for soft tissue modeling is that it correctly represents the deformability of real tissue. More precisely, we need to control the accuracy of deformation and therefore compare the computed soft tissue model with the actually deformed tissue. The level of accuracy typically depends on the application. For surgery training for instance, realistic visual and haptic display is more important than the accuracy of deformation. However, if the difference of behavior is too great under large deformations, it could result in learning inappropriate procedures.

In all cases, it is vital to have quantitative knowledge of the biomechanical behavior of soft tissue. We distinguish between knowledge of soft tissue deformable properties and knowledge of interaction with surrounding tissues.

The study of soft tissue deformability belongs to the field of biomechanics. There exists a large bibliography on the study of soft tissue deformation [Fun93]. Those studies range from the determination of qualitative behavior of tissues to the recovery of quantitative parameters governing their deformation. The tissues that have been mainly studied include skin [Lar86], bones [Cow01], vessels [Fun93], muscles [McM87], brain [CM96] and heart [HS88].

Many mathematical models of soft tissue deformations have been proposed [Fun93], and an extensive description would fall outside the scope of this paper. However, we briefly present the simplified models that are commonly encountered for surgery simulation purposes.

The simplest model of static reversible elastic deformation corresponds to the *linear elastic* model. According to this model, a cylindrical tissue sample of height H and diameter L under a load τ (see figure 16) follows Hooke's law :

$$\varepsilon_1 = \frac{\Delta H}{H} = \frac{1}{E}\tau \quad \varepsilon_2 = \frac{\Delta L}{L} = \frac{\sigma}{E}\tau \tag{1}$$

where :

- ε_1 and ε_2 represent the longitudinal and transverse elongation of the tissue under the load τ.
- E and σ are the Young modulus and the Poisson coefficient that characterize the nature of the material.

Therefore a linear elastic material has a linear *stress/strain* relationship and the trajectories of physical points under a load of varying intensity are straight lines. Because of this linear relationship, linear elastic materials have been used for fast surgery simulation [CDA99a, PRZ92, KGG96].

For most materials, the linear elastic model is only valid for small displacements. For large displacements, more complex non-linear models have been been introduced such as the Mooney-Rivlin model [BN96, SBM*94, TBHF03], the Veronda materail [HFS*01] or the St. Venant Kirchoff [TF88c, BNC96, Kai96, PDA03] where the stress/strain

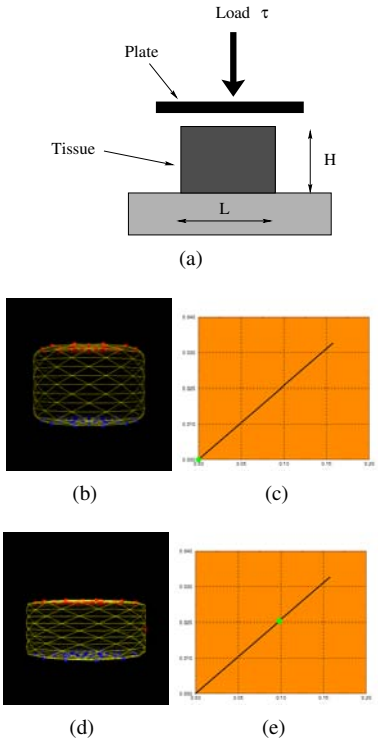

Figure 16: *Behavior of linear elastic materials.*

or stress/displacement relationships are no longer linear. Additional physical constraints may be considered, such as isotropy, transverse isotropy (e.g. for fibrous tissues) and incompressibility.

Furthermore, for many tissues, plastic deformations, where the strain does not reverse to zero after unloading (see figure 17 (a)), occur when the material reaches its elastic limit. Similarly, the material may have a non-reversible elastic behavior (see figure 17 (b)).

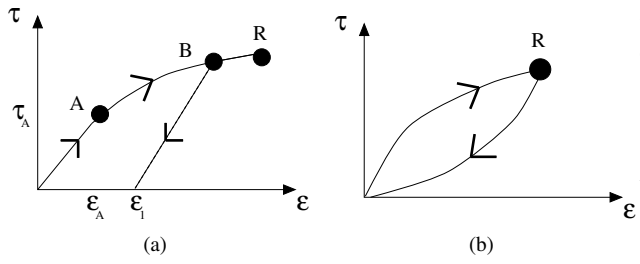

Figure 17: *Stress/strain relationships for plastic material (a) and hysteresis elastic material (b).*

The stress/strain relationship characterizes the static behaviors of a material. However, in general, the stress τ is also related to the speed of deformation $\dot{\varepsilon}$ (the *strain rate*) and therefore tissue material can be considered *viscous*. This implies that the deformation of a viscous material depends on the history of the applied forces and not only on their instantaneous values. The simplest viscous material is the *Newtonian* viscous fluid, where the stress is proportional to the strain rate (Equation of Newton) :

$$\tau = \eta\dot{\varepsilon} \tag{2}$$

Most soft tissues are *viscoelastic*, combining elastic and viscous behaviors. For instance, the Maxwell (serial dashpot and spring) and the Kelvin-Voigt (parallel dashpot and spring) viscoelastic models associate linear elasticity with constant viscosity. Those materials are characterized by their *creep* and *relaxation* functions occuring when the material is subjected to a constant load or deformation (see figure 18). Finally, many anatomical tissues can be considered as biphasic, where the tissue is a mixture of a solid porous matrix and an incompressible fluid [MKLA80].

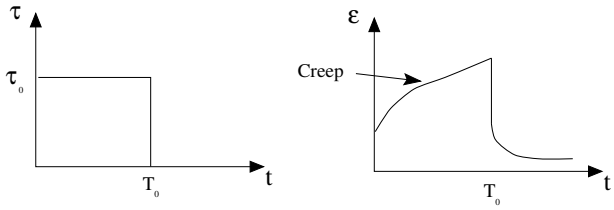

Figure 18: *Deformation response of a linear viscoelastic material under loading and unloading*

2.1.2. Interaction with rigid or soft bodies

The interaction between a soft tissue and surrounding bodies (surgical instruments, bones, soft tissue,..) can be decomposed into two different tasks : collision detection and computation of interaction forces (or displacements). Many algorithms have been proposed for detecting the collision between a moving point object and a static mesh. They are sufficiently efficient to sustain a 30Hz refresh rate with meshes of reasonable size (between 10 000 to 40 000 polygons).

However, in surgical simulation, those algorithms are not applicable because soft tissue models cannot be considered as static since several surgical instrument interact at the same time. Furthermore, for many surgical instruments, the interaction with tissue not only occurs at their extremity, but also along their length. Therefore, it is necessary to consider the surgical instruments as rigid moving meshes and soft tissue as deformable moving meshes. When detecting the collision between two moving objects, pre-computation is not possible and the computational time rises sharply. greater and it constitutes a bottleneck for

When considering the collisions between two deformable

meshes or the self-collisions of a deformable mesh, the computational complexity becomes prohibitive with current algorithms for real-time processing. However, recent techniques [TKH*05] offer possible research directions to overcome those major limitations. Distance fields (or implicit models), well suited for rigid objects, allow to quickly check if primitives lie inside/outside an object. Spatial subdivision methods store primitives/ grid correspondances within a hash table, and collisions are detected by comparing object entries. Image-space techniques use graphics hardware to project (render) images of the scene, and check collisions in 1D within the layered depth image. Bounding volume hierarchies (probably the most popular technique) use simple enclosing volumes for fast overlapping tests. Finally, stochastic methods are inexact methods that ensure control over computational time and collision detection quality.

Once a collision has been detected, the second task consists of computing the resulting interaction forces (collision response) to reach a non-penetration state. We usually distinguish *constraint dynamics* methods that incorporate reaction forces into the system [WGW90], *penalty methods* that enforce non-penetration through stiff springs [HFS*01], *exact methods* [Bar92] that analytically compute penalty terms by resolving non-penetration equations involving the relative acceleration between bodies, and finally *impulse-based methods* [MC94] that apply velocity impulses to guaranty non-penetration. The main limitation is that very little is known about the real interaction between biological tissues. Unlike material characterization, the experiments on tissue interaction must be performed *in vivo*.

2.1.3. Real Time Deformation

The real-time deformation of soft tissue is an important constraint for medical virtual reality systems. It has been clearly established that the immersion of the operator, therefore its ability to learn from a computer simulated system, is directly linked to the bandwidth of the simulator. An acceptable bandwidth for visual display is in the range of 20-60Hz while the acceptable bandwidth for haptic display is on the range if 300-1000Hz (300Hz is the free hand gesture frequency).

Two numbers are particularly important for accurate perception by the user : *latency* and *computation time* . Latency measures the time between sensor acquisition (for instance the position of the surgical instrument) and action (visual or haptic display). The computation time is that needed to update the geometric model. On multi-processor computers, those two numbers are not necessarily correlated.

Latency is critical for the user immersion. The hardware configuration of the system can greatly influence latency since communication between elements may be responsible for additional delays. In figure 19, we show the architecture of the simulation system used at INRIA [CDA96]. It is composed of one haptic display, a PC and a graphics workstation. There are several contributing causes to latency: communication between the haptic display and the PC, communication between the PC and the graphics workstation, the time taken by the graphics display, the computation time for collision detection, force feedback and deformation. Since much of the communication between elements are asynchronous, the total latency is not the sum of those delays but it is important to reduce them to their minimum values. The latency depends greatly on hardware, specifically on computation and graphics performance.

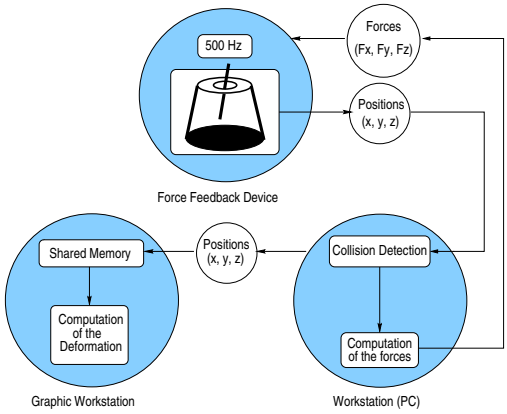

Figure 19: *Example of simulation systems.*

The computation time depends on the choice of the geometric and physical model of soft tissue (see section 2.2). If we write X_t as the position of the tissue model at iteration t, we will write C_t the computation time needed to compute the new position X_{t+1}. The computation time must be bounded ($C_t < C^\star$) in order to guarantee minimal bandwidth.

We distinguish between a *static* equilibrium equation and a *dynamic* law of motion depending whether X_{t+1} depends on the previous position X_t. When using static equations of the form $F(X) = 0$, the computed shape corresponds to the state of equilibrium, so takes no account of inertia or viscoelasticity.

With dynamic laws of motion, the current position of the tissue has an influence on its future position. Mostly, this can be modeled using a Newtonian law of motion having the corresponding differential equation :

$$m\frac{\partial^2 X}{\partial t^2} = -\gamma\frac{\partial X}{\partial t} + F(X) \tag{3}$$

Depending on the complexity of the deformation model, the computation of X_{t+1} may be performed interactively (see Section 1.3.4). The number of iterations must be bounded since the computation time C_t must be less than $C^\star$. A major difficulty consists of ensuring *synchronicity*, *i.e.* that the *numerical time* used for computation matches the *user time*

used for interaction. For instance to correctly estimate the speed of a vertex x_t, one should divide the position difference $x_t - x_{t-1}$ by the real time spent between two iterations. Adaptive approaches have been proposed to ensure bounded computational load for computing deformations [DDCB01].

Static equations have two advantages. First, they are faster to compute because no time integration is needed. Second, they are well suited for parallel algorithms or for asynchronous computation. For instance, in Cotin *et al.* [CDA96], we use a static formulation motion that enable us to decouple the force with the deformation computation. This is highly advantageous because, as noted before, the force computation must run at least at 300Hz but requires few operations (product of a matrix with a vector) whereas the deformation computation must run at 30Hz with a large amount of computing.

However, static laws are not able to model realistic deformations such as inertia or viscoelasticity. Dynamic laws of motion on the other hand, can model more accurate deformations but are more difficult to handle and computationally more expensive.

2.1.4. Tissue Cutting and Suturing

The ability to cut and suture tissue is of primary importance for designing a surgery simulation system. The impact of those operations in terms of tissue modeling is considerable. In fact, they imply that the geometric representation of tissue must change its topology over time [CDA99a]. The cost of such a topological change depends largely on the chosen representation (see section 2.2).

In addition, the behavioral model of the tissue must be adapted at parts where cutting or suturing occurs. Little is known about the stress/strain relationship occuring during and after cutting. The basic assumption is that the physical properties of tissue are only modified locally. However, in practise, cutting can greatly modify the boundary conditions between tissue and the surrounding organs which entails considerable change in terms of deformability.

Finally, when cutting volumetric or surface models, it is very likely that the new geometric and physical representation of tissue leads to self-intersections. The detection of self-intersections is computationally extremely expensive, therefore repulsive force between neighboring vertices are sometimes added to prevent self-intersections.

2.1.5. Force-Feedback Computation

Haptic display serves at least two purposes in a surgical simulator : kinesthetic and cognitive. First, it provides the sensation of movement to the user and therefore it greatly enhances it surgical performance. Second, it is used to distinguish between tissues by testing their mechanical properties.

However, the addition of a haptic display in a simulation system increases by a large factor its complexity and the required computational power [MRF*96] : an increase by a factor 10 of the required bandwidth, synchronisation between visual and haptic displays, force computation, ... Few papers have assessed the importance of haptic feedback in surgery [Mar96] [CDA99a] [DDCB01]. In general, it is accepted that the combination of visual and haptic displays are optimal for surgery training or pre-planning.

In video-surgery, the surgical instruments slide inside a trocard and are constrained to go through a fixed point. This entails substantial friction, specifically in laparoscopy where airtightness must be enforced. The friction of the instruments inside trocards perturbes the sensing of forces by the end-user. Despite those perturbations, it appears that it is still necessary to provide force-feedback for realistic user immersion.

The force computation depends on the chosen deformation model. Deformation models based on biomechanics naturally lead to physically meaningful forces and are very likely to provide an intuitive feeling to the end-user. On the other hand, with more *ad hoc* deformation models, for instance spring models, the force computed at each iteration may not correspond to intuitive sensation, especially if there is no continuous representation of the tissue.

Another problem with the use of haptic displays is the small number of commercially available systems. They often correspond to modified versions of joysticks dedicated to different applications (teleoperation, video games, general user interface,...). Therefore, they rarely meet the specific constraints of open and video surgery, in terms of workspace or encoder resolution. Current commercial equipment read the position of the end-effector and are force-controlled.

2.1.6. Visualization

Visual feedback is the most powerful perception channel. The quality of the visual rendering greatly influences user immersion and therefore the effictiveness of the simulator. In the past few years, substantial technological advances have enabled sharp decreases in the price of efficient graphics boards. With the existence of standard programming environment (such as OpenGL, Direct3D,...) and standard platform configurations, this trend should grow over the next five years to deliver high performance at reasonable price. It will have the effect of increasing the number of polygons that can be drawn at each frame. By combining improved graphics performance with anti-aliased texture mapping with large texture memory, a major improvement of the realism can be foreseen.

However, there are many problems in visualization for surgery simulation that cannot be solved simply by improving the raw graphics processing. For instance, after cutting or suturing the tissue model, modification of the texture coordinates must occur to reflect the change of topology. The development of volumetric texture images could be used instead of two-dimensional texture images. When simulating

video-surgery, the optical aberration of large angled endoscopes, as well as the intense light of the optical fiber, must be modeled for more realistic display. Finally, techniques for depicting bleeding, and the rendering of semi-translucid and filandrous structures, must be improved.

2.2. Deformable Tissue Models

In this section, we review the main existing models of soft tissue.

2.2.1. Surface or Volumetric Tissue Models

The geometric representation of deformable tissue may consist of surfaces or volumes. The choice between surface and volume based models is governed by two factors: computer efficiency and physical accuracy. In terms of computation, surface models are advantageous because they have less vertices than volumetric models for representing the same shapes.

However, most biomechanical models such as those presented in section 2.1.1, naturally call for volumetric representations rather than surface models. Surface models tend to give physically invalid deformations especially in regions that are thin. In figure 20, we show the deformation of a volumetric plate model. Similar experiments with a surface plate model would have entailed self-intersections. Furthermore, the behavior of volumetric models may take into account physical inhomogeneities, for instance due to the presence of lesions.

Finally, volumetric models are better-suited to the simulation of cutting or suturing operations. This is because those operations changes the geometrical and the physical nature of the model.

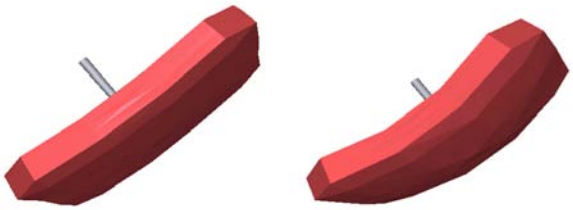

Figure 20: *Volumetric deformable model behavior.*

However, surface models may be relevant for modeling cavernous tissues such as vessels or the gallblader. In this case, the physical model of deformation can incorporate a representation of a liquid or of a gaseous pressure combined with surface tensions.

Here, we briefly review the main existing approaches, focusing of solid object expressed from the Lagrangian point of view. For more details, please refer to the recent exhaustive review on physically based deformable models by Nealen et al. [NMK*05].

2.2.2. Springs and Particles

The discretisation of the spatial continuous domain leads to a set of nodes that can be considered as particles with a certain mass, evolving in space under forces. As opposed to continuum-based methods, particle systems solve motion equation independently for each particle. This is a critical difference, allowing particle systems to be much more efficient in terms of computational speed. Indeed, particle systems, as they rely on a local description, lead to a much more sparse system that FEM, thus to a simpler resolution.

