Towards a realistic echographic simulator with force feedback

D. d'Aulignac and C. Laugier

M. C. Cavusoglu

GRAVIR/INRIA Rhône Alpes
38330 Montbonnot Saint-Martin, France
{Diego.D_Aulignac, Christian.Laugier}@inrialpes.fr

Robotics and Intelligent Machines Laboratory
University of California, Berkeley
Berkeley, CA 94720

Abstract

This paper proposes a mass-spring model of a human thigh based on real data acquired. It addresses both the difficulties of determining the parameters of this model to fit the measurements and the computational demands. Implicit integration is used to update the model through time. The motivation behind this work is to provide accurate force-feedback for an echographic simulator that could be used to train practitioners to detect a thrombosis.

1 Introduction

Echography, in general terms, exploits the information which an acoustic signal provides when it is reflected off a structure to determine the position and shape of the latter. The frequency of these acoustic signals lie within the range of 1MHz and 12 MHz, thus higher then the sounds which are perceptible by humans. Therefore, they are also refered to as *ultrasounds*.

Even though they had already been discovered by P. Curie in 1880, they would have to wait until the 1970's to find an application in the field of medicine. Since, they have been widely used as an inexpensive and non-traumatic means of diagnosis. A common exam is the echography of the thigh to detect a thrombosis in the vein. A healthy vein will compress under the influence of an external force while a vein affected by thrombosis will only partially or even not at all compress, depending on the stage in the evolution of the illness. Depending on the pressure the practitioner applies with the echographic probe on the thigh, he will get an image from which the current state of the vein can be deduced, and hence a possible thrombosis diagnosed.

However, the learning process of this procedure is somehow long and only after approximately 1000 echographic exams an acceptable competence in acquired. The first 500 exams will have to be carried out under the supervision of an experienced practitioner. Virtual environments present an alternative to the conventional medical training scheme. It is possible to create an interactive 3D simulation environment, where the doctors can manipulate or cut dynamically and geometrically correct models of organs and tissues with an haptic interface. The idea is similar to using flight simulators to train pilots. Virtual environments give an environment where there is no risk to a patient, and therefore less stressful. They are interactive and three dimensional contrary to books. Virtual environments also give a unique advantage, as it is possible to generate arbitrary anatomies and pathologies, so that the doctors can be trained for cases that are not frequently encountered.

The goal of this work is to lay the groundwork for the development of an echographic simulator with force feedback. In the final system, the trainee will be looking at artificially generated echographic images, and interacting with a computer simulated dynamical thigh model through a haptic interface. Constructing realistic but computationally efficient models is the main challenge in developing a virtual reality training simulator. In this application it is necessary to have models for deformable tissue being manipulated by the doctor as well as models to construct artificial echographic images.

In this paper, a dynamic model of the human thigh based on experimentally determined deformation characteristics will be presented. The simulation of the model is performed using AlaDyn3d, a dynamical modeling and simulation package, which will also be briefly introduced, followed by a discussion of the results and future directions.

2 Previous Work

Henry [4] examined how a set of echographic images could be used to construct a more general model that would take into account the orientation and pressure exerted with the probe on the thigh. The approach is based on an interpolation method; from a real set of echographic images taken, it will build an echographic image for *any* position and orientation of the probe (see Figure 1 and 2)[1]. Further it is possible to deform the interpolated image obtained with respect to the pressure that is applied to the probe. The criteria on which this deformation depends includes important factors such as the arterial and venous pressure. By modifying these criteria we can simulate a set of pathologies on which medical students could be trained for the identification of the latter.