A particular implementation of particle systems is the well known mass-spring model, in which internal forces are straightforwardly computed from independent springs between particles. Spring models consist of a set of points linked by springs and dampers. In the simplest formulation, the equation of motion of a point i is :

$$\mu \frac{d^2\mathbf{r}_i}{dt^2} = -\gamma \frac{d\mathbf{r}_i}{dt} + \sum_{j \in N(i)} K_{i,j} \frac{(l^0_{i,j} - \|\mathbf{r}_i\mathbf{r}_j\|)\mathbf{r}_i\mathbf{r}_j}{\|\mathbf{r}_i\mathbf{r}_j\|} \quad (4)$$

where μ is the mass, γ the damping factor and $K_{i,j}$ the stiffness of the spring connecting point i and points j in the neighborhood $N(i)$ of point i.

Spring models have been used extensively for simulating the elasticity of soft tissue. Waters [Wat92] has defined springs on regular lattices for modeling facial tissue. He derived the two stiffness parameters of biphasic springs from stress/strain curves described in a biomechanical study [KGEB75]. Similarly, Delingette *et al.* [DSCP94] represented fat tissue elasticity as a network of springs on a 3-simplex mesh. Keeve *et al.* [KGG96] proposed a similar approach for modeling fat tissue in a craniofacial surgery simulation system but with the addition of a volume-preservation force intended to model the incompressibility of human tissue. Koch *et al.* [KG96] combined a finite element model for representing the skin surface, with a spring model to represent the fat tissue. The stiffnesses of the springs are derived from the intensity of voxels in a CT-scan image. The underlying assumption is that stiffness is proportional to tissue density and therefore to the Hounsfield units.

The main advantage of spring models is their ease of implementation since they do not require continuous parameterization. They have been used for static as well as for dynamic computation. Another advantage is their ability to model cutting or suturing simply by removing or adding connections between vertices. Mass-spring simulation is fast and can be easily parallelised and GPU accelerated [Mr05] (e.g. 100 000 particles at 100 frames/s).

However, for the soft tissue simulation, they suffer from the following problems :

- **Topological design**. The topology of springs and masses

is of great importance. Since a spring constrains the length between two vertices, the number of springs per vertex conditions the global behavior of the system. If the system is under-constrained, several rest positions are possible, and the system can fall into unwanted local minima. If the system is over-constrained, it tends to decrease the range of deformation. For tetrahedral meshes, the number of springs per vertex should be as close as possible to 6. Because of this difficulty in designing topologically a network of springs and masses, authors have usually organized the springs on regular lattices [Wat92] or on prisms with triangular base [KG96,KGG96]. Those choices however, impose the restriction of the geometric representation that it be organized in sets of parallel layers.

- **Validity of deformations**. The deformation induced by springs cannot easily be compared with those given by biomechanical studies because springs do not rely on continuum mechanics. However, for small deformations, a spring model behaves similar to a linear elastic finite element model, as verified by Keeve *et al.* [KGG96]. By linearizing equation 4, the stiffness parameter can be identified with the stiffness of a linear elastic model. However, for large deformations, spring models do not behave like a linear elastic material and comparison with nonlinear elastic finite element models is difficult. In order to identify spring parameters, given a behavior model, several algorithms have been proposed. In [JL95], Louchet uses a genetic algorithm for identifying the spring parameters to represent a cloth animation model. Deussen *et al.* [DKT95] bases the search for optimal parameters on simulated annealing.
- **Dynamic Behavior**. For dynamic spring models consisting of n nodes, for a given time step Δt and a given mass $\mu = m_{total}/n$, there is a critical stiffness K_c above which the numerical system is divergent. The relationship between K_c and the time step Δt is :

$$K_c \approx \frac{\mu}{\pi^2(\Delta t)^2} \approx \frac{m_{total}}{n\pi^2(\Delta t)^2} \tag{5}$$

  This relation, which is also valid for explicit linear elastic models, implies that in order to increase the stiffness of the model, it is necessary to decrease the time step. Since the computation time C_t is independent of the time step Δt, to maintain the refresh rate as high as possible, it is necessary either to decrease the stiffness K or to decrease the number of nodes n. Stiffness controls the propagation of constraints along the tissue model. High stiffness models tend to exhibit global behavior which is desirable for most soft tissues. In practise, we have found that greater time steps could be used with explicit linear elastic finite element models than for spring models. This implies that the range of possible dynamic behaviors of spring models is more limited than those of finite element models.
- **Visualization** Spring models are composed of mass-points and edges represented by springs. To visualize the tissue surface, it is necessary to define polygons from the set of edges and vertices. However, when cutting or suturing occurs it is necessary to update the set of visible faces. This suggests that the network of springs and masses should be built upon a manifold (surface or volume). A complete data structure of the manifold must be used in a way similar to finite element models.

Luciani [LJF*91] developed an animation system (CORDIS/ANIMA) based on particle systems for modeling complex physical phenomenon. The system is coupled with a force-feedback device for real-time interaction. Provot *et al.* [Pro95] defined super-elongated springs to increase the stiffness of its cloth model but with no guarantee of convergence. Cover *et al.* [CEO93] used springs to represent the surface of the gall bladder. They combined *home* forces with internal forces to enforce shape constraints. Similarly, Kuhn *et al.* [KKKN96], in the KISMET project, developed a surface model of the gall bladder based on springs, dampers and plastic elements where mass-points are connected to parent nodes. In that system, a NURBS representation is attached to points, giving a realistic rendering of the tissue. Stone *et al.* [RRF96] added a slip and a split threshold to a spring network, in order to model plasticity and fracture. Meseure *et al.* [MC97] proposed a surface representation of tissue similar to the hybrid model proposed by Terzopoulos [TF88a]. Springs of zero length are attached to a virtual rigid component. During deformation, both the positions of mass-points and the translation/rotation parameters of the rigid component are updated. This framework has the advantage of decoupling stiff behavior modeled by the rigid component and local deformation modeled by springs and dampers. Similarly, to reduce distortions due to linear strain formulation, Müller et al. [MG04] use non-rotated configuration to allow large deformations through simple computation. They also present a meshless implementation [MHTG05a] that makes use of point explicit registration (Section 1.3.4) to derive geometric internal forces and a stable integration scheme.

Gibson *et al.* [GSM*97] proposed a "ChainMail" model that does not derive from the equation of dynamics 3. Instead, the deformation model governs the displacement of all nodes, given a displacement on its boundary. Tissue is represented as a set of deformable voxels linked to their 6 nearest neighbors. When a node is pulled or pushed, neighboring links absorb the movement by moving slightly. If a link between two nodes is stretched or compressed to its limit, displacements are transferred to neighboring links. Within this framework, stiff behavior can be modeled for large displacements whereas compliant motion will be observed for small displacements. The "ChainMail" model is well-suited for real-time deformation and cutting operations, However, it is not clear how realistic the deformation is nor how feasible the real-time rendering of the tissue model is.

Other methods consider forces acting on elements such as tetrahedrons [CDA99a], or triangles [BW98] [VMT00], taking into account more that particle pair interactions, in

the evaluation of forces and force derivatives. They have shown promising results, in reproducing the various deformation modes of real materials, thanks to the formulation of the strain (small deformation strain or linearized version of the Green-Lagrange strain) at the element level.

2.2.3. Continuum-based methods

The finite-element method (FEM), the finite-difference method (FDM) and the boundary-element method (BEM) are well-established mechanical engineering techniques that have been widely used since the 70's for simulating tissues deformation. Their main idea is to discretise the continuous equation describing energy equilibrium (virtual work principle) inside an object [Zie77]. The spatial domain is subdivided into elements such as triangles, bricks or more frequently tetrahedrons [CDA99b] [OH99] [HFS*01] [TSB*05], and boundary conditions are defined (Dirichlet or Neumann boundary conditions). While the change of variables across a regular spatial grid is straightforwardly evaluated by discrete derivatives with FDM [TPBF87], FEM makes use of an interpolation shape function with limited support (e.g. B-splines, piecewise polynomial, linear basis, Bernstein-Bézier etc.), which is integrated in the energy equation. Finite elements with C^0 continuity, where the shape node consists of a vertex position are similar to finite difference methods. In the general case, FEM exploits higher order derivatives than FDM and requires fewer nodes, not necessarily regularly sampled (better approximation of organ boundaries). This leads to continuous representations with varying levels of continuity. A finite element model is fully defined by the choice of its elements, its shape function and its global parameterization between parameter space Ω and $\mathbb{R}^3$ ($\Omega \subset \mathbb{R}^2$ for surfaces and $\Omega \subset \mathbb{R}^3$ for volumes). For surfaces that are neither topologically planar or cylindrical, the parameterization can be problematic.

A finite element model is represented by the node vector $\mathbf{X}$. For static computation, the stress-strain relationships leads to $f(\mathbf{X}) = 0$, whereas for dynamic computation, the following Newtonian formulation is often used :

$$m\ddot{\mathbf{X}} + \gamma\dot{\mathbf{X}} + f(\mathbf{X}) = \mathbf{0} \tag{6}$$

Those equations can equivalently be derived by minimization of the bending energy through the principle of virtual work. The stiffness matrix expresses the strain-stress relationship and is consequently based on the underlying constitutive mechanical model. The integration of this differential equation can be performed using an semi-implicit or explicit scheme (see Section 1.3.4). In general, implicit schemes are unconditionally stable whereas explicit schemes are conditionally stable. This implies that smaller time steps must be used with explicit schemes. However, explicit schemes are simpler to compute and may not require the matrix inversion.

Membrane and thin plate energies have been largely used in computer vision [TWK87b], computer graphics [TPBF87] and for modeling elastic tubular surface tissue [Meg96]. These two belong to a family of regularizing energies : the controlled-continuity generalized spline kernels [Ter86]. Those quadratic energies have been used extensively because of their numerical properties (they lead to linear elastic forces). However they do not correspond to physical elastic energies. The membrane energy is a linearized version of the surface tension energy on soap films [Hil44]. The thin plate energy is a linearized version of the isotropic thin shell flexion energy [Ber96]. In particular, those energies are not invariant with respect to a change of parameterization. However, for small deformations, they can be considered as valid approximations. In [TF88a], Terzopoulos proposes a more general bending energy consisting of the sum of the square of the metric tensor and the curvature variation. He defines a hybrid formulation of deformation that includes a rigid component and a deformable component. With this new parameterization of deformation, invariant to rigid transformation, thin plate bending energy is used to model large deformations. In [CYTT92], a dynamic model similar to [TF88a] is proposed but the damping factor is replaced by the time derivative of the strain tensor. In [TF88c] a Maxwell and Voigt viscoelastic model has been implemented using a semi-implicit scheme with finite differences on a regular grid. The addition of plastic units enabled non reversible behavior to be modeled. However, those results required substantial computational power for relatively small grids (around 30×30).

Biological tissues are, in general, non-linear, anisotropic, hyperviscoelastic and quasi-incompressible materials. For simplicity and because of time-step restriction, linear elastic models have been widely used to model the deformation of soft tissues. In such cases, the stress/strain relationship is represented by a linear equation : $\mathbf{F} = \mathbf{KX}$. The rigidity matrix $\mathbf{K}$ depends on the rest shape geometry, the Young modulus E and the Lamé parameter λ. In most cases, only C^0 elements are used, leading to simple shape functions. However, for large deformations (Green strain), strain non-linearities are paid by high computational cost (the stiffness matrix depends on the displacement).

Pieper [PRZ92] and Keeve [KGG96] simulated fat tissue elasticity for plastic surgery with prismatic finite elements. Chen [CZ92] built a sophisticated muscle model based on biomechanical data where a linear elastic muscle is submitted to non-linear tendon forces. Despite the approximation on the muscle elastic behavior of the muscle, good correspondence with biomechanical has been observed. Gourret [GMTT89] computes the skin deformation of human fingers during the grasping of a soft object. The simulation models the interaction between two deformable bodies.

More complex elastic behaviors have been proposed by Terzopoulos. In [TF88c], he defines the square norm of

the metric tensor as the potential energy, which can be shown [BN96] to be equivalent to a St. Venant Kirchoff deformation model. Bro-Nielsen [BN95] used this non-linear elastic model for simulating craniofacial surgery with a finite difference scheme defined on cubic lattices. In [Kai96], Kaiss *et al.* proposes a complete model of the eye/trepan contact for predicting eye deformations under surgery. The sclera and cornea are modeled as St. Venant-Kirchoff materials and transversely isotropic materials. The penetration of the trepan requires the re-meshing of the cornea represented by hexahedrical elements with 27 nodes. The computation time was of the order of an hour on a powerful workstation.

To simplify FEM formulation, the boundary element method (BEM) has been used [JP99]. Its priciple is to transform the differential operator defined in the domain to integral operators defined on the boundary (divergence theorem) by assuming that the material is linear and homogeneous. Using this idea for individual tetrahedrons, the finite volume method (FVM) [TBHF03] [TSB*05] makes use of the divergence theorem to simplify force and strain computation by handling tetrahedron triangles instead of the boundary of the Voronoi region associated to each node. Assumption in FVM is the constant strain tetrahedra and linear variables change (linear basis).

Continuum-based models have been used widely to compute soft tissue deformations under mechanical constraints. In the past years, real-time models have been developed as an alternative to spring models. Sagar [SBM*94] developed a virtual environment for eye surgery simulation where the cornea deformation is modeled as a non-linear elastic material (Mooney-Rivlin material). The finite element solver computed the cornea deformation every second while the graphics module was able to provide a 10Hz refresh rate. In [CDA96], Cotin *et al.* describes a hepatic surgery simulator where the liver is represented as a linear elastic volumetric model with static constraints. By precomputing the response of surface vertices to position constraints, the liver model can be deformed in real-time. Furthermore, the force-feedback computation and the liver deformation computation can be decoupled to achieve optimal haptic display (500Hz) and visual display (30Hz). Similarly, Bro-Nielsen [BNC96, BN96] decreased the computation time of a linear elastic model with a semi-implicit scheme, by condensing and explicitly inverting the reduced stiffness matrix in a preprocessing stage. Video frame-rates of 15-20 frames/second were obtained with this method. Koch et al. [KRG*02] have used a high order finite element model for maxillofacial surgery planning from medical images. By taking advantage of the linear nature of the static or dynamic equation, the methods of Cotin [CDA99b] and Bro-Nielsen [BNC96] decreases the computation time of finite element models by at least factor of 100.

In figure 21 we show an example of real-time hepatic simulator developed at INRIA [CDA96, Cot97]. The user manipulates a force-feedback simulation platform and feels the contact between the virtual instrument and a liver model.

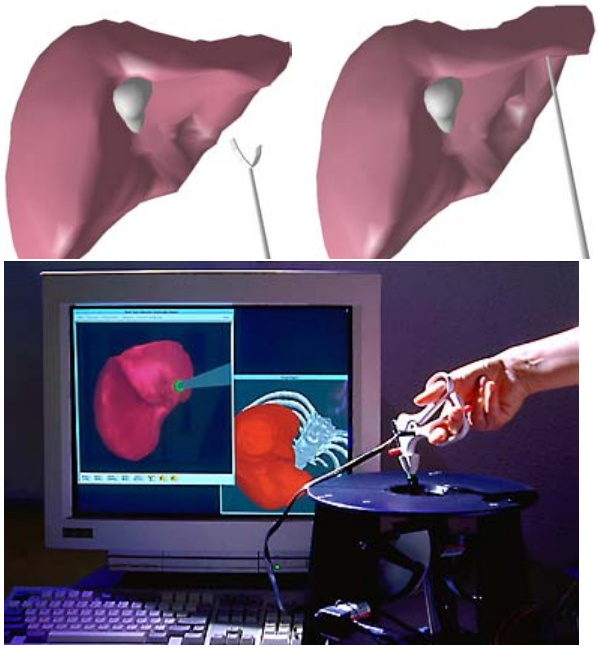

Figure 21: *Hepatic surgery simulation platform at INRIA*

As a conclusion, finite element models are well suited to compute accurate and complex deformation of soft tissue. However, it is extremely difficult to get real time performance on a moderately powerful workstation using finite element models. But for linear elastic models, only valid for small displacements, it is possible to achieve real-time deformations. Unlike spring models (see section 2.2.2), there is no restriction on the stiffness value of the model with respect to the time step Δt when using semi-implicit or static schemes.

The cutting or suturing operation requires the finite element model to be re-meshed. When using structured elements such as rectangular, prismatic or hexaedral elements, the cutting is often constrained to occur along a given direction [Kai96, SR95]. With unstructured elements such as triangular or tetrahedral elements, more general cut planes may be designed at the cost of greater complexity. Finally, visualization of finite elements is well suited for graphics hardware since it consists a rendering visible elements. After cutting or suturing a volumetric model, it is necessary to update the list of visible facets. Figure 22 shows the real-time cutting of a tetrahedral finite element mesh of a kidney [Cot97]. The kidney has a linear elastic behaviour and an explicit computation scheme prevents the re-computation of the global stiffness matrix. To handle cuts and mesh topology changes, Bro-Nielsen and Cotin combine their method with a particle

system ("mass-tensor" network), where internal forces are formulated with respect to the linear elasticity framework (as opposed to mass-spring networks) [CDA99b]. This work has been extended by Pincibono et al. [PDA03] for handling anisotropy and non-linear elasticity.

Another approach by Debunne et al. [DDCB01] is to use multi-resolution (space/time adaptive sampling) with a weighted finite difference integration technique to approximate the Laplacian and the gradient of the divergence operators of the linear elasticity operator. They use a linearized version of the Green-Lagrange strain (invariant to rigid motion) allowing large deformations. Similarly, a local approximation of the resolution is performed in [OH99], through mass lumping, to reduce computational charge. This is known as *explicit FEM*. All these techniques are encouraging since they try to merge the simplicity of particle systems and the accuracy of finite element methods. However, no method currently allows realistic simulation (anisotropic non-linear viscoelastic materials under large deformations) of large multi-organs complexes (neighbourhood constraints). This is enhanced by the difficulty of validating and comparing such methods.