Figure 1: Creating an image from a set of real echographic images. The highlighted image was obtained by interpolation.
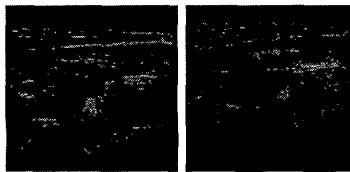

Figure 2: On the left we can see the real image, and on the right the image obtained by interpolation

However, for a meaningful echographic simulation it is of paramount importance to consider the forces involved in such a procedure since they will dictate the deformation of the tissue, and therefore, the subsequent image that is acquired. Laffont [7] was the first to study the implications of the dynamics of such a system which are essential in the development of a realistic simulator coupled to a haptic interface. He based his model on a system of inter-connected springs. Such a type of approach has already been successfully applied to areas such as facial animation

---

[1] Figures kindly provided by the TIMC-GMCAO project

[8] and have the decisive advantage of being easy-to-implement and faster than finite-element simulation, even though lately considerable speed-ups have been achieved in this domain through the use of pre-calculated matrices. This has resulted in, for example, the real-time simulation with force-feedback of the human liver [2] based on the *Visible Human* dataset. However, the model used is inconvenient in the sense that the assumption is made that the tissue is hyperelastic and linear in its deformation. This is clearly not the case. As biomechanical experiments confirm the stress-strain curve of human tissue is non-linear. [8] tackles this problem by using a three-layer model which takes into account the properties of the epidermis, the fatty sub-cutaneous tissue, the muscle, and the bone. The great difficulty for such a system lies in the identification of the parameters of the individual springs that will give the same results to an external force being applied as the measurements in the real world.

3 Modeling of the Thigh

We have constructed a dynamical model of the human thigh based on experimental measurements of its elasticity.

3.1 Experimental Setup and Data Acquisition

In order to model an object such that its behavior corresponds to reality, measurements must be taken on the real object. In this case we are interested in the deformation of the thigh with respect to an external force which is applied. Intuitively we can affirm that the deformation of the thigh is not the same depending on the shape of the object used to provoke this deformation, or more precisely, the contact surface of that object. Since our aim is to build a generic simulation which will allow a physically correct behavior which is almost independent of the object we choose to deform the thigh with, two different objects have been used to measure the behavior of the thigh in terms of penetration distance with respect to the external force being applied.

The first of these has a tip of pyramidal shape to provoke a punctual force response, while the second one has the same contact surface area as a typical echographic probe. These pseudo-probes are then mounted on a force sensor which in turn is mounted on a PUMA articulated arm. The probe is then positioned perpendicularly to the surface of the thigh at

each of 64 points where measurements will be taken. These points are regularly distributed over the area where the echography is performed when trying to detect a thrombosis in the vein. The robotic arm then advances $2mm$ using the reference of the end effector, i.e. the probe pushes along the axis which is normal to the surface of the thigh at the given point. The force is recorded and the procedure is repeated up to an upper force limit.

Figure 3: The two probes used for measuring the behavior of the thigh.

Figure 4 shows the non-linear relationship between the penetration distance and the reaction force at 11 different points along the thigh. The difference in the curves can be accounted for due to the fact that the thigh is not homogeneous (e.g. in some regions fatty tissue might be pre-dominant, while in others there may be very little separating the epidermis from the bone).
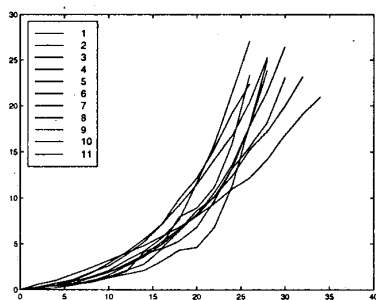

Figure 4: Plot of the reaction force in function of the penetration distance at 11 different points on the thigh using a probe with a punctual contact.

Figure 5 plots the values of the forces for the same

points as above but using the second probe with a larger contact surface. As one might expect this results in a larger force for the same penetration distance since the external force applied is distributed over a much larger area.
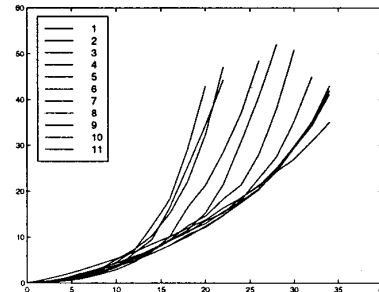

Figure 5: Plot of the reaction force in function of the penetration distance at 11 different points on the thigh using a probe where a surface is in contact.