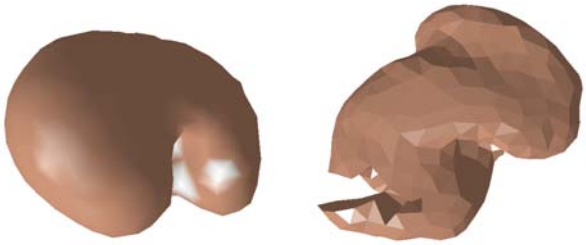

Figure 22: *Cutting of a kidney finite element model.*

2.2.4. Other Deformable Models

Other models of deformable bodies have been proposed in computer graphics. Even if such methods have seldom been applied to medical simulation, they are relevant to the development of real time simulation of soft tissue.

For instance, implicit surfaces defined by potential fields attached to skeletons is a representation well-suited for collision detection computation and the modeling of very soft objects. In [CGD97], Gascuel shows that material stiffness is related to the gradient of potential field and that such models can support cutting or suturing. Other global transformations such as extended free-form deformations [Coq90] or modal analysis [SP91] are potentially of interest, for instance for modeling the effect on abdominal tissues of breathing.

3. Medical visualization and applications

3.1. Introduction

The large amount of data provided by the modern clinical modalities can give a lot of diagnostic information to radiologists and clinicians. Data usually encode maps of physical parameter values in 2D, 3D or 4D and it is not obvious in general, how to show them to the doctors allowing a clear understanding of organs morphology or functionality and eventually obtain quantitative diagnostic values.

To obtain visual representations really helping diagnosis and useful for a variety of new applications like image guided surgery, augmented reality in operating room, surgical training etc. is a big challenge for the Computer Graphics community. In the same way, usability and the easiness of use of new interfaces implementing these applications is another big challenge for experts in Human Computer Interaction.

The simplest data analysis method, still widely applied by radiologists, consists of just showing 2D images (printed or on screen), acquired or resliced, using ad hoc color mapping and windowing to enhance the structures of interest.The intrinsic 3D or 4D nature of most of the data, however, suggests the necessity of finding effective ways of displaying the whole dataset in order to give an immediate interpretation of the organ morphology and functionality.

Standard Computer Graphics techniques, like volume rendering, surface extraction/rendering, can be applied to reach this goal, obtaining photo-realistic representation of organs in 3D. Several visualization enhancements can then be applied to give the physicians useful diagnostic information in addition to the simple morphology. It is possible, for example, to show selectively different image features and to add to the scene results of segmentation procedures, or additional functional parameters related to the morphological information (i.e. pathological region localization, motion information, functional information) and represented using classical visualization methods (color codes, transparency, vectors, streamlines, etc.). The importance of medical data rendering is not, however, limited to a simple diagnostic use. 3D/4D visualizations can be applied in surgical planning, to increase the understanding of organs morphology and or to perform accurate measurements. 3D virtual objects can be also superimposed to real images to provide the physicians with relevant information during surgical procedures (Augmented Reality).

Another relevant application of advanced visualization is related to human body simulation, and to the realization of surgical simulators for medical training. Especially in this last case, the visualization task is particularly critical due to the fact that the model rendered is interactively changed by the user and therefore it is not possible, for example, to use pre processing methods often applied to speed up the rendering process.

Furthermore, diagnostic images are often used to build anatomy teaching tools for students and doctors, requiring user friendly visualization interfaces. For this kind of applications, and also for simple simulations, web based 3D data visualization tools have also been proposed and tested [Joh07]. Standard web3D technologies, like VRML and X3D have been used for the visualization of medical models and also to realize simple clinical applications. The Medical Working Group of the X3D consortium is developing an open interoperable standard for human anatomy representation.

All these specific application of 3D data visualization (pre-operative planning, augmented reality in operating rooms, surgical simulation, teaching, etc.) require not only the application of advanced visualization techniques, but also a careful design of the interaction in order to give the end user the possibility of easily navigating the scene, selecting the views and performing and qualitative and quantitative evaluation.

The section reviews the main techniques and trends in medical visualization and a few example applications. It is organized as follows: subsection 3.2 presents the state of the art of the 3D applications currently available in clinical departments. Subsection 3.3 describes the basic rendering techniques (volume rendering, surface rendering, etc.) used ti display medical data. Time critical visualization, rendering acceleration and large model visualization are discussed as well. Subsection 3.4 shows how a new generation of 3D displays can help the analysis of medical data. Subsection 3.5 presents applications of real time visualization in surgical simulators. Subsection 3.6 discusses generic issues related to performing quantitative measurements on volume data and an example application for measuring reconstructed vascular structures for interventional planning. Subsection 3.7 discusses recent advances in virtual endoscopy techniques and related clinical applications. Subsection 3.8 deals with visualization issues related to medical applications of augmented reality.

3.2. Visualization of diagnostic imaging in clinical centers

CT and MRI scanners are since their introduction computerized digital machines, but for several years (and even now) radiologists used to print the images to sheets of film in order to hang them to the light-boxes. The recent introduction of PACS archives and visualization workstations, however, gives the possibility of displaying 3D information and performing more complex data analysis.

Several advanced tools are already available in clinical environments, even if their use in real applications is still limited.

Medical workstations provided by company like GE, Siemens, Barco/Voxar, Vital Images, Viatronix allow the visualization of 2D slices (with windowing and color mapping options and measurement tools) and several pre-processing and volume reconstruction methods. Different volume rendering methods (see section 3.3.1) are usually implemented, supporting 3D windowing, i.e. the interactive mapping of voxel values to transparency values. 2D analysis can be applied not only on original slices (acquired once, usually in axial direction, but also, with the Multi Planar Reconstruction (MPR), along arbitrary cutting planes (standard coronal, sagittal views, oblique MPR. Also Curved Planar Reconstructions are sometimes available, mapping curved surfaces to viewing planes, in order, for example, to show arteries, or the human spine or teeth. Of course this option requires particular procedures to define the surfaces, usually based on human interaction. MPR can be extremely useful to avoid measurement errors due to tortuosity of vessels and organs. Ideally, it would, of course preferable to perform quantitative evaluation during the 3D data analysis; however the interaction with 3D world, especially if not performed with immersive Virtual Reality environments, presents relevant problems. Problems related on measurements in 3D are discussed briefly in subsection 3.6. Clinical examples of 3D quantitative evaluation in vascular analysis and virtual endoscopy are also described.

Other useful visualization methods available in clinical workstations are those related to multi-modal analysis. Different image registration methods can be applied (see [MV98a]for a review), and relevant problems may arise, especially if markers are not available and nonrigid transformation are required. From the visualization point of view the problem is then the graphical representation of the multiple sources of information (data fusion).

Finally, packages available to clinicians offer specialized packages for a selected number of clinical studies (e.g. Vascular analysis, Virtual Endoscopy, Heart). They implement several advanced visualization methods and segmentation/reconstruction tools (i.e. vessel tracking, surface reconstruction, etc). The potential applications of the 3D visualization and analysis are, however, quite limited by the necessity of specific education and training for the effective use of the 3D segmentation/visualization/analysis systems. Diagnostic applications also require a careful analysis of the medical work-flow and of the user interaction in order to have an effective use of the systems. It is difficult, in fact, to realize easy to use segmentation procedures and it is difficult to implement optimal 3D interaction paradigms, especially if standard monitors and not Virtual Reality immersive displays (see subsection 3.4) are used.

3.3. Medical data rendering techniques

3.3.1. Volume rendering

Volume rendering in general covers the creation of images from volumetric data. These three-dimensional data sets can

be said to represent one or more functions over a three-dimensional space. A volumetric data set can, for example, be created from a stack of images from an image acquisition system, such as a medical Computed Tomography (CT) scanner. This technique is computationally demanding and has only recently begun to be applied in medical visualization tasks. Generally speaking, algorithms for rendering of volume data can be divided into two main categories: Image order and object order approaches. In the past, both of these techniques have proven to be useful for specific applications, though image order approaches seemed to be more generally applicable because the computational complexity scales with image rather than object size, which often makes it a better choice for applications where the exact size and structure of the volume is unknown.

Image order approaches As the name suggests, image order approaches iterate through the parts of the image that should finally be generated and try to find all possible contributions to this part. With these parts being screen pixels most of the time, the algorithm tries to find all objects that change the appearance of a certain pixel. The most widely known image order approach is ray-casting [Lev88], where a ray is cast through each pixel and sampled at regular intervals. The contributing samples of the volume are composited to a final pixel color, which in the end is a (more or less exact) approximation of the volume rendering integral along that ray. Primary advantage of this algorithm is the fact that it is much more dependent on screen resolution than on object size, which also makes it easily scalable by reducing the resolution for quick in-between renders. Also, the number of objects in the volume does not influence rendering speed as much as with object order approaches (this situation changes a little bit with the implementation of techniques like early ray termination and empty space skipping). Obvious disadvantage of image order approaches is that, if no further measures are taken, very sparsely populated volumes will be rendered a lot slower than with object order approaches, because a lot of pixels will be checked (and thus a lot of rays started) that never even hit an object. On the other hand, techniques like early ray termination make sure that in the case of very dense volumes, only those objects are rendered that really contribute to the final image.

Object order approaches In contrast to image order approaches, object order approaches iterate through all parts of the object - in most cases voxels - and determine their contribution to the final image. The most popular object order approach is splatting [LH91, Wes90], where a footprint of the current object is generated and splatted onto the image plane. This works particularly well if there are only very few non-empty voxels inside a large volume. In all other cases, the probability is quite high that a substantial part of the objects will not even be visible in the final image, thus wasting a lot of computational effort. Apart from speed concerns, object order approaches can have a substantial advantage in terms of memory consumption. This is mainly due to the fact that virtually all image based approaches rely on some kind of regular grid to store the data, thus reserving the same amount of memory for empty and non-empty voxels. In the case of object order approaches, usually only non-empty voxels are stored, which can reduce memory demands significantly if the volume is not heavily populated. However, a drawback of this approach is that the appearance of the footprint limits the quality of the final image - zooming in on a splatted image will reveal the structure of the footprint quickly, making the quality of the precomputed kernel essential for good results [Wes90]. Though various extensions of this algorithm have been published [XX06, CRZP04], the quality of rendering is still inferior to image-based approaches in most cases.

Ray-casting fundamentals The most common software ray-casting algorithm, was proposed by Marc Levoy in his initial publication on ray-casting [Lev88]. It divides the process of image generation into six distinct steps:

- preparation of volume densities along a regular grid, resulting in voxel values for each discrete position;
- classification of voxels, mapping each voxel density to a respective opacity value;
- resampling of sample opacities at the discrete sampling positions along the ray;
- shading, mapping each voxel density to a color value;
- resampling of voxel colors at the discrete sampling positions along the ray;
- compositing step, by computing a final pixel color from the vector of shaded samples and respective opacities.

The first step is to prepare the volume densities for further processing, arranging the acquired values at certain positions inside the volume along a regular grid. This step might include correction for nonorthogonal sampling grids, patient motion while scanning or even contrast enhancements, interpolation of additional samples or pre-filtering of noisy data. The output of this step is an array of prepared values which is again used as input for the shading and classification steps. In the case of shading, phong shading is used regularly because it represents a good trade-off between speed and quality. The classification performs the essential step of assigning each voxel a respective opacity value. This opacity value can be a function of various parameters, like voxel density, normal vector direction or gradient magnitude. Standard ray-casting modes would include setting the opacity above a certain threshold to 1. This results in rendering of the first intersection with a value above the threshold along the ray, commonly referred to as iso-surface ray-casting or first-hit-ray-casting. Another common classification strategy is the simple definition of opacities for all density values via a transfer function, resulting in a visualization of translucent tissue that is used primarily for direct volume rendering. Including the normal vector or gradient magnitude into the classification function is primarily used for nonphotorealistic renderings and thus mostly found in specialized ap-

plications where these strategies provide a better insight into certain structures. With shading and classification strategies defined, the actual algorithm is performed by casting rays into the volume and resampling the voxel densities at evenly spaced locations along that ray. The color and opacity values are usually trilinearly interpolated from the eight voxels closest to each sample location. This provides a good trade-off between simple nearest-neighbor interpolation (always take the value from the closest voxel) and more complex filter kernels like tricubic interpolation, that yield better results at higher computational demands.

Finally, these color and opacity values have to be composited to the final pixel color. In order to exploit strategies like early ray termination front-to-back compositing is usually used, starting the ray at the viewing plane and casting rays through every single pixel until a certain alpha threshold near 1 is reached. Front-to-back compositing calculates the summed pixel color by adding further samples according to the following formula:

$$C_{out} = C_{in} + (1 - O_{in}) C_v O_v \quad (7)$$

$$O_{out} = O_{in} + (1 - O_{in}) O_v \quad (8)$$

where C_{in} and O_{in} are the input color and opacity values before adding the current sample, $Cout$ and O_{out} are the output color and opacity values after adding the current sample, and C_v and O_v are the color and opacity values of the sample point (i.e. the result of the trilinear interpolation of classified and shaded samples). Depending on the compositing strategy, different rendering modes can be achieved, as shown in figure 23. For medical visualization, the most common compositing strategies are Maximum Intensity Projection and X-Ray simulation.

Maximum Intensity Projection Maximum Intensity Projection (MIP) is a compositing strategy of direct volume rendering, where, instead of composing optical properties, the maximum value encountered along a ray is used to determine the color of the corresponding pixel [ME05] (see fig. 23).

An important application area of such a rendering mode, are medical datasets obtained by MRI (magnetic resonance imaging) or CT (Computer Tomography) scanners. Such data sets usually exhibit a significant amount of noise that can make it hard to extract meaningful iso-surfaces, or define transfer functions that aid the interpretation. MIP is considered very useful for visualizing angiography data sets, since the data values of vascular structures are higher than the values of the surrounding tissues. The biggest drawback is that intensity channel completely unmasks depth information; so, in normal 2D displays users are not able to disambiguate features and recognize which objects are at the front and which are at the back. Figure 24 shows a rotational angiography

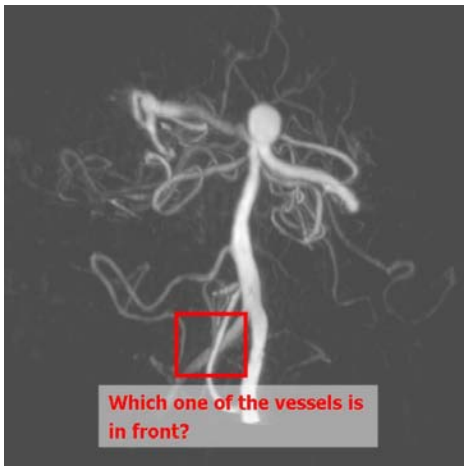

Figure 24: ***Depth oblivious MIP angiography rendering.*** *MIP volume rendering of a rotational angiography scan of a head with aneurysm. The positions and the crossings of vascular structures are not detectable, or wrongly interpreted.*

scan of a head with aneurysm. By using MIP volume rendering technique, only the contrasted vessels are visible. In a 2D view, the positions and the crossings of vascular structures are not detectable, or can be wrongly interpreted. This is because the technique does not provide any depth information in a single image and does not support occlusion. In subsection 3.4 we show how new generation spatial 3D displays can help solving these kind of ambiguities.

X-ray In classical X-ray volume rendering, a viewing ray is cast through the center of each pixel and the line integral of the intensity is evaluated along the given ray [ME04].

Figure 25: ***Depth oblivious X-ray CT rendering.*** *X-ray volume rendering of a CT Scan of abdomen and pelvis. In 2D views, it is not possible to distinguish between front and rear.*

In this case internal parts of the volume are visible, but

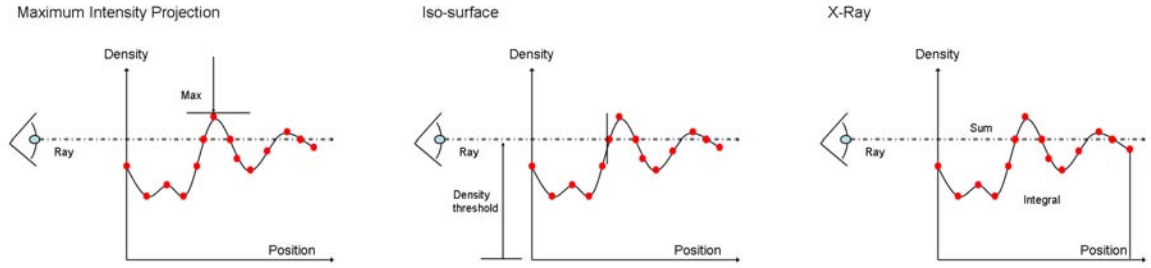

Figure 23: ***Various compositing strategies****. In ray casting different schemes can be chosen, in order to compose color and opacities to final pixel color. Maximum Intensity Projection consists of taking the maximum density along the ray, while iso–surface or first–hit strategy consists of taking the first density above a given threshold along the ray; finally x–ray is obtained by integrating voxel contribution along the ray.*

depth information is not maintained. Hence, view disambiguation tasks are very hard to accomplish using 2D. Figure 25 shows a CT scan of abdomen and pelvis, containing also a stent in the abdominal aorta. In this case no contrast agent was used to enhance blood vessels. A X-ray volume visualization is able to highlight bone and vascular structures, but in 2D views it is not possible to distinguish between front and rear.