The assumption is that from these two distinct sets of data it will be possible to make a model which will respond correctly to not only to the two probes used for measurement purposes, but also other probes of a different shape.

3.2 Model construction

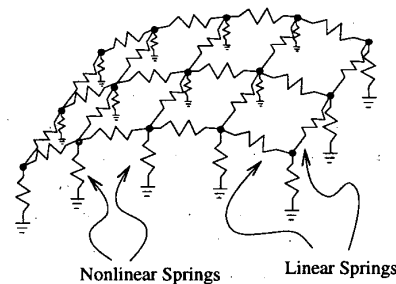

Nonlinear Springs   Linear Springs

Figure 6: Two layer model of the thigh.

Based on the experimental data and the computational requirements, a two layer lumped element model is chosen. The two layer model is composed of a surface mesh of masses, linear springs and dampers, and a set of nonlinear springs orthogonal to the surface to model volumetric effects by giving normal support to the surface mesh (See Figure 6). Two layer models were also used by several authors in the literature, for example in [8]. The deformation-force relation of the nonlinear springs are of the form

$$f(x) = \frac{x}{ax + b} \qquad (1)$$

The nonlinear spring response is chosen to model the incompressibility of the thigh after a certain deformation. In the model, the values of the surface elements are chosen uniform whereas the parameters of the nonlinear spring vary around the mesh to model heterogeneous nature of the thigh mentioned above, while keeping the number of unknown parameters small.

Parameter Estimation

Estimation of the model parameters from experimental measurements is a critical part of the modeling. We have used a two step optimization approach based on nonlinear least squares estimation. In the two step approach, the experimental data is first fit to a simple model without the surface elements, and the results of this fit is used as the initial conditions for the estimation of the parameters of the complete model. This approach is chosen to avoid problems with local minima.

One thing to note here is that the parameter estimation is based on a simplified interaction model between the tissue and indentor, making the assumption that the probe will principally act on one spring, since the contact area is small. In this simplified interaction model, the nodes that are not in contact with the probe are kept stationary. This simplifies the force calculation as only a small number of springs are deformed for any given dataset.

The mean absolute error between the measured values and the values estimated by the model is 1.05 Newtons, with standard deviation 1.84 Newtons, over the whole dataset. This is equivalent to an approximate error of 5 percent. The distribution of the error can be seen in Figure 7, showing how many of the values estimated by the model exibit a given error with respect to the real measurements.

4 Dynamic Simulation of the Thigh

The dynamic simulation of the deformation of the thigh under the pressure of the probes has been achieved using the AlaDyn3D system developed by the SHARP research team.

4.1 Outline of the AlaDyn3d System

The AlaDyn3D system allows the construction of complex physical scenes involving various interacting
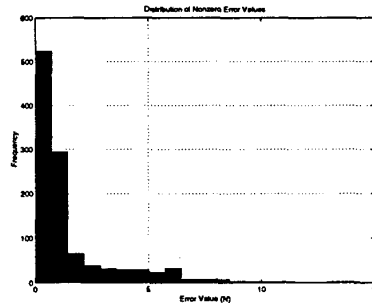

Figure 7: Frequency distribution of the error between the measured values and the values estimated by the model.

objects with different physical properties (deformable, rigid, friction, etc.) [5].

Object modeling The geometrical properties of an object are represented by a set of points and polygons which describe the surface of the object (N.B. the simulator also allows to discretise the interior of an object into tetrahedrons, but this, in our case is irrelevant). From this data the inertial matrix of rigid objects can be constructed through classical rigid body mechanics. For deformable objects the inertial properties of an object are implicitly defined by the distribution of the mass on its surface. These masses are interconnected by spring-damper connectors along the edges of the polygons, describing the constraint between pairs of masses. Typically these are visco-elastic *linear* constraints governed by the equation

$$F = -\lambda \Delta d - \mu \dot{d}$$

where d is the distance between the two particles and $\dot{d}$ is the relative speed of the two particles. However, this gives rise to a linear relationship between the force and the difference of the distance of the particles with respect to their resting length. This disagrees with the non-linear behavior of human tissue. This problem, and the proposed solution have been addressed in Section 3.2.