Transfer functions Transfer functions are fundamental to direct volume rendering because their role is essentially to make the data visible: by assigning optical properties like color and opacity to the voxel data, the volume can be rendered with traditional computer graphics methods. Good transfer functions reveal the important structures in the data without obscuring them with unimportant regions. Commonly, in medical practice transfer functions are limited to one-dimensional (1D) domains, meaning that the 1D space of scalar data value is used to assign color and opacity. Often, there are features of interest in volume data that are difficult to extract and visualize with 1D transfer functions. Many medical datasets created from CT or MRI scans contain a complex combination of boundaries between multiple materials. This situation is problematic for 1D transfer functions because of the potential for overlap between the data value intervals spanned by the different boundaries. Nowadays, the usage of multi-dimensional transfer functions is considered, but the application to volume rendering is really complicated. Even when the transfer function is only 1D, finding an appropriate transfer function is generally accomplished by trial and error. This is one of the main challenges in making direct volume rendering an effective visualization tool. Levoy first introduced two styles of transfer functions, both two-dimensional, and both using gradient magnitude for the second dimension [Lev88]. One transfer function was intended for the display of interfaces between materials, the other for the display of isovalue contours in more smoothly varying data. Other work involving multi-dimensional transfer functions uses various types of second derivatives in order to distinguish features in the volume according to their shape and curvature characteristics [KKH02].

Accelerating techniques Real-time methods to directly convey the information contents of large volumetric scalar fields are still considered a challenge to the computer graphics community. Based on the observation that the capability of a single general-purpose CPU is not sufficient to achieve interactivity or even real-time for large data sets in general, considerable effort has been spent on the development of acceleration techniques for CPU-based volume rendering and on the design and exploitation of dedicated graphics hardware. Among others, one research direction has lead to volume rendering techniques that exploit hardware assisted texture mapping. Fundamentally, these systems re-sample volume data, represented as a stack of 2D textures or as a 3D texture, onto a sampling surface or so-called proxy geometry [EKE01, RSEB*00]. The most common surface is a plane that can be aligned with the data, aligned orthogonal to the viewing direction, or aligned in other configurations (such as spherical shells). The ability to leverage the embedded trilinear interpolation hardware is at the core of this acceleration technique. Furthermore, ray-casting approaches can be accelerated by employing GPU implementations supporting perspective projection without decreasing performance [KW03]. The entire volume is stored in a single 3D texture, and resampling is performed by fetching samples from this texture with trilinear texture filtering. For each pixel, a fragment program steps from sample location to sample location in a loop and performs compositing until the volume is exited or full opacity is reached. For further information related to GPU based volume rendering techniques, refer to [EHK*04].

Non-photorealistic techniques Historically, most volume rendering techniques are based on an approximation of a realistic physical model. It was noticed, however, that tradi-

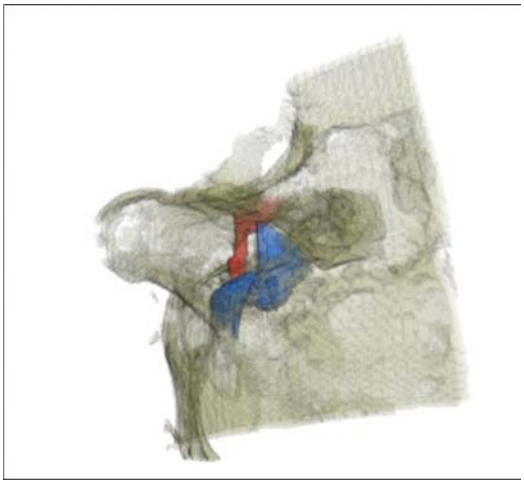

Figure 26: ***Illustrative rendering.*** *Illustrative rendering of temporal bone CT scan dataset. Data courtesy of Prof. Sellari Franceschini, University of Pisa.*

tional depictions of the same types of data as found in medical textbooks, for example, deliberately use non-realistic techniques in order to focus the viewer attention to important aspects. Using abstraction, visual overload is prevented leading to a more effective visualization. Recent approaches have considered this fact, leading to an increased interest in illustrative volume visualization [BG07]. Approaches for illustrative volume visualization frequently employ nonphotorealistic rendering techniques to mimic the style of traditional illustrations [RE01]. They take advantage of the illustrators century-long experience in depicting complex structures in an easily comprehensible way [SES05, LMT*03]. Figure 26 shows an example of illustrative rendering obtained by employing transparent direct volume rendering and exaggerated color map in order to highlight vascularity inside a temporal bone CT scan dataset.

3.3.2. Surface rendering

Surface rendering consists of visualizing data-sets by generating polygonal surfaces representing the anatomical part of interest. This option simplifies the rendering complexity, but requires pre-processing and causes loss of relevant data. Triangulated surface rendering with textures and lighting effects is optimally implemented by the graphics hardware, so the major problems of surface rendering applications are related to the extraction of the polygons representing the model (segmentation) and the simplification of the representation that is necessary when the number of triangles becomes huge (large scale visualization). In medical applications huge surface meshes are not usual; several techniques are however available for mesh simplifications and to build multi-resolution data structures allowing efficient visualization of complex surfaces [CGG*04]. For the creation of surfaces from medical data variants of "marching cubes" algorithm [LC87] are commonly employed. They are able to extract iso-surfaces representing ideally organs boundaries. To recover more accurate and topologically controlled surface representation of organs from diagnostic images, several advanced 3D segmentation techniques are applied (deformable surfaces, front propagation, model based methods), usually specialized for each particular image type and medical task (i.e. vascular segmentation, liver segmentation etc.). Current state of the art in image segmentation techniques is described in section 1 of this tutorial.

3.4. Non-standard visualization on future displays

The representation of 3D scenes on standard 2D monitors limits the possible applications of advanced visualization. This is the reason why 3D displays have been proposed for clinical applications.

Stereo viewers using standard technology (polarized glasses or head mounted displays) are used to improve depth perception in particular medical data visualization tasks and in clinical applications like Virtual Endoscopy. The overhead for the visualization software only consists in the fact that two images should be rendered at interactive frame rates. Some displays use head/eye tracking to automatically adjust the two displayed images to follow the viewer eyes as he moves the head.

Auto-stereoscopic displays do not require googles or head mounted devices. They are usually based on the use of lenticular lenses or parallax barriers that allow the user who keeps his head in a particular viewing positions, to perceive a different image with each eye giving a stereo image. 3D displays of such kind are already available also in commercial medical systems: the SpatialView 3D Medical Workstation offers, in fact an auto-stereoscopic monitor with head tracking coupled with a vision based freehand intuitive gesture based 3D navigation and analysis (SpaceTool). New generations of volume displays are becoming available and are expected to be widely used in the near future. One way to obtain 3D images is to display 2D images on rapidly rotating surfaces. The Perspecta Display by Actuality systems [Fav05], creates 3-D imagery by projecting a sequence of 2x198 images onto a screen rotating at or above 900 rpm. The company offers to research institutions a medical application of the display that is able to visualize DICOM datasets (limited in depth and size) with the 3D method. The LightSpaces Technology DepthcCube uses a similar principle but with a different geometry: different LCD panels stacked in a cube and switching from transparency to images.

A powerful kind of volumetric display can be however realized similarly to auto-stereoscopic screens, but synthesizing with multiple projections 3-D light fields nearly equivalent to volumetric images. Continuous parallax monitors

have been obtained in this way by Holografika (HoloVizio), which realized high resolution real 3D spatial displays allowing users to see a 3D scene moving freely (horizontally) within the viewing angle (up to 70 degrees), and supporting real time operation.

3D spatial display in details The largest display Holografika display [ABF*06] uses an array of 64 XGA projectors and a holographic screen with a diagonal of 72 inches. Smaller screens(32, 26 inches) are also produced.

In these systems, module images generated by the projectors, are determined by geometry and not associated with specific view directions. Each module emits light beams toward a subset of the points of the holographic screen and each point of the holographic screen is hit by more light beams arriving from different modules. The light beams propagate to address fixed spatial positions that are independent from the viewer's position. Many modules contribute to each view of the 3D image, thus no sharp boundary occurs between views, and the display offers continuous and smooth change at different image areas, resulting in a truly 3D experience with continuous horizontal parallax. The holographic screen transforms the incident light beams into an asymmetrical pyramidal form. The cut of this light distribution is a long rectangle, where the vertical size of the rectangle is the vertical field of view, while the horizontal size corresponds to the neighboring emitting directions. This configuration corresponds to the horizontal-only-parallax capability of the display. With the same principles it would to possible to provide vertical parallax as well, at the cost, however, of another order of magnitude increase in data size, rendering times, and system complexity, for little gain in the visual performance in standard settings. The screen is a holographically recorded, randomized surface relief structure that enables high transmission efficiency, controlled angular distribution profile and eliminates moiré and chromatic effects.

The angular light distribution profile introduced by the holographic screen, with a wide plateau and steep Gaussian slopes precisely overlapped in a narrow region, results in a highly selective, low scatter hat-shaped diffuse characteristics. The result is a homogeneous light distribution and continuous 3D view with no visible crosstalk within the field of depth determined by the angular resolution (see figure 27).

Multiple video streams are generated by an array of 16 PCs, connected to the display through four DVI connections each. Each PC runs a server that controls a graphics frame buffer. The server is responsible for generating images associated to a fixed subset of the display rendering modules. In the 72 inches prototype, each PC generates 4 XGA images using two double head NVIDIA boards controlling a 4096x768 frame buffer. In order to support legacy graphics programs and to simplify the development of new holographic applications, an OpenGL compatible front-end has been developed. The front-end runs on a client PC and

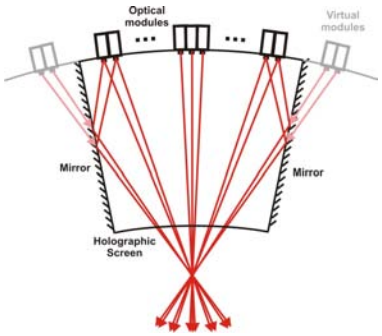

Figure 27: *In the Holovizio display a large number of light beams create a spatial point.*

looks to applications like an ordinary OpenGL library which transparently broadcasts the graphics command stream to the dedicated cluster driving the holographic display. The client PC is connected to the cluster through dual Gbit ethernet links. The back-end servers listen to the network and decode the stream of multicast graphics commands coming from the client. Once decoded, commands are interpreted and sent to the local renderer. The interpretation of the graphics commands involves modifying the way they are generated according to parameters available from the local configuration service, in order to transform the original central view into the view associated with each of the associated optical modules. For each of the optical module views, the graphics commands of the current frame are re-executed, with the following modifications: the original perspective matrix is replaced with a matrix that matches the module's specific position and viewing frustum; a geometrical calibration is performed, to correct nonlinearities in the display/optical geometry; a light calibration is performed to correct the intensity and contrast differences response of the optical modules; an angular resolution correction is performed for depth dependent anti-aliasing.

Volume rendering of diagnostic data on the holographic display New trends in medical visualization research focus on advancing medical visualization by combining 3D rendering techniques with novel spatial 3D displays able to provide all the depth cues exploited by the human visual system [ABGP07].

For example, we can consider as test cases two depth oblivious techniques commonly employed in medical visualization: Maximum Intensity Projection and X-Ray volume rendering. These techniques were originally developed for visualizing images on 2D displays, and have the drawback that 2D views are really ambiguous, since no depth cues are present in the intensity channel, the occlusion is not considered, and shading is not present. The addressed question is whether a 3D display is able to recover lost 3D info, and it was shown that it is possible. Figure 28 shows a MIP ren-

Figure 28: ***MIP angiography rendering.*** *Direct capture from the spatial 3D display taken from different positions in order to disambiguate vascular structures.*

Figure 29: ***X-ray CT rendering.*** *Direct capture from the spatial 3D display taken from different positions in order to immediately understand spatial information.*

dering of an angiography CT scan dataset visualized on the Holovizio spatial display, while figure 29 shows a sequence of pictures taken from direct observation of an X-ray rendering of the CT abdomen and pelvis dataset in the Holovizio spatial 3D display. The set of pictures were taken from different positions in the display workspace, and try to convey part of the visual information provided by the display. The pictures here are merely illustrative, by the way they show that 3D spatial displays are very effective in resolving depth relationships of overlapping, if combined with direct order-indipendent volume rendering approaches (like X-ray and MIP).

An important issue related to the use of immersive 3D displays is that it is necessary to develop new interaction techniques in order to make their use simple and effective. While in simulation tasks the interaction design is simple, because the interaction is designed mimicking the real procedure, generic paradigms for 3D model analysis (that can be based, for example on special devices like spacemouses, gloves or on vision based tracking) did not emerge clearly.

3.5. Application: time-critical rendering for surgical simulation

A real time interactive surgical simulator presents different challenges respect to simple interactive rendering with changing viewpoints. The scene, in fact, changes with time and, even if the frequency constraints are not strict as for the haptic response, it could be difficult to reproduce a reasonable reproduction of the operation scene in real time. As anticipated in the description of basic techniques, specialized approaches can be applied for accelerated volume rendering and surface rendering of dynamically changing data. We describe here two examples of virtual reality training systems employing these techniques: the first is a training system for temporal bone surgery developed at CRS4 in the context of the European project IERAPSI (IST-1999-12175)), while the second is a simulator for the Phaco-emulsification and capsulorhexis tasks in the context of cataract extraction training.

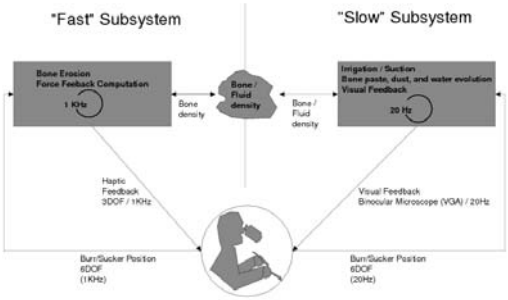

Figure 30: ***Temporal bone system architecture****. The system is divided in a "fast" subsystem, responsible for the high-frequency tasks (surgical instrument tracking, force feedback computation, and bone erosion), and a "slow" one, which is essentially dedicated to the production of data for visual feedback.*

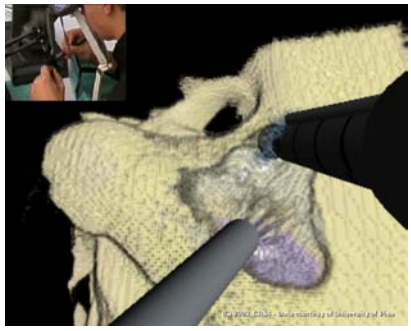

Figure 31: *Snapshot taken from IERAPSI surgical system, where a ENT trainee is performing a mastoidectomy. The inset shows the hardware setup, composed by two PHANToM haptic interfaces and a NVision binocular display*

3.5.1. IERAPSI temporal bone simulator

The IERAPSI prototype virtual reality training system has been developed to simulate completely the temporal bone dissection procedure. Force feedback on haptic devices (Sensable Phantom) is realized using a physically motivated burr-bone interaction model, loosely based on Hertz contact theory, that includes haptic forces evaluation, the bone erosion process and the resulting debris and a multi-resolution representation of the bone characteristic function to adaptively trade simulation quality with speed [AGG*02]. Visual rendering (showing the bone erosion and a particle system based simulation of blood, water and mud) is obtained through a dynamic direct volume rendering technique based on OpenGL register combiners that is able to render shaded representations of dynamically changing rectilinear scalar volume in parallel to simulation threads. The system is based on patient-specific volumetric object models derived from 3D CT and MR imaging data. The performance constraints dictated by the human perceptual system are met by exploiting parallelism via a decoupled simulation approach on a multi-processor PC platform. Specifically, the system is divided in a "fast" subsystem, responsible for the high-frequency tasks (surgical instrument tracking, force feedback computation, and bone erosion), and a "slow" one, which is essentially dedicated to the production of data for visual feedback (see figure 30). The system has been evaluated by experts (Senior ENT Surgeons) and demonstrated realistic haptic and visual rendering, and the ability of simulating complete mastoidectomy procedures [AGG*04]. Figure 31 shows a snapshot taken from IERAPSI surgical system, where a ENT trainee is performing a mastoidectomy. The inset shows the hardware setup, composed by two PHANToM haptic interfaces and a NVision binocular display.

3.5.2. The Cataract surgical simulator

Cataract is a clouding of the eye's natural lens, normally due to natural aging changes, and involving at least half of the population over 65 years. Cataract extraction is the only solution for restoring a clear vision, and nowadays is probably the most frequently practiced surgical procedure. To train surgeons to perform this task, a specific virtual reality simulation system for cataract surgery, involving the capsulorhexis and phaco-emulsification tasks has been developed. The simulator runs on a multiprocessing PC platform and provides realistic physically-based visual simulations of tools interactions. The current setup employs SensAble PHANToM for simulating the interaction devices, and a binocular display for presenting images to the user [AGP*06, AGG*06]. Similarly to IERAPSI temporal bone simulator, the system is divided in a fast subsystem, responsible for the surgical instrument tracking (100 Hz), and a slower one, essentially dedicated to the production of data for visual feedback. The slow subsystem is responsible for the global simulation of the eye, and for the interaction of the devices with the cornea, the anterior camera membrane and the crystalline lens. The algorithms used to control the simulations are local in character, leading naturally to a further break-up of the slow subsystem in components, each dedicated to the generation of a specific visual effect, and thus possibly to a parallel implementation on a multiprocessor architecture. The system runs on two multiprocessor machines connected with a 100 Mbit Ethernet link. The first machine is dedicated to the high-frequency devices tracking tasks (100 Hz), while the second machine concurrently runs the low frequency task (20-25 Hz): eye simulation, anterior camera simulation, and crystalline lens simulation. Since the low-frequency tasks do not influence high-frequency ones, the two machines are synchronized using one-way message passing. Differently from the bone drilling case, since force feedback returned by surgical instruments is here nearly imperceptible and the main way the surgeon have to under-

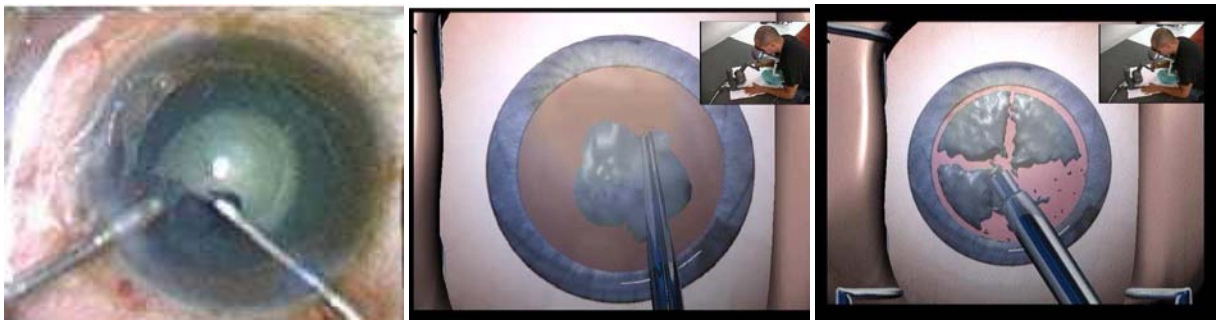

Figure 32: *Comparison between real (left picture) and virtual (center picture) capsulorhexis, and a snapshot from a virtual emulsification (right picture)*

stand the effort really exerted on the patient eye is the visual feedback provided by the tool position related to the environment and the eye globe displacement. This fact enhance the need for a realistic real time visualization. Figure 32 shows a comparison between a picture taken from a real capsulorhexis and a snapshot taken from the simulator (left and center pictures), as well as a snapshot taken from the virtual emulsification task (right picture). The inset shows the hardware setup for the virtual reality training simulator system. In this case, visualized data have no relationships with images acquired on real patients, but accurate reproduction of realistic textures have been superimposed to the dynamic polygonal scene rendered by exploiting hardware acceleration. However, the time critical part of the system is the physically based simulation of the interaction between anatomy and devices during the virtual intervention. The anterior camera for the capsulorhexis was modeled as triangular faceted mesh moving under forces generated by a spring-mass model joining different particles with several effects (following the spring-mass technique described in section 2.2). The emulsification task was simulated by modeling the eye lens as a collection of simplexes built from a tetrahedral mesh. Particles are eroded during the procedure and a shape matching approach [MHTG05b] is used to animate the residual part. The model parameters are tuned by expert surgeons, and overall realism is considered sufficient for training purposes. The system is currently used in specialization courses.

3.6. Model measurement for diagnosis and pre-operative planning

Diagnostic images are not only used for qualitative analysis, but can also be quantitatively measured in order to define the relevance of a pathology or to plan a surgical procedure. Measure can be static (i.e. tumor size, vessel calcification volumes, etc.) or dynamic (i.e. left ventricle ejection fraction) and can be performed on 2D slices or on 3D volumes. The use of 3D measurements is extremely important to avoid errors that are due to slicing, for example on vessel length and sections (fig. 33). The use of dynamic and 3D quantita-

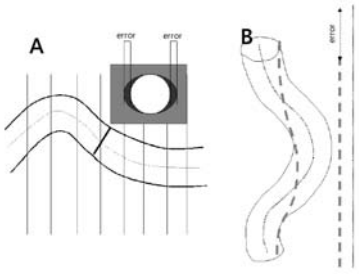

Figure 33: *Large errors are often caused by the use of 2D measurements on wrong planes to evaluate vascular lengths, diameters and tortuosity*

tive measurements is indeed important to improve the quality of diagnosis operative planning, but, in order to introduce such new techniques in the work-flow, two important things should be taken into account:

- A clinical validation of the reconstruction/measurement procedure is required
- The interaction with 3D scenes is not simple and specific interaction modes should be designed to obtain an effective use
- It is often necessary to change the approach of radiologist and surgeons to the particular procedure and make them appreciate the benefits of the new method

Commercial 3D systems are widely used only for very few precise tasks, e.g. intra-cranial or surgery. One of the problems for using 3D quantitative analysis is the difficulty in interacting with 3D scenes.

User interfaces based on the classic WIMP (Windows, Icons, Menus, Pointing devices) paradigm are not well suited for navigating/exploring a 3D scene and a post-WIMP 3D interaction general paradigm has not yet clearly emerged.

3D measurement modes supported by advanced medical workstation modules (i.e Voxar Cardia Metrix, Vessel

Metrix or Colon Metrix) usually limit the interactivity and perform pre defined analysis and visualizations given some input hints in 2D images.

It is not simple to develop, as would be required for the creation of a fully immersive virtual archive environment, user friendly tools avoiding complex user interaction modes, keeping at the same time the advantages of showing 3D paths and spatial relationships between brain areas. Furthermore, it is not easy for all users to move in the 3D world with situational awareness without using immersive virtual environments.

Discussion and general tests on 3D measurement and interaction methods have been presented in recent works. In [PTSP02] 3D measurement tools like distance lines, rulers, angular measurements, interactive volume approximations are presented. In [RSBB06] the interaction style when 3D measurements are performed is discussed and the importance of stereoscopic view and haptic feedback is discussed. Authors present a 3D measurement toolkit developed for the Virtual Liver Surgery Planning and also measured accuracy and time necessary to perform particular tasks (angles, volumes and distances estimation) and conclude that "3D measurements which build an integrated part of a surgical planning environment can be better carried out in a VR environment rather than on a 2D desktop-based application".

3.6.1. Pre-operative measurements with Web3D technology: an example

In the AQUATICS European project (IST 1999-20226), a tool for the web based measurement of Abdominal Aortic Aneurysms morphology has been developed, tested and validated with the help of interventional radiologists. Measurements were performed on a simple 2D web interface with ad hoc solutions making easy to perform 3D length volume and angles measurements without the use of 3D displays. The measurements tool developed have been, however, also modified to be used with the HoloVizio 3D collaborative displays as a test of the new technology.

Endovascular repair of abdominal aortic aneurysms requires a preliminary accurate assessment of patient's specific anatomy. The measurements of the geometry of the aorta is therefore extremely important. The true aneurysm diameter and its growth ratio are fundamental parameters to evaluate the risk of a rupture and to compare it with the risk of a surgical intervention. Fig. 34 shows distances and angles that have to be measured with sufficient precision for the planning of the surgical interventions and the selection/design of prostheses.

Reconstruction of the aortic model The model used for aortic measurement was a complex one defined by a watertight surface mesh representing the vascular lumen, other

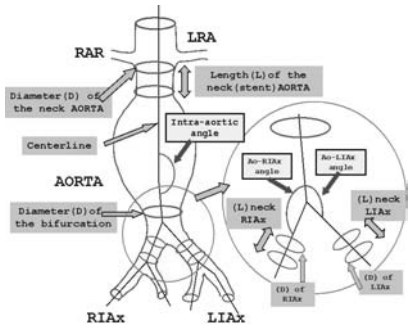

Figure 34: *A large number of measurements are necessary, including the length and diameter of the proximal aortic "neck", the length and diameter of the distal "cuffs", and the length of the two graft limbs. Measurements for an endovascular tube graft are only slightly less complex.*

Figure 35: *Example of skeleton extraction. White line: Shortest path at high resolution. Black line: Final snake-based centerline extracted.*

surface meshes representing calcification and external vascular walls including thrombus and a curve skeleton with a tree structure. The vascular lumen was extracted using a deformable surface algorithm (fast simplex mesh balloon). The user of the reconstruction tool just initialize the surface as a small sphere inside the lumen and the forces applied make it grow to detect the borders. A reparametrization of the surface with topological changes can be obtained using marching cubes in the case the surface becomes auto-intersecting. The curve skeleton extraction is realized with a voxel coding method [ZT99] improved by using a snake-based regularization.

Let us describe shortly the procedure, that is similar to those used to recover centerline paths in Virtual Endoscopy (section 3.7. It is based on the computation of two distance maps (or potential fields), one computing the distance from borders (BSC), and the other from a seed point (SSC). Taking as starting points local maxima of the SSC with high values, shortest paths reaching the seed or the previously extracted branches are found and then centered using the BSC map. In our implementation we also applied a regularization step driven by the BSC map, making the lines centered

and smooth. With this method we obtained results compliant with the requirements: continuous curves connected in a tree structure and locally centered in the volume. Fig. 35 shows an example of the procedure, that was sufficiently fast and reliable for our needs. The external surface of the aorta was recovered by slicing perpendicularly to the vessel direction the data set and using constrained snakes initialized outside the lumen in order to reach the external boundaries. Shape constraints have been introduced due to the presence of tissues with HU value close to the thrombus outside the vessel. To recover calcifications we applied therefore the well known "marching cubes" algorithm [LC87] with a threshold chosen to represent the calcification boundaries. Even in this case, however, it is necessary to customize the procedure and to consider some peculiarities of the problem. Details on different methods to recover vascular trees ad centerlines can be found in [GZ06]

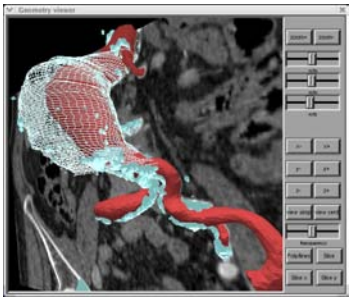

Figure 36: *Hybrid visualization of different vascular components(lumen, thrombus, calcium) superimposed to the CT data.*

Measurement interface Patient specific models built from CT data could then be remotely and collaboratively measured with a web based tool realized with open standards (VRML, External Authoring Interface, now replaceable with X3D and Scene Authoring Interface). Arterial models have been coded in VRML, including Javascript code to support guided navigation, measurement procedures, image visualization, etc. The VRML scene was included in dynamic

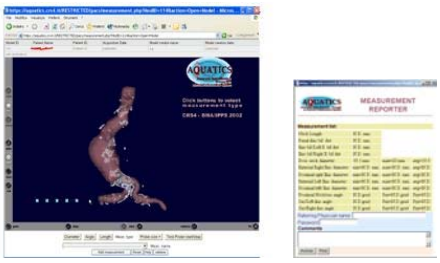

Figure 37: *Example of model visualization and measurements on the web interface.*

Figure 38: *Aneurysm phantom realized at University of Innsbruck and used to validate the aortic measurement system.*

web pages, including the measurement interface, the Electronic Patient record and the measurement reporting interface. The Aquatics Measurement application shows the segmented vessel and give the possibility of interactively hide or show its part, navigate the scene or choose pre defined views and perform 3D measurements of lengths and centerline paths, vessel diameters and volumes. The interface allows the interventional radiologist to perform easily all the quantitative measurements required to design the endovascular prosthesis to be inserted in the aorta 37.

The tool is a good examples of integration of 3D contents in standard web applications. Despite the large diffusion of Web applications, the use of Web 3D standards and of, in general of 3D applications on the web is still limited due to several reasons, including compatibility problems between different browser implementations and platforms. However, the example show the possibility of creating 3D visualization and analysis systems that can be used collaboratively by different remote users to obtain better diagnosis and procedure planning.

Validation The web based aortic measurement tool has been validated in several ways. Different technicians performed reconstructions using the segmentation tool and three clinical specialists used the web interface to measure the parameters for endovascular procedure planning. More than 40 patient specific aortic models have been reconstructed and models of a synthetic phantom 38 have been recovered from CT scans for validation. AQUATICS measurements resulted compatible with phantom's true data and patient data measurements done manually by radiologists using standard methods. The t-test showed a very good correlation between the measurements obtained on phantom with the Aquatics system and the true measurements of the phantom ($p < 0.0001$), demonstrating the reliability of the system. The correlation between different observers and reconstructing operator was also tested with the Spearman rank test and again a statistical significant correlation was proved ($p < 0.0001$), proving the reproducibility of the measurements.

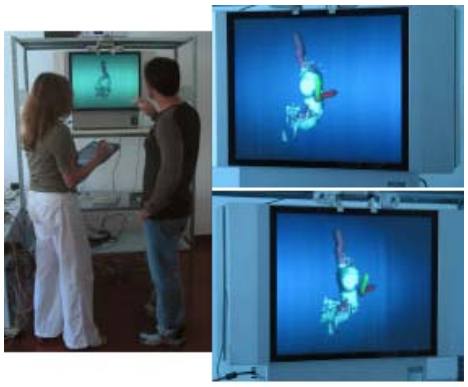

Figure 39: *Left: the collaborative aneurysm analysis session: an operator selects and check measurement types, while the other interacts with the 3D model. Right: Aneurysm rendered on the holographic display, imaged at the same time from two different angles.*

3.6.2. Collaborative measurement/evaluation on the 3D spatial display

The analysis of abdominal aortic aneurysms has been chosen also as a demonstration of the possible collaborative analysis performed with the holographic display. We realized a test application where the measurement interface is implemented as a dynamic web page, and the user can access it through a PDA [BFG*06]. When a patient-specific reconstruction is selected with the touchscreen, the rendering application (messages are passed through wireless connection and HTML protocol) loads the model and display it on the holographic screen. Since objects rendered on the holographic display appear floating in fixed positions, it is possible to naturally interact with them with a 3D user interface that supports direct manipulation in the display space. This is achieved by using tracked 3D cursors manipulated by users. Multiple cursor control interfaces have been developed, using both commercial 3D trackers (Logitech 3D mouse) and custom-made wireless solutions (camera based tracking of pointers, using a wireless USB interface for buttons). With this interface the same measurements for the endovascular repair planning of the AQUATICS application are supported, and a collaborative analysis is made possible.

3.6.3. Diagnostic use of 3D models: Wall stress and blood flow simulations

The enhancement of morphological models is not necessarily limited to selective structural visualization or geometrical measurement support. Even without acquiring functional data through the use of MRI or PET/CT scanners, it is possible to recover functional information with diagnostic relevance through the use of physical simulation. In the vascular field, for example, it is possible to simulate, given geometrical models and material behaviors, the deformation of

Figure 40: *Wall stress mapped on the aneurysm surface.*

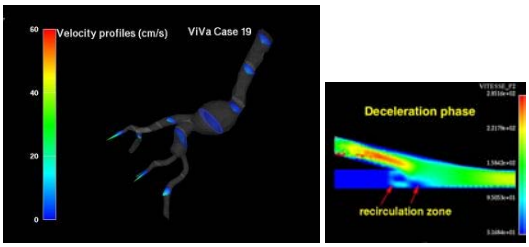

Figure 41: *Blood flow simulation results, displayed on 3D models or on 2D projections with standard methods like profiles, streamlines or color codes, can be used to show anomalous regions correlated with high risk of rupture or plaques formation*

the vascular boundaries when the blood pressure is applied. Fig. 40 shows the result of a finite element structural simulation showing the wall stress when a blood pressure cycle was applied on a geometry reconstructed for the AQUATICS project.

Figure 41 show results of some test blood flow simulations realized with a Navier Stokes finite element solver developed at CRS4 for the ViVa project [GSG*98]. The solver was parallelized with a domain decomposition technique.

Recent advances in simulation methods should give the possibility of modeling also the fluid-structure interaction in order to reconstruct a more reliable flow field and accurate wall stress estimation. For example in [FLBW06] a turbulent fluid-structure interaction simulation is then run on a commercial solver where the blood is treated as a non-Newtonian fluid. The transient velocity profile at the inlet boundary is taken from Doppler-ultrasound measurements. Other values, like density, Young's modulus, and Poisson coefficient, are based on material-specific data that have been adapted during test simulations.

Local simulation can be also coupled wit 1D full body circulation models recently developed and simulating the propagation of pressure and velocity waveforms in 1-D arterial models [KQFL03]. Accurate flow and wall shear stress

evaluation have a relevant diagnostic value, because they have been proved to be an indicator of vessel rupture and plaque formation. Furthermore, pre-operative simulation can be used to chose an optimal geometry for an endovascular prosthesis. In [FLBW06] the simulation results have been used to investigate the pressure and flow patterns in patient specific models of AAAs before and after stent-graft implantation.

3.7. Interaction techniques for medical applications: Virtual Endoscopy

Virtual endoscopy is one of the few well established application where virtual navigation in the 3D acquired data with clinical evaluation and measurement is effectively applied in the clinical work-flow. The basic idea is to simulate the clinical procedure of inserting a catheter inside the tubular structure lumen rendering the scene viewed by the real camera.

Current imaging modalities (i.e. Helical CT scanners, MRI) provide data at a sufficient resolution to make the reconstructed scene sufficiently accurate for several purposes. It should, however, considered that the accuracy of the rendered scene depends also on the quality of the segmentation, often user dependent, and the risk of imaging artifacts should be taken always into account. Different goals can be obtained through the use of Virtual endoscopy, from education, diagnosis, interventional planning and also enhancement of real endoscopy by registering volume data and virtual views with the real endoscopic view. A detailed review on Virtual Endoscopy techniques applied in research and clinical practice is presented in [Bar05].

Applications of Virtual Endoscopy have been mainly developed for the investigation of the colon, blood vessels (virtual angioscopy), of the trachea (virtual bronchoscopy). Similar applications have also been tested for the planing of endonasal interventions.

The advantages of virtual procedures are the non-invasivity, the lower cost, the complete control of lighting and orientation, the absence of access limitations. The real procedure, however, presents the advantage of offering a better resolution, texture information, and the possibility of interacting with tissues and directly perform the actual intervention. Furthermore, it is possible that for virtual procedures requiring a subsequent real intervention planned on their basis, changes happen between the image acquisition and the intervention.

Especially for Virtual Colonoscopy, that already applied in clinical environments, advanced processing techniques offer then several add-ons like the possibility of performing automatic polyp detection and highlighting in colonoscopy, or the possibility of visualizing unfolded surfaces [VWKG01].