Processing physical interactions Further, to detect the interaction between objects in a scene, a collision detection algorithm must be used. It is well known that collision detection is a computationally expensive task and a potential bottleneck in every application that aims to achieve real-time performance, such as in this case. For this reason it is essential to optimize the routines to their maximal extent. Gilbert

et al. have developed an algorithm which calculates, in linear time $O(n)$, the positive distance between the two convex bodies, and gives an approximation of the negative distance when 'small' inter-penetrations occur. [6] extended this algorithm in order to obtain, in linear time, the points of contact on two deformables concave objects, as well as the direction of the contact and the volume of interpenetration.

Once a collision has been detected an appropriate response must be computed. There exist several collision response models but for deformable objects the penalty method seems to be the most appropriate (N.B. impulse based methods make the assumption that a collision happens in a infinitely small interval of time, this clearly not being the case when dealing with deformable objects; see [3] for more details). Using this approach the force applied at a given point of an object where a collision has taken place is given by

$$\vec{F_c} = \begin{cases} (-\lambda v - \mu \dot{v} v)\vec{k} & \text{if } v < 0 \\ 0 & \text{otherwise} \end{cases} \qquad (2)$$

where λ is the rigidity factor of the collision, μ is a damping factor (which represents the dissipation of energy), v the volume of inter-penetration, and $\vec{k}$ the contact direction.

4.2 Implicit Integration

The simplest integration method for motion update is the well-known explicit Newton-Euler. Using this method, particle i of an object has the following basic update formulae.

$$\begin{aligned} v_i^{t+\Delta t} &= v_i^t + \Delta t a_i^t \\ x_i^{t+\Delta t} &= x_i^t + \Delta t v_i^t \end{aligned} \qquad (3)$$

However, this integration method takes no notice of wildly changing derivatives. In our case, forces change very rapidly when the position of point on the surface of the thigh changes, and the only solution to stop the integration from diverging is to reduce the time step. This is very bad news for a real-time application since at every timestep collision detection algorithms must be rerun and the forces of all constraints reevaluated. This problem has been addressed by [1] in the context of cloth simulation. They propose an *implicit* integration method, also known as *backward* Euler method. In its general from it has the following formulation.

$$\Delta Y \left[\frac{1}{\Delta t} I - \lambda (\nabla f)|_{Y=Y_0} \right] = f(Y_0) \qquad (4)$$

where Y is the state vector expressing the positions and velocities of all particles. Thus ΔY represents the change in state and $f(Y_0)$ the derivative at the current state. Further, I is the identity matrix, ∇f the jacobian matrix derived from the spring configuration of the object, and λ is a constant where $\lambda \in [0, 1]$. The resulting square matrix is composed of n^2 elements, where n is the size of the state vector. However, since particles in most objects are nearly never fully interconnected, this matrix is in the overwhelming majority of cases sparse. Thus for large objects it is essential to store these matrices in a sparse format if we want to keep memory requirement within reasonable limits.

The second problem is that at each timestep the resulting matrix on the right-hand-side of the equation has to be inverted to find the change of state ΔY. Inverting a matrix is a computationally expensive task with a complexity $O(n^3)$. For small objects inverting the matrix is still feasible and has been applied with success. However, once we are dealing with a large number of particles in our system we have to find an alternative solution. Thus we use the conjugate-gradient method [9] to solve Equation 4. This works well in practice, achieving stability at much larger timesteps than an explicit method would allow. Of course the computational complexity *of each timestep* is much larger than an explicit method, but the ability to take vastly bigger steps outweighs this disadvantage in cases where springs are stiff.