While for other applications only research prototypes are available, several commercial systems are available for CT based colonoscopy (Siemens Syngo, Philips EasyVision Endo 3D, Ge Avantage Windows, Viatronix V3D, Vital Images Vitrea, Voxar 3D ColonMetrix, Tiani Jvision).

The examination is extremely useful due to the high incidence of colorectal cancer and the necessity of early detection in order to have successful treatment and the fact that the optical procedure is invasive and expansive. Virtual colonoscopy consists of pumping air into the colon and acquiring multiple high resolution CT scans in different positions to remove the effects of fecal residuals.

Once the elaborated volume are rendered, it is possible to interact with the virtual colon simulating the real procedure using automatic, guided or free navigation along the tubular structure centerline path, usually extracted semi-automatically with algorithms similar to that described in section 3.6.

Different rendering methods have been applied for generic virtual endoscopy. Surface rendering would be the simplest solution for efficiency, but the algorithm has disadvantages such as the low quality of visualized image and the loss of the volume data. All the non-segmented data are not represented, and this means that if the procedure is used for a complex surgical planning activity, all the interested structures must be reconstructed (with difficult, user dependent and time consuming techniques) before the navigation. Clinical validation showed also the superiority of Volume Rendering techniques over Surface Rendering for diagnostic purposes [HLS*03].

Fortunately, algorithmic solutions, like the use of accelerated empty space transversal [WKB99] and technological evolution of current generation of GPUs with advanced pixel shaders allow the realization of GPU-based ray-casting delivering high frame rates and interactivity for either perspective or orthogonal projections. These cards allow the realization of real time perspective volume rendering and hybrid volume/surface representations [SHN*06].

Older systems used also 2D/3D texture mapping for fast rendering, but could not exploit shading effects. Figure 42 show a snapshot of the research virtual endoscopy system developed at CRS4 in 1998, where the visualization was realized in this way [GPZT98]. The system implemented also stereoscopic rendering and measurement tools to quantify vessel stenoses. A particularity of the system was the physically based virtual camera control: we interpreted input from the six-degree-of-freedom device used to control the camera (spaceball) as forces and torques applied to a virtual camera model. A viscous friction force field proportional to volume opacity is used to avoid penetration in opaque areas. This way, the endoscopic camera was always confined in the interior of a cavity. To detect walls during ray-casting, we used an accumulated-opacity algorithm tracing the ray from front to back and stopping when the accumulated opacity was larger than a user settable threshold.

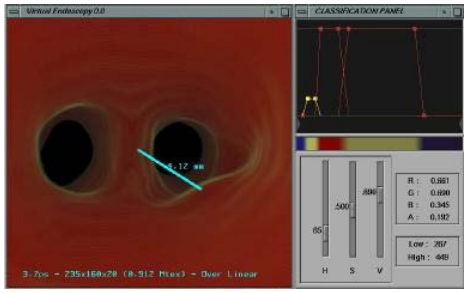

Figure 42: *The Virtual endoscopy prototype realized at CRS4.*

This method is effective and can overcome the difficulty of freely navigating without constraints. Another approach for effective navigation allowing real time interaction and full colon view are use centerline constraint for camera path an free camera orientation. Nonphotorealistic rendering methods can improve the efficacy of Virtual Colonoscopy. For example, to remove false detection of fecal residuals as polyps a recent technique consists of using a translucent rendering mapping colors related to the density over the colon walls, reducing the interpretation time required for the virtual analysis (one of the drawbacks of the 3D analysis) [KLKB05]. Recent studies showed that with Virtual colonoscopy it is possible to measure polyps better than with optical colonoscopy [PCL*07].

3.8. Image guided surgery and augmented reality in operating rooms

Another relevant application of advanced visualization techniques consists in using the 3D data to help the surgeon in performing his tasks. Some surgical procedures, in orthopaedics, or neurosurgery are indeed already planned and executed on the basis of image analysis and reconstruction and commercial system to plan and realize these interventions are available. For example, in neurosurgical tumor resection, surgeons can decide how to access the region through image analysis, and can then use devices guiding the instrument on the basis of the images thanks to the rigid registration of the patients skull with respect to the devices and the image reconstructed volume. Devices can be tracked with magnetic or optical systems. The rigid registration is usually obtained with fiducial markers. Deformations of tissues with respect to the acquired morphology limits the use of such techniques (Brain-shift problem).

In operating room surgeons can also be helped by the use of intra operative imaging (Ultrasound, MRI or angioscopy, C-Arm) These systems are extremely useful, even if the visualization tasks related to them are not particularly challenging for computer scientists.

Different is the case of the clinical use of Augmented Reality. Augmented Reality (AR) consists in overlying the real scene and the virtual image based one, with an interactive real time visualization requiring a perfect registration of the two scene and the tracking of the operator egocentric view. Different kinds of displays are applied for this effect: the most popular are probably head mounted displays (HMD), integrating virtual objects into the real environment in the original field of view of the user. With this technique hand-eye coordination is easy, in particular for keyhole surgery. Drawbacks can be found in the weight and display resolution. Another possibility for in-situ visualization is to use a display panel mounted on a swivel arm and placed between the surgeon and the patient, as for example done in the Medarpa project [SS04]. Accurate tracking of patient user and display position (performed with magnetic or IR systems) is in this case necessary to have an accurate result. If operating microscopes or binoculars are used in the surgical procedure, virtual images can be inserted into the optical paths of the devices and no additional display is necessary. A head mounted see-trough displays like the Varioscope AR [BFH*00] is a generic tool that adds two VGA displays to optical binoculars, and tracking the position of the head can realize the projection transformation (with Tsai's calibration method) of the virtual scene over the real one in real time. In other procedures it is possible to fix cameras and mirrors to the gantry of an imaging device obtaining augmentation without a tracking system. Intra-operative devices which can be augmented this way are endoscopes or laparoscopes used in the majority of minimally invasive surgeries.

Video see-through is an alternative to optical see through that overlay the virtual scene to a video acquisition, registering in this case camera position and orientation with the patient and the data [FNFB04].

The main technological challenge of AR is the accurate real time calibration of the scenes that is necessary to avoid latency or jitter of the virtual objects. Different focus for real world and computer graphics scenes can also cause eye strain and irritation. Due to these problems, augmented reality is still a research field and there are few clinical applications of widespread use. Several initiatives are, however, carried on by research groups to solve problems and to introduce AR in the clinical practice. Among these it is of particular interest the activity of the Marie Curie action ARISER (http://www.ariser.info/), involving eight relevant European partners with strong experience in the field.

3.9. Discussion

We presented a brief description of 3D medical data visualization methods and an overview of example applications. From this overview it appears clear that the application are heterogeneous and the visualiztion methods are quite different according to the different end users. Surgical simulators presents time critical aspects, related to the necessity of re-

alizing real time volume rendering and/or real time tracking and calibration. Other applications usually presents weaker time constraints and their problems rely mainly in their usability, or in the use of new generations of volumetric displays. An interesting remark on the diffusion of volume data in clinical centers is that it created a crisis in the radiology work-flow that requires a complete reorganization of the activity of radiologists and physicians. 3D processing and visualization are necessary to exploit and correctly interpret volume data. The role of technologists in this process can be relevant also in the clinical practice. 3D laboratories where radiologists and experts in image processing and visualization work together are already active (and billing for 3D reconstructions). Radiologists probably prefer to have the full control on the diagnostic process, but the complexity of the 3D reconstruction procedures (reformations, segmentation, visualization) requires new competences and often the use of expensive hardware. Advances in visualization algorithms and new generations of graphic hardware will help to develop easy to use workstations allowing simple 3D analysis and measurements of volume data, moving to the 3D labs only complex procedures like surgical planning [Rid07].

Finally, for the realization of Virtual Humans and physical simulators, that is the main subject of the tutorial, it is expected that the complexity of the reconstructed models will grow relevantly in the future and the large size of the data sets will require specific solution for visualization. Multi resolution schemes have been recently proposed to speed up the rendering of huge data sets on commodity hardware, also for generic triangular meshes [CGG*04]. For large arbitrary surface models, a multiresolution scheme for out-of-core coding and interactive inspection is presented in [GM05]. These kind of techniques, already applied with success in industrial applications, could be used to implement methods for interactive visualization of large and complex human models (static or with pre computed animations).

4. References

References

[ABF*06] AGOCS T., BALOGH T., FORGÁCS T., BETTIO F., GOBBETTI E., ZANETTI G.: A large scale interactive holographic display. In *Proc. IEEE VR 2006 Workshop on Emerging Display Technologies* (Conference Held in Alexandria, VA, USA, March 26 2006, 2006). CD ROM Proceedings.

[ABGD02] ASAKAWA D., BLEMKER S., GOLD G., DELP S.: In vivo motion of the rectus femoris muscle after tendon transfer surgery. *J. of Biomechanics 35*, 8 (2002), 1029–1037.

[ABGP07] AGUS M., BETTIO F., GOBBETTI E., PINTORE G.: Medical visualization with new generation spatial 3d displays. In *Eurographics Italian Chapter Conference* (Conference held in Trento, Italy, February 2007), Eurographics Association. To appear.

[ACP02] ALLEN B., CURLESS B., POPOVIC Z.: Articulated body deformation from range scan data. *ACM Transactions on Graphics 21*, 3 (2002), 612–619.

[AFP00] AUDETTE M., FERRIE F., PETERS T.: An algorithmic overview of surface registration techniques for medical imaging. *Medical Image Analysis 4*, 3 (2000), 201–217.

[AGG*02] AGUS M., GIACHETTI A., GOBBETTI E., ZANETTI G., ZORCOLO A.: Real-time haptic and visual simulation of bone dissection. In *IEEE Virtual Reality Conference* (Feb. 2002), pp. 209–216.

[AGG*04] AGUS M., GIACHETTI A., GOBBETTI E., ZANETTI G., ZORCOLO A.: Hardware-accelerated dynamic volume rendering for real–time surgical simulation. In *Workshop in Virtual Reality Interactions and Physical Simulations (VRIPHYS 2004)* (September 2004). Conference held in Colima, Mexico, September 20-21, 2004.

[AGG*06] AGUS M., GIACHETTI A., GOBBETTI E., ZANETTI G., ZORCOLO A.: Real time simulation of phaco-emulsification for cataract surgery training. In *Workshop in Virtual Reality Interactions and Physical Simulations (VRIPHYS 2006)* (November 2006), Eurographics Association. Conference held in Madrid, Spain, November 6-7.

[AGP*06] AGUS M., GOBBETTI E., PINTORE G., ZANETTI G., ZORCOLO A.: Real-time cataract surgery simulation for training. In *Eurographics Italian Chapter Conference* (Conference held in Catania, Italy, 2006), Eurographics Association.

[AHB87] ARUN K., HUANG T., BLOSTEIN S.: Least-squares fitting of two 3-d point sets. *IEEE Transactions on pattern analysis and machine intelligence 5* (1987), 698–700.

[AT01] AUBEL A., THALMANN D.: Interactive modeling of the human musculature. *Proc. of Computer Animation* (2001).

[AWJ90] AMINI A., WEYMOUTH T., JAIN R.: Using dynamic programming for solving variational problems in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence 12*, 9 (1990), 855–867.

[Bar92] BARAFF D.: Dynamic simulation of non-penetrating rigid bodies. *PhD thesis, Department of Computer Science, Cornell University* (1992).

[Bar05] BARTZ D.: Virtual endoscopy in research and clinical practice. *Computer Graphics Forum 24*, 1 (March 2005), 111–126.

[BCA96] BARDINET E., COHEN L., AYACHE N.: Tracking and motion analysis of the left ventricle with deformable superquadrics. *Medical Image Analysis 1*, 2 (1996), 129–149.

[BD05] BLEMKER S., DELP S.: Three-dimensional representation of complex muscle architectures and geome-

tries. *Annals of Biomedical Engineering 33*, 5 (2005), 661–673.

[Ber96] BERNADOU M.: *Finite element methods for thin shell problems*. Masson, 1996. ISBN 0-471-95647-3/2-225-84940-4.

[BFG*06] BETTIO F., FREXIA F., GIACHETTI A., GOBBETTI E., PINTORE G., ZANETTI G., BALOGH T., FORGACS T., AGOCS T., BOUVIER E.: A holographic collaborative medical visualization system. In *Medicine Meets Virtual Reality 2006* (January 2006), Westwood J. D., (Ed.), IOS, Amsterdam, The Netherlands.

[BFH*00] BIRKFELLNER W., FIGL M., HUBER K., WATZINGER F., WANSCHITZ F., HANEL R., WAGNER A., RAFOLT D., EWERS R., BERGMANN H.: The varioscope ar - a head-monted operating microscope for augmented reality. In *MICCAI '00: Proceedings of the Third International Conference on Medical Image Computing and Computer-Assisted Intervention* (London, UK, 2000), Springer-Verlag, pp. 869–877.

[BG07] BRUCKNER S., GRÖLLER M. E.: Style transfer functions for illustrative volume rendering. *Computer Graphics Forum (accepted for publication) 26*, 3 (Sept. 2007), to be presented at Eurographics 2007.

[Blu67] BLUM H.: A transformation for extracting new descriptors of shape. *Models for the Perception of Speech and Visual Form* (1967).

[BM92] BESL P., MCKAY N.: A method for registration of 3-d shapes. *IEEE Trans. PAMI 14*, 2 (1992), 239–256.

[BN95] BRO-NIELSEN M.: Modelling elasticity in solids using active cubes - application to simulated operations. In *Computer Vision, Virtual Reality and Robotics in Medecine* (Apr. 1995), vol. 905 of *Lecture Notes in Computer Science*, Springer, pp. 535–541.

[BN96] BRO-NIELSEN M.: Medical image registration and surgery simulation. *Ph.D. Thesis, Deptartment of Mathematical Modelling, Technical University of Denmark* (1996).

[BNC96] BRO-NIELSEN M., COTIN S.: Real-time volumetric deformable models for surgery simulation using finite elements and condensation. *Computer Graphics Forum (Proc. of Eurographics Š96)* (1996), C57–C66.

[Bro92] BROWN L.: A survey of image registration techniques. *ACM Computing Surveys 24*, 4 (Dec. 1992), 325–376.

[BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. *Proc. of SIGGRAPHŠ98, Computer Graphics 32* (1998), 106–117.

[CDA96] COTIN S., DELINGETTE H., AYACHE N.: Real Time Volumetric Deformable Models for Surgery Simulation. In *Vizualisation in Biomedical Computing* (1996), Hohne K., Kikinis R., (Eds.), vol. 1131 of *Lecture Notes in Computer Science*, Springer, pp. 535–540.

[CDA99a] COTIN S., DELINGETTE H., AYACHE N.: A hybrid elastic model allowing real-time cutting, deformations and force-feddback for surgery training and simulation. *Proc. of Computer Animation Š99* (1999).

[CDA99b] COTIN S., DELINGETTE H., AYACHE N.: Real-time elastic deformations of soft tissues for surgery simulation. *IEEE Transactions on Visualization and Computer Graphics 5*, 1 (1999), 62–73.

[CEO93] COVER S. A., EZQUERRA N. F., O'BRIEN J. F.: Interactively Deformable Models for Surgery Simulation. *IEEE Computer Graphics and Applications* (1993), 68–75.

[CGD97] CANI-GASCUEL M.-P., DESBRUN M.: Animation of deformable models using implicit surfaces. *IEEE Transactions on Visualization and Computer Graphics 3*, 1 (Mar. 1997).

[CGG*04] CIGNONI P., GANOVELLI F., GOBBETTI E., MARTON F., PONCHIO F., SCOPIGNO R.: Adaptive TetraPuzzles – efficient out-of-core construction and visualization of gigantic polygonal models. *ACM Transactions on Graphics 23*, 3 (August 2004), 796–803. Proc. SIGGRAPH 2004.

[Chr94] CHRISTENSEN G.: Deformable shape models for anatomy. *Washington University Ph.D. Thesis* (1994).

[CM96] CHINZEI K., MILLER K.: Compression of swine brain tissue ; experiment in vitro. *Journal of Mechanical Engineering Laboratory* (July 1996), 106–115.

[CMD*95] COLLIGNON A., MAES F., DELAERE D., VANDERMEULEN D., SEUTENS P., G.MARCHAL: Automated multimodality image registration using information theory. *in Information Processing in Medical Imaging* (1995), 263–274.

[Coh91] COHEN L.: On active contour models and balloons. *Computer Vision, Graphics, and Image Processing: Image Understanding 53*, 2 (1991), 211–218.

[Coq90] COQUILLART S.: Extended free-form deformation: A sculpturing tool for 3D geometric modeling. In *Computer Graphics (SIGGRAPH '90 Proceedings)* (Aug. 1990), Baskett F., (Ed.), vol. 24, pp. 187–196.

[Cot97] COTIN S.: *Modèles déformables temps-réel pour la simulation de chirurgie avec retour d'effort*. PhD thesis, Université de Nice Sophia-Antipolis, Nice, France, Nov. 1997. in French.

[Cow01] COWIN S.: Bone mechanics handbook, second ed. *Cambridge University Press* (2001).

[CRZP04] CHEN W., REN L., ZWICKER M., PFISTER H.: Hardware-accelerated adaptive ewa volume splatting. In *VIS '04: Proceedings of the conference on Visualization '04* (Washington, DC, USA, 2004), IEEE Computer Society, pp. 67–74.

[CT01] COOTES T., TAYLOR C.: Statistical models of ap-

pearance for medical image analysis and computer vision, in medical imaging. *Proc. of SPIE 4322* (2001), 238–248.

[Cut93] CUTTS A.: Muscle physiology and electromyography. *Mechanics of Human Joints: Physiology, Pathophysiology, and Treatment, ed. by V. Wright, E.L. Radin* (1993).