5 Results

Figure 8 shows the model as it has been built in our simulation system. A force is being applied on the thigh using a probe which provokes a deformation which is in accordance with the measurements taken.
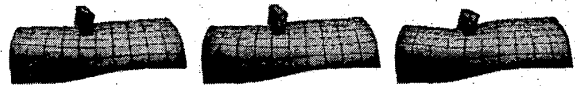

Figure 8: Model viewed on the simulation system's graphical interface.

The computational speed on a Silicon Graphics R10000 machine is in the order of 100 frames per second of animation, which lets us envisage our simulation system as part of a working echographic simulator which operates in real-time.

6 Discussion

We are at the stage where we can simulate accurately the deformation of the thigh under a static load for a variety of objects used to provoke this deformation. However, what is missing are measurements about the damping and friction coefficients. The first is important in the sense that it defines how fast the tissue will react to a change in external force being applied to it, and the latter due to the fact that the practitioner *slides* the probe down the inner side of the thigh to capture the echographic images. This friction coefficient is influenced by the gel that is typically applied on the thigh to avoid the presence of air between the probe and the tissue, which in turn might influence the quality of the echographic images that are acquired.

We believe that these coefficients may either be adjusted manually using the feedback provided by an experienced practitioner or that additional measurements may be necessary. Once these last problems resolved the primary goal will be to integrate the resulting model with the echographic image creator presented in Section 2 and a force feedback device of type PHANToM built by SensAble Technologies.

7 Conclusions

In this paper we have presented and constructed a model of a human thigh from measurements taken in the real world. We emphasize the point that the human thigh is not linear in its elasticity and introduce the use of non-linear springs for this application and propose a possible solution for the identification of the parameters of the springs. Further, from experimental results of the simulation using the proposed model we believe in the feasibility of our approach as a part of a real-time system which will allow realistic force feedback and the automatic generation of echographic images which takes into account the current deformation.

Related web sites For more information on the project visit our project home-page: http://www.inrialpes.fr/sharp/modelisation/index-en.html

References

[1] D. Baraff and A. Witkin. Large steps in cloth simulation. In *Computer Graphics (Proc. SIG-GRAPH)*, pages 43–54, 1998.

[2] M. Bro-Nielsen and S. Cotin. Real-time volumetric deformable models for surgery simulation using finite elements and condensation. In *Proceedings of Eurographics*, volume 15, pages 57–66, 1996.

[3] A. Deguet, A. Joukhadar, and C. Laugier. Models and algorithms for the collision of rigid and deformable bodies. In P. K. Agarwal, L. E. Kavraki, and M. T. Mason, editors, *Robotics: the algorithmic perspective*, pages 327–338. A K Peters, 1998. Proc. of the Workshop on the Algorithmic Foundations of Robotics. Houston, TX (US). March 1998.

[4] D. Henry. *Outils pour la modélisation de structure et la simulation d'examens échographiques*. PhD thesis, Université Joseph Fourier, Grenoble (FR), 1997. (in french).

[5] A. Joukhadar and C. Laugier. Dynamic simulation: Model, basic algorithms, and optimization. In *Proc. of the Workshop on the Algorithmic Foundations of Robotics*, Toulouse (FR), July 1996.

[6] A. Joukhadar, A. Wabbi, and C. Laugier. Fast contact localisation between deformable polyhedra in motion. In *Proc. of the IEEE Computer Animation Conf.*, pages 126–135, Geneva (CH), June 1996.

[7] P. Laffont. Simulation dynamique pour le diagnostic de thromboses veineuses. Mémoire de Diplôme d'Etudes Approfondies, Université de Savoie, Chambéry (FR), 1997.

[8] Y. Lee, D. Terzopoulos, and K. Waters. Realistic facial modeling for animation. In *Computer Graphics* Proceedings, Annual Conference Series, Proc. SIGGRAPH '95 (Los Angeles, CA), pages 55–62. ACM SIGGRAPH, August 1995.

[9] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C*. Cambridge Univ. Press, 2 edition, 1992.