[CYTT92] CARIGNAN M., YANG Y., THALMANN N. M., THALMANN D.: Dressing animated synthetic actors with complex deformable clothes. In *Computer Graphics (SIGGRAPH '92 Proceedings)* (July 1992), Catmull E. E., (Ed.), vol. 26, pp. 99–104.

[CZ92] CHEN D., ZELTZER D.: Pump it up: Computer animation of a biomechanically based model of muscle using the finite element method. *Computer Graphics (Proc. of SIGGRAPH Š92)* (1992), 89–98.

[DDCB01] DEBUNNE G., DESBRUN M., CANI M.-P., BARR A.: Dynamic real-time deformations using space and time adaptive sampling. *Computer Graphics (Proc. of SIGGRAPH Š01)* (2001).

[DeL97] DELUCA C.: the use of surface electromyography in biomechanics. *J. App. Biomech. 13* (1997), 135–163.

[Del99] DELINGETTE H.: General object reconstruction based on simplex meshes. *International Journal of Computer Vision 32*, 2 (1999), 111–146.

[DHW*04] DENIS K., HUYSMANS T., WILDE T. D., FORAUSBERGER C., RAPP W., HAEX B., SLOTEN J. V., AUDEKERCKE R. V., DER PERRE G. V., HEITMANN K., DIERS H.: A 4d-optical measuring system for the dynamic acquisition of anatomical structures. *Proc. of MICCAI'04 2* (2004), 259–266.

[DKA95] DERICHE R., KORNPROBST P., AUBERT G.: Optical flow estimation while preserving its discontinuities: A variational approach. *Proc. of the 2nd Asian Conference on Computer Vision* (1995), 71–80.

[DKT95] DEUSSEN O., KOBBELT L., TUCKE P.: Using simulated annealing to obtain a good approximation of deformable bodies. In *Proc. Eurographics Workshop on Animation and Simulation* (Maastricht (NL), Sept. 1995), springer.

[DL95] DELP S., LOAN J.: A graphics-based software system to develop and analyze models of musculoskeletal structures. *Comput. Biol. Med. 25*, 1 (1995), 21–34.

[DP94] DRACE J., PELC N.: Measurement of skeletal muscle motion in vivo with phase-contrast mr imaging. *J. of Magnetic Resonance Imaging 4* (1994), 157–163.

[DP02] DEROY K., PEERAER L.: Visualising the pressure distribution on the plantar surface of the foot during stance, walking and running using a plantar pressure measurement system. *the encyclopedia of imaging science and technology, chapter 8* (2002).

[DSB99] DESBRUN M., SHRÖDER P., BARR A.: Interactive animation of structured deformable objects. *Proc. of Graphics Interface* (1999), 1–8.

[DSCP94] DELINGETTE H., SUBSOL G., COTIN S., PIGNON J.: A Craniofacial Surgery Simulation Testbed. In *Visualization in Biomedical Computing (VBC'94)* (octobre 1994).

[DSMK01] DAVATZIKOS C., SHEN D., MOHAMED A., KYRIACOU S.: A framework for predictive modeling of anatomical deformations. *IEEE Trans. Med. Imaging 20*, 8 (2001), 836–843.

[EEH00] EBERHARDT B., ETZMUSS O., HAUTH M.: Implicit-explicit schemes for fast animation with particles systems. *Eurographics workshop on Computer Animation and Simulation* (2000).

[EHK*04] ENGEL K., HADWIGER M., KNISS J. M., LEFOHN A. E., SALAMA C. R., WEISKOPF D.: Real-time volume graphics. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Course Notes* (New York, NY, USA, 2004), ACM Press, p. 29.

[EKE01] ENGEL K., KRAUS M., ERTL T.: High-quality pre-integrated volume rendering using hardwareaccelerated pixel shading, 2001.

[Fav05] FAVALORA G. E.: Volumetric 3d displays and application infrastructure. *Computer 38*, 8 (2005), 37–44.

[FDC01] FENSTER A., DOWNEY D., CARDINAL H.: Three-dimensional ultrasound imaging. *Phys. Med. Biol. 46* (2001), 67–99.

[FHM00] FITZPATRICK J., HILL D., MAURER C.: Image registration, chapter 8. *Handbook of Medical Imaging - Volume 2, Medical Image Processing and Analysis* (2000), 447–513.

[FLBW06] FRAUENFELDER T., LOTFEY M., BOEHM T., WILDERMUTH S.: Computational fluid dynamics: Hemodynamic changes in abdominal aortic aneurysm after stent-graft implantation. *CardioVascular and Interventional Radiology 29* (2006), 613–623.

[FNFB04] FISCHER J., NEFF M., FREUDENSTEIN D., BARTZ D.: Medical Augmented Reality based on Commercial Image Guided Surgery. In *Eurographics Symposium on Virtual Environments (EGVE 2004)* (2004), Coquillart S., Göbel M., (Eds.).

[Fun93] FUNG Y. C.: *Biomechanics - Mechanical Properties of Living Tissues*, second ed. Springer-Verlag, 1993.

[FWG*99] FERRANT M., WARFIELD S., GUTTMANN C., JOLESZ F., KIKINIS R.: 3d image matching using a finite element based elastic deformation model. *Medical Image Computing and Computer Assisted Intervention (MICCAI) 1679* (1999), 202–209.

[GM05] GOBBETTI E., MARTON F.: Far Voxels – a multiresolution framework for interactive rendering of huge

complex 3d models on commodity graphics platforms. *ACM Transactions on Graphics 24*, 3 (August 2005), 878–885. Proc. SIGGRAPH 2005.

[GMMT06] GILLES B., MOCCOZET L., MAGNENAT-THALMANN N.: Anatomical modelling of the musculoskeletal system from mri. *Proc. of MICCAI'06 4190* (2006), 289–296.

[GMTT89] GOURRET J., MAGNENAT-THALMANN N., THALMANN D.: Simulation of object and human skin deformations in a grasping task. *Proc. of SIGGRAPH'89, Computer Graphics 23* (1989), 21–30.

[GPMTV04] GILLES B., PERRIN R., MAGNENAT-THALMANN N., VALLÉE J.-P.: Bones motion analysis from dynamic mri: acquisition and tracking. *Proc. of MICCAI'04 2* (2004), 942–949.

[GPZT98] GOBBETTI E., PILI P., ZORCOLO A., TUVERI M.: Interactive virtual angioscopy. In *Proceedings IEEE Visualization* (Conference held in Research Triangle Park, NC, USA, October 1998), IEEE Computer Society Press, pp. 435–438.

[Gra00] GRAY H.: Anatomy of the human body. *20th edition. Philadelphia: Lea and Febiger 2* (2000).

[GSG*98] G.ABDULAEV, S.CADEDDU, G.DELUSSU, M.DONIZELLI, GIACHETTI L. A., E.GOBBETTI, LEONE A., MANZI C., P.PILI, A.SCHEININE, TUVERI M., A.VARONE, A.VENEZIANI, AND.ZORCOLO G.: Viva: The virtual vascular project. *IEEE trans. on Information Technology in medicine 14*, 1 (1998), 34–48.

[GSM*97] GIBSON S., SAMOSKY J., MOR A., FYOCK C., GRIMSON E., KANADE T., KIKINIS R., LAUER H., MCKENZIE N.: Simulating arthroscopic knee surgery using volumetric object representations, real-time volume rendering and haptic feedback . In *Proceedings of the First Joint Conference CVRMed-MRCAS'97* (Mar. 1997), Troccaz J., Grimson E., Mosges R., (Eds.), vol. 1205 of *Lecture Notes in Computer Science*, pp. 369–378.

[GSZW98] GHANEI A., SOLTANIAN-ZADEH H., WINDHAM J.: A 3d deformable surface model for segmentation of objects from volumetric data in medical images. *Computers in Biology and Medicine 28* (1998), 239–253.

[GZ06] GIACHETTI A., ZANETTI G.: Vascular modeling from volumetric diagnostic data: A review. *Current Medical Imaging Reviews 2*, 4 (November 2006), 415–423.

[HE01] HAUTH M., ETZMUSS O.: A high performance solver for the animation of deformable objects using advanced numerical methods. *Proc. of Eurographics* (2001).

[HFDKR00] HERMENS H., FRERIKS B., DISSELHORST-KLUG C., RAU G.: Development of recommendations for semg sensors and sensor placement procedures. *J Electromyography and Kinesiology 10* (2000), 361–374.

[HFS*01] HIROTA G., FISHER S., STATE A., LEE C., FUCHS H.: An implicit finite element method for elastic solids in contact. *Proc. of Computer Animation* (2001).

[HHD*00] HOLDEN M., HILL D., DENTON E., JAROSZ J., COX T., ROHLFING T., GOODEY J., HAWKES D.: Voxel similarity measures for 3-d serial mr brain image registration. *IEEE Transactions on Medical Imaging 19*, 2 (2000), 94–102.

[Hil44] HILBERT D.: *Geometry and the Imagination*. Dover, 1944.

[HLS*03] HEUSCHMID M., LUZ O., SCHAEFER J. F., STUECKER D., VONTHEIN R., LUBOLDT W., CLAUSSEN C. D., SEEMANN M. D.: Comparison of volume-rendered and surface-rendered mr colonography. *Technology in Cancer Research and Treatment 2*, 1 (2003), 13–17.

[HP78] HOROWITZ S., PAVLIDIS T.: A graph-theoretic approach to picture processing. *CGIP 7*, 2 (1978), 282–291.

[HPP*03] HÖHNE K., PFLESSER B., POMMERT A., PRIESMEYER K., RIEMER M., SCHIEMANN T., SCHUBERT R., TIEDE U., FREDERKING H., GEHRMANN S., NOSTER S., SCHUMACHER U.: Voxel-man 3d navigator: Inner organs. regional, systemic and radiological anatomy. *Springer-Verlag Electronic Media, Heidelberg, ISBN 3-540-40069-9* (2003).

[HS81] HORN B., SCHUNCK B.: Determining optical flow. *Artificial Intelligence 17* (1981), 185–203.

[HS88] HUNTER P., SMAIL B.: The analysis of cardiac function : A continuum approach. *Journal of Prog. Biophysic. molec. Biology 52* (1988), 101–164.

[JL95] JEAN LOUCHET XAVIER PROVOT D. C.: Evolutionary identification of cloth animation model. In *Workshop on Computer Animation and Simulation (Eurographics'95)* (1995), pp. 44–54.

[Joh07] JOHN N. W.: The impact of web3d technologies on medical education and training. *Comput. Educ. 49*, 1 (2007), 19–31.

[JP99] JAMES D., PAI D.: Artdefo: Accurate real time deformable objects. *Proc. of SIGGRAPH Š99, Computer Graphics* (1999).

[Kai96] KAISS, M. AND LE TALLEC, P.: La Modélisation numérique du contact œil-trépan. *Revue Européenne des Eléments Finis 5*, 3 (Feb. 1996), 375–408.

[KdH04] KAPTEIN B., DER HELM F. V.: Estimating muscle attachment contours by transforming geometrical bone models. *Journal of Biomechanics 37* (2004), 263–273.

[KG96] KOCH R., GROSS M.: Simulating facial surgery using finite element models. *Proc. Siggraph '96, Computer Graphics 30* (1996), 421–428.

[KGEB75] KENEDI R., GIBSON T., EVANS J., BARBENEL J.: Tissue mechanics. *Physics in Medecine and Biology 20*, 5 (1975), 163–169.

[KGG96] KEEVE E., GIROD S., GIROD B.: Craniofacial surgery simulation. *Proc. of VBC'96* (1996), 541–546.

[KKH02] KNISS J., KINDLMANN G., HANSEN C.: Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics 8*, 3 (2002), 270–285.

[KKKN96] KUHN C., KÜHNAPFEL U., KRUMM H.-G., NEISIUS B.: A 'virtual reality' based training system for minimally invasive surgery. In *Proc. Computer Assisted Radiology (CAR '96)* (Paris, June 1996), pp. 764–769.

[KLKB05] KAUFMAN A. E., LAKARE S., KREEGER K., BITTER I.: Virtual colonoscopy. *Commun. ACM 48*, 2 (2005), 37–41.

[KMTM*98] KALRA P., MAGNENAT-THALMANN N., MOCCOZET L., SANNIER G., AUBEL A., THALMANN D.: Real-time animation of realistic virtual humans. *Computer Graphics and Applications 18*, 5 (1998), 42–56.

[KQFL03] KUMAR B. V. R., QUATERONI A., FORMAGGIA L., LAMPONI D.: On parallel computation of blood flow in human arterial network based on 1-d modelling. *Computing 71*, 4 (2003), 321–351.

[KRG*02] KOCH M., ROTH S., GROSS M., ZIMMERMANN A., SAILER H.: A framework for facial surgery simulation. *Proc. of ACM SCCG* (2002).

[KW03] KRUGER J., WESTERMANN R.: Acceleration techniques for gpu-based volume rendering. In *VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03)* (Washington, DC, USA, 2003), IEEE Computer Society, p. 38.

[KWT88] KASS M., WITKIN A., TERZOPOULOS D.: Snakes: Active contour models. *International Journal of Computer Vision 1* (1988), 321–331.

[Lar86] LARRABEE W.: A finite element model of skin deformation: Biomechanics of skin and tissue : a review. *Laryngoscope 96* (1986), 399–405.

[LC87] LORENSEN W., CLINE H.: Marching cubes: A high resolution 3d surface construction algorithm. *Proc. of SIGGRAPH'87, Computer Graphics 21*, 39 (1987), 163–169.

[Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Comput. Graph. Appl. 8*, 3 (1988), 29–37.

[LH91] LAUR D., HANRAHAN P.: Hierarchical splatting: a progressive refinement algorithm for volume rendering. In *SIGGRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1991), ACM Press, pp. 285–288.

[LJF*91] LUCIANI A., JIMENEZ S., FLORENS J., CADOZ C., RAOULT O.: Computational physics: a modeler simulator for animated physical objects. In *Eurographics Workshop on Animation and Simulation* (Vienna, 1991), pp. 425–437.

[LM01] LÖTJÖNEN J., MÄKELÄ T.: Segmentation of mr images using deformable models: Application to cardiac images. *International Journal of Bioelectromagnetism 3*, 2 (2001), 37–45.

[LMT*03] LU A., MORRIS C. J., TAYLOR J., EBERT D. S., HANSEN C., RHEINGANS P., HARTNER M.: Illustrative interactive stipple rendering. *IEEE Transactions on Visualization and Computer Graphics 9*, 2 (2003), 127–138.

[LO99] LU T., O'CONNOR J.: Bone position estimation from skin marker co-ordinates using global optimisation with joint constraints. *J. of Biomechanics 32* (1999), 129–134.

[LRMK99] LÖTJÖNEN J., REISSMAN P., MAGNIN I., KATILA T.: Model extraction from magnetic resonance volume data using the deformable pyramid. *Medical Image Analysis 3*, 4 (1999), 387–406.

[Mar96] MARCUS B.: Hands on : Haptic feedback in surgical simulation. In *Proc. of Medecine Meets Virtual Reality IV (MMVR IV)* (San Diego, CA, Jan. 1996), pp. 134–139.

[MC94] MIRTICH B., CANNY J.: Impulse-based dynamic simulation. *The Algorithmic Foundations of Robotics* (1994).

[MC97] MESEURE P., CHAILLOU C.: Deformable Body Simulation with Adaptative Subdivision and Cuttings. In *Proceedings of the WSCG'97* (Pizen (Czech Republic), Feb. 1997), pp. 361–370.

[McM87] MCMAHON T.: Muscle mechanics. *Handbook of Bioengineering 8*, 2 (1987), 187–203.

[MD05] MONTAGNAT J., DELINGETTE H.: 4d deformable models with temporal constraints: application to 4d cardiac image segmentation. *Medical Image Analysis 9*, 1 (2005), 87–100.

[ME04] MORA B., EBERT D. S.: Instant volumetric understanding with order-independent volume rendering. In *Eurographics 2004* (2004), vol. 23.

[ME05] MORA B., EBERT D. S.: Low-complexity maximum intensity projection. *ACM Transactions on Graphics, 24*, 4 (October 2005), 1392–1416.

[Meg96] MEGLAN D.: Making surgical simulation real. In *Computer Graphics* (Nov. 1996), Addison Wesley, pp. 37–39.

[Men00] MENACHE A.: Understanding motion capture for computer animation and video games. *Academic Press, San Diego* (2000).

[MET*04] MORI S., ENDO M., TSUNOO T., KANDATSU S., TANADA S., ARADATE H., SAITO Y., MIYAZAKI H., SATOH K., MATSUSHITA S., KUSAKABE M.: Physical performance evaluation of a 256-slice ct-scanner for four-dimensional imaging. *Medical Physics 31*, 6 (2004), 1348–1356.

[MG04] MÜLLER M., GROSS M.: Interactive virtual materials. *Proc. of Graphics Interface'04* (2004), 239–246.

[MH06] MCINTOSH C., HAMARNEH G.: Genetic algorithm driven statistically deformed models for medical image segmentation. *Proc. of MedGEC and GECCO* (2006).

[MHHR06] MÜLLER M., HEIDELBERGER B., HENNIX M., RATCLIFF J.: Position based dynamics. *Proc. of Virtual Reality Interactions and Physical Simulations (VRIPhys)* (2006), 6–7.

[MHTG05a] MÜLLER M., HEIDELBERGER B., TESCHNER M., GROSS M.: Meshless deformations based on shape matching. *Proc. of SIGGRAPH'05* (2005), 471–478.

[MHTG05b] MULLER M., HEIDELBERGER B., TESCHNER M., GROSS M.: Meshless deformations based on shape matching. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers* (New York, NY, USA, 2005), ACM Press, pp. 471–478.

[MKLA80] MOW V., KUEI S., LAI W., ARMSTRONG C.: Biphasic creep and stress relaxation of articular cartilage: Theory and experiments. *ASME J Biomech Eng 102* (1980), 73–84.

[Mr05] MOSEGAARD J., RENSEN T. S.: Gpu accelerated surgical simulators for complex morphology. *Proc. of IEEE Virtual Reality '05* (2005), 147–153.

[MRF*96] MARK W., RANDOLPH S., FINCH M., VAN VERTH J., TAYLOR II R. M.: Adding force feedback to graphics systems: Issues and solutions. In *SIGGRAPH 96 Conference Proceedings* (Aug. 1996), Rushmeier H., (Ed.), Annual Conference Series, ACM SIGGRAPH, Addison Wesley, pp. 447–452.

[MRL*96] MUTHUPILLAI R., ROSSMAN P., LOMAS D., GREENLEAF J., RIEDERER S., EHMAN R.: Magnetic resonance imaging of transverse acoustic strain waves. *Magn. Reson. Med 36* (1996), 266–274.

[MT95] MCINERNEY T., TERZOPOULOS D.: A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4d image analysis. *Comp. Med. Imag. Graph. 19*, 1 (1995), 69–83.

[MT96] MCINERNEY T., TERZOPOULOS D.: Deformable models in medical image analysis: a survey. *Medical Image Analysis 1*, 2 (1996), 91–108.

[MT99] MCINERNEY T., TERZOPOULOS D.: Topology adaptive deformable surfaces for medical image volume segmentation. *IEEE Transactions on Medical Imaging 18*, 10 (1999), 840–850.

[MV98a] MAINTZ J., VIERGEVER M.: A survey of medical image registration. *Medical Image Analysis 2*, 1 (April 1998), 1–36.

[MV98b] MAINTZ J., VIERGEVER M.: A survey of medical registration. *Medical image analysis 2*, 1 (1998), 1–36.

[MZ02] MORI S., ZIJL P. V.: Fiber tracking: principles and strategies, a technical review. *NMR Biomed 15* (2002), 468–480.

[NLM99] NLM T. V. H. P.: National library of medicine, http://www.nlm.nih.gov/research/visible/visible_human.html.

[NM65] NELDER J., MEAD R.: A simplex method for function minimization. *Computer Journal 7* (1965), 308–313.

[NMK*03] NAVA A., MAZZA E., KLEINERMANN F., AVIS N. J., MC-CLURE J.: Determination of the mechanical properties of soft human tissues through aspiration experiments. *Proc. of MICCAI'03 1* (2003), 222–229.

[NMK*05] NEALEN A., MÜLLER M., KEISER R., BOXERMAN E., CARLSON M.: Physically based deformable models in computer graphics. *In Eurographics State of-the-Art Report (EG-STAR) 3* (2005), 71–94.

[NTH00] NG-THOW-HING V.: Anatomically-based models for physical and geometric reconstruction of humans and other animals. *Ph.D. Thesis, Department of Computer Science, University of Toronto* (2000).

[OH99] OŠBRIEN J., HODGINS J.: Graphical models and animation of brittle fracture. *Proc. of SiggraphŠ99, Computer Graphics* (1999), 137–146.

[OS88] OSHER S., SETHIAN J.: Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulation. *Journal of Computational Physics 79* (1988), 12–49.

[PBHS05] PAPAZOGLOU S., BRAUN J., HAMHABER U., SACK I.: Two-dimensional waveform analysis in mr elastography of skeletal muscles. *Phys. Med. Biol. 50* (2005), 1313–1325.

[PCL*07] PARK S. H., CHOI E. K., LEE S. S., BYEON J.-S., JO J.-Y., KIM Y. H., LEE K. H., HA H. K., HAN J. K.: Polyp measurement reliability, accuracy, and discrepancy: Optical colonoscopy versus ct colonography with pig colonic specimens. *Radiology online* (2007).

[PDA03] PICINBONO G., DELINGETTE H., AYACHE N.: Non-linear anisotropic elasticity for real-time surgery simulation. *Graph. Models 65*, 5 (2003), 305–321.

[PFJ*03] 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.: Deformable m-reps for 3d medical image segmentation. *International Journal of Computer Vision 55*, 2 (2003), 85–106.

[PHSE91] PELC N., HERFKENS R., SHIMAKAWA A., ENZMANN D.: Phase contrast cine magnetic resonance imaging. *Magnetic Resonance Quarterly 7* (1991), 229–254.

[PM90] PERONA P., MALIK J.: Scale-space and edge detection using anisotropic diffusion. *IEEE trans. on Pattern Analysis and Machine Intelligence 12*, 7 (1990), 629–639.

[PMTK01] PARK J., MCINERNEY T., TERZOPOULOS D., KIM M.: A non-self-intersecting adaptive deformable surface for complex boundary extraction from volumetric images. *Computers and Graphics 25*, 3 (2001), 421–440.

[Pro95] PROVOT X.: Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Graphics Interface '95* (May 1995), Davis W. A., Prusinkiewicz P., (Eds.), Canadian Information Processing Society, Canadian Human-Computer Communications Society, pp. 147–154. ISBN 0-9695338-4-5.

[PRZ92] PIEPER S., ROSEN J., ZELTZER D.: Interactive graphics for plastic surgery: A task-level analysis and implementation. *Computer Graphics (1992 Symposium on Interactive 3D Graphics) 25*, 2 (Mar. 1992), 127–134.

[PTSP02] PREIM B., TIETJEN C., SPINDLER W., PEITGEN H. O.: Integration of measurement tools in medical 3d visualizations. In *VIS '02: Proceedings of the conference on Visualization '02* (Washington, DC, USA, 2002), IEEE Computer Society, pp. 21–28.

[PTVF92] PRESS W., TEUKOLSKY S., VETTERLING W., FLANNERY B.: Numerical recipies in c (2nd ed.). *Cambridge University Press* (1992).

[QLH*02] QUICK H., LADD M., HOEVEL M., BOSK S., DEBATIN J., LAUB G., SCHROEDER T.: Real-time mri of joint movement with truefisp. *Journal of Magnetic Resonance Imaging 15*, 6 (2002), 710–715.

[RAD03] ROHDE G., ALDROUBI A., DAWANT B.: The adaptive bases algorithm for intensity-based nonrigid image registration. *IEEE Trans. Med. Imag. 22* (2003), 1470–1479.

[RE01] RHEINGANS P., EBERT D.: Volume illustration: Nonphotorealistic rendering of volume models. *IEEE Transactions on Visualization and Computer Graphics 7*, 3 (2001), 253–264.

[RFS99] ROHR K., FORNEFETT M., STIEHL H.: Approximating thin-plate splines for elastic registration: Integration of landmark errors and orientation attributes. *Lecture Notes in Computer Science 1613* (1999), 252–265.

[Rid07] RIDLEY E. L.: Radiologists may be better off taking hands-on approach to 3d workflow. AuntMinnie.com, 2007.

[RMA00] ROCHE A., MALANDAIN G., AYACHE N.: Unifying maximum likelihood approaches in medical image registration. *International Journal of Imaging Systems and Technology: Special Issue on 3D Imaging 11*, 1 (2000), 71–80.

[RRF96] R.J.STONE, R. F M.: Virtual environment training systems for laparoscopic surgery; activities at the uk's wolfson centre for minimally invasive therapy. *The journal of Medicine and Virtual Reality 1*, 2 (1996), 42–51.

[RSBB06] REITINGER B., SCHMALSTIEG D., BORNIK A., BEICHEL R.: Spatial analysis tools for virtual reality-based surgical planning. In *3DUI '06: Proceedings of the 3D User Interfaces (3DUI'06)* (Washington, DC, USA, 2006), IEEE Computer Society, pp. 37–44.

[RSEB*00] REZK-SALAMA C., ENGEL K., BAUER M., GREINER G., ERTL T.: Interactive volume on standard pc graphics hardware using multi-textures and multi-stage rasterization. In *HWWS '00: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware* (New York, NY, USA, 2000), ACM Press, pp. 109–118.

[RSH*99] RUECKERT D., SONODA L., HAYES C., HILL D., LEACH M., HAWKES D.: Nonrigid registration using free-form deformations: Application to breast mr images. *IEEE Trans. Med. Imag. 18* (1999), 712–721.

[RSS*96] ROHR K., STIEHL H., SPRENGEL R., BEIL W., BUZUG T. M., WEESE J., KUHN M.: Point-based elastic registration of medical image data using approximating thin-plate splines. *Visualization in Biomedical Computing* (1996), 297–306.

[SBM*94] SAGAR M. A., BULLIVANT D., MALLINSON G., HUNTER P., HUNTER I.: A Virtual Environment and Model of the Eye for Surgical Simulation. In *Computer Graphics (SIGGRAPH'94)* (1994), Annual Conference Series, pp. 205–212.

[SC04] SHELLOCK F., CRUES J.: Mr procedures: Biologic effects, safety, and patient care. *Radiology 232* (2004), 635–652.

[SD92] STAIB L., DUNCAN J.: Deformable fourier models for surface finding in 3d images. *Proc. of Visualization in Biomedical Computing (VBC)* (1992), 90–104.

[SES05] SVAKHINE N., EBERT D. S., STREDNEY D.: Illustration motifs for effective medical volume illustration. *IEEE Comput. Graph. Appl. 25*, 3 (2005), 31–39.

[SGT98] SINGH A., GOLDGOF D., TERZOPOULOS D.: Deformable models in medical image analysis. *IEEE Computer Society* (1998).

[She94] SHEWCHUK J.: An introduction to the conjugate gradient method without the agonizing pain. *Technical report, Carnegie Mellon University* (1994).

[SHN*06] SCHARSACH H., HADWIGER M., NEUBAUER A., WOLFSBERGER S., BUHLER K.: Perspective isosurface and direct volume rendering for virtual endoscopy applications. In *EUROVIS - Eurographics /IEEE VGTC Symposium on Visualization* (2006), pp. 315–322.

[SL96] SZELISKI R., LAVALLÉE S.: Matching 3-d anatomical surfaces with non-rigid deformations using octree-splines. *J. Comp. Vision 18* (1996), 171–186.

[SP86] SEDERBERG T., PARRY S.: Free form deformations of solid geometric models. *Proc. of SIGGRAPH'86, Computer Graphics 20*, 4 (1986), 150–161.

[SP91] SCLAROFF S., PENTLAND A.: Generalized implicit functions for computer graphics. In *Computer Graphics (SIGGRAPH '91 Proceedings)* (July 1991), Sederberg T. W., (Ed.), vol. 25, pp. 247–250.

[SR95] SONG G. J., REDDY N. P.: Tissue Cutting In Virtual Environment. In *Medecine Meets Virtual Reality IV* (1995), IOS Press, pp. 359–364.

[SS04] SCHWALD B., SEIBERT H.: Registration tasks for a hybrid tracking system for medical augmented reality, 2004.

[TBHF03] TERAN J., BLEMKER S., HING V. N. T., FEDKIW R.: Finite volume methods for the simulation of skeletal muscle. *Proc. of Eurographics/SIGGRAPH Symposium on Computer Animation* (2003).

[TBP03] TSAO J., BOESIGER P., PRUESSMANN K.: k-t blast and k-t sense: dynamic mri with high frame rate exploiting spatiotemporal correlations. *Magnetic Resonance in Medicine 50*, 5 (2003), 1031–1042.

[Ter86] TERZOPOULOS D.: Regularization of inverse visual problems involving discontinuities. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1986), 413–424.

[TF88a] TERZOPOULOS D., FLEISCHE K.: Deformable models. *The Visual Computer 4* (1988), 306–311.

[TF88b] TERZOPOULOS D., FLEISCHER K.: Deformable models. *The Visual Computer 4*, 6 (1988), 306–331.

[TF88c] TERZOPOULOS D., FLEISCHER K.: Modeling inelastic deformation: Viscoelasticity, plasticity, fracture. *Computer Graphics (SIGGRAPH'88) 22*, 4 (1988), 269–278.

[Thi95] THIRION J.: Fast non-rigid matching of 3d medical images. *INRIA Technical report 2547* (1995).

[TKH*05] TESCHNER M., KIMMERLE S., HEIDELBERGER B., ZACHMANN G., RAGHUPATHI L., FUHRMANN A., CANI M.-P., FAURE F., MAGNETAT-THALMANN N., STRASSER W., VOLINO P.: Collision detection for deformable objects. *Computer Graphics Forum 24*, 1 (2005), 61–81.

[TKZ*04] TESCHNER M., KIMMERLE S., ZACHMANN G., HEIDELBERGER B., RAGHUPATHI L., FUHRMANN A., CANI M.-P., FAURE F., MAGNENAT-THALMANN N., STRASSER W., VOLINO P.: Collision detection for deformable objects. *Proc. of Eurographics State-of-the-Art Report* (2004), 119–135.

[TM91] TERZOPOULOS D., METAXAS D.: Dynamic 3d models with local and global deformations: Deformable superquadrics. *IEEE Transactions on Pattern Analysis and Machine Intelligence archive 13*, 7 (1991), 703–714.

[TPBF87] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. *Computer Graphics 21*, 4 (1987), 205–214.

[TSB*05] TERAN J., SIFAKIS E., BLEMKER S., HING V. N. T., LAU C., FEDKIW R.: Creating and simulating skeletal muscle from the visible human data set. *IEEE TVCG 11* (2005), 317–328.

[TW88] TERZOPOULOS D., WITKIN A.: Physically based models with rigid and deformable components. *IEEE Computer Graphics and Applications 8*, 6 (1988), 41–51.

[TWK87a] TERZOPOULOS D., WITKIN A., KASS M.: Symmetry-seeking models and 3d object reconstruction. *International Journal of Computer Vision 1* (1987), 211–221.

[TWK87b] TERZOPOULOS D., WITKIN A., KASS M.: Symmetry-seeking models for 3d object reconstruction. *International Journal of Computer Vision 1(3)* (1987), 211–221.

[VGW06] VERESS A., GULLBERG G., WEISS J.: Measurement of strain in the left ventricle during diastole with cine-mri and deformable image registration. *ASME Journal of Biomechanical Engineering* (2006).

[Vio95] VIOLA P.: Alignment by maximization of mutual information. *Ph.D. thesis, Massachusetts Inst. Technol.* (1995).

[VM04] VALTORTA D., MAZZA E.: Dynamic measurements of soft tissue viscoelastic properties with a torsional resonator device. *Proc. of MICCAI'04 2* (2004), 284–292.

[VMT00] VOLINO P., MAGNENAT-THALMANN N.: Implementing fast cloth simulation with collision response. *Proc. of the International Conference on Computer Graphics* (2000), 257–268.

[VMT05] VOLINO P., MAGNENAT-THALMANN N.: Implicit midpoint integration and adaptive damping for efficient cloth simulation: Collision detection and deformable objects. *Computer Animation and Virtual Worlds 16* (2005), 163–175.

[VWKG01] VILANOVA BARTROLÍ A., WEGENKITTL R., KÖNIG A., GRÖLLER E.: Virtual colon unfolding, 2001.

[Wat92] WATERS K.: A physical model of facial tissue and muscle articulation derived from computer tomography data. In *Visualization in Biomedical Computing (VBC'92)* (Chappel Hill, NC, 1992), vol. 574.

[WC03] WILSSENS J., COCK A. D.: Barefoot, insole and ground pressure measurement comparison during 5m/s running. *Proc. of the 6th Symposium on footwear biomechanics* (2003).

[WCM92] WOODS R., CHERRY S., MAZZIOTTA J.: Rapid automated algorithm for aligning and reslicing pet images. *J. Comp. Assisted Tomogr. 16*, 4 (1992), 620–633.

[Wes90] WESTOVER L.: Footprint evaluation for volume rendering. In *SIGGRAPH '90: Proceedings of the 17th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1990), ACM Press, pp. 367–376.

[WG01] WEISS J., GARDINER J.: Computational modeling of ligament mechanics. *Critical Reviews in Biomedical Engineering 29*, 4 (2001), 1–70.

[WGW90] WITKIN A., GLEICHER M., WELCH W.: Interactive dynamics. *Computer Graphics 24*, 2 (1990), 11–22.

[WHQ05] WANG K., HE Y., QIN H.: Incorporating rigid structures in non-rigid registration using triangular b-splines. *VLSM* (2005), 235–246.

[WKB99] WAN M., KAUFMAN A., BRYSON S.: High performance presence-accelerated ray casting. In *VIS '99: Proceedings of the conference on Visualization '99* (Los Alamitos, CA, USA, 1999), IEEE Computer Society Press, pp. 379–386.

[WS00] WANG Y., STAIB L.: Physical model-based non-rigid registration incorporating statistical shape information. *Medical Image Analysis 4* (2000), 7–20.

[XX06] XIAO H., XIAO D.-G.: An accelerating splatting algorithm based on multi-texture mapping for volume rendering. In *GRAPHITE '06: Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia* (New York, NY, USA, 2006), ACM Press, pp. 205–208.

[YSAT01] YOU B., SIY P., ANDERST W., TASHMAN S.: In vivo measurement of 3-d skeletal kinematics from sequences of biplane radiographs: Application to knee kinematics. *IEEE Transactions on medical imaging 20*, 6 (2001), 514–524.

[Zie77] ZIENKIEWICZ O.: The finite element method. *McGraw-Hill, London, 3rd edition* (1977).

[ZPR*88] ZERHOUNI E., PARISH D., ROGERS W., YANG A., SHAPIRO E.: Human heart: Tagging with mr imaging, a method for noninvasive assessment of myocardial motion. *Radiology 169* (1988), 59–63.

[ZT99] ZHOU Y., TOGA A. W.: Efficient skeletonization of volumetric objects. *IEEE Transactions on Visualization and Computer Graphics 5*, 3 (/1999), 196–209.