1-1-2000

Automatic modeling of knee-joint motion for the virtual reality dynamic anatomy (VRDA) tool

Yohan Baillot
University of Central Florida

Jannick P. Rolland
University of Central Florida

Kuo-Chi Lin
University of Central Florida

Donna L. Wright

Yohan Baillot
Jannick P. Rolland
School of ECE and Computer Science, and
School of Optics/CREOL
University of Central Florida,
Orlando, FL 32816-2700

Kuo-Chi Lin
Institute for Simulation and Training
University of Central Florida,
Orlando, FL 32826-0544

Donna L. Wright
Division of Radiologic Science
Departments of Allied Health Science and Radiology
University of North Carolina, Chapel Hill, NC 27599-7130

Automatic Modeling of Knee-Joint Motion For The Virtual Reality Dynamic Anatomy (VRDA) Tool

Abstract

This paper presents a method and algorithms for automatic modeling of anatomical joint motion. The method relies on collision detection to achieve stable positions and orientations of the knee joint by evaluating the relative motion of the tibia with respect to the femur (for example, flexion-extension). The stable positions then become the basis for a look-up table employed in the animation of the joint. The strength of this method lies in its robustness to animate any normal anatomical joint. It is also expandable to other anatomical joints given a set of kinematic constraints for the joint type as well as a high-resolution, static, 3-D model of the joint. The demonstration could be patient specific if a person's real anatomical data could be obtained from a medical imaging modality such as computed tomography or magnetic resonance imaging. Otherwise, the demonstration requires the scaling of a generic joint based on patient characteristics. Compared with current teaching strategies, this Virtual Reality Dynamic Anatomy (VRDA) tool aims to greatly enhance students' understanding of 3-D human anatomy and joint motions. A preliminary demonstration of the optical superimposition of a generic knee joint on a leg model is shown.

1 Introduction

The Virtual Reality Dynamic Anatomy (VRDA) tool is a device that will allow the visual superimposition of internal anatomy on corresponding external anatomy. We are currently working on the visualization of anatomical joints. In this case, as the user manipulates the real joint of a participant as illustrated in Figure 1(a), a tracking sensor measures the attitude of the joint. A synthetic image of the inside of the joint is then shown stereoscopically to the user at the location of the real joint and in the same attitude. The internal and external components of the joint appear registered to create for the user the impression of 3-D "X-ray" vision (Wright, Rolland, & Kancherla, 1995). An illustration of a view that the user should have is shown in Figure 1(b). The internal components may correspond to the real joint if they are preacquired from the participant via medical imaging, or they may come from a generic joint scaled to approximate the size of the participant's internal anatomy. The VRDA tool application will be further detailed in section 2.3.

To assemble such a tool, we are developing several technologies:

The custom design of head-mounted displays (HMDs) (Rolland, 1998; Rolland et al., 1998; Rolland, 2000).

Presence, *Vol. 9, No. 3, June 2000,* 223–235

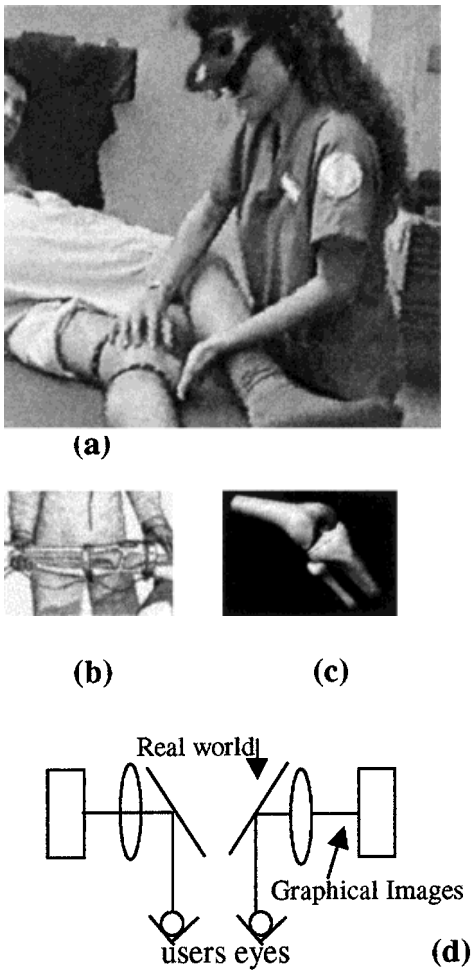

Figure 1. *(a) The VRDA tool (in development) allows superimposition of virtual anatomy on a model patient. (b) An illustration of the view of the HMD user (courtesy of Andrei State). (c) A rendered frame of the knee bone structures that will be integrated in the tool. (d) Schematic of the optical superimposition inside the HMD.*

Trackers and tracking methods that yield accuracy of at least 0.1 mm for rigid body motions (such as the head) (Davis, Peuchot, & Rolland, 1999).

A method for modeling anatomical joint motion so that, when combined with actual tracking data of joint motion, 0.1 mm accuracy is also obtained for joint motion capture.

A method for registering the graphical model of inner joint anatomy with the external joint anatomy in motion.

A set of methods to assess the tool.

Within this research framework, the main contribution of this paper is the presentation of a method for the automatic modeling of joint motion. The method was applied to the knee joint, and the results are presented. Moreover, some early results of the registration of a graphical anatomical joint on a static mannequin leg obtained via simple methods are also included to convey a vision for what the final superimposition may look like. Advanced methods for motion capture and registration that are currently being developed will be reported elsewhere.

In this paper, we shall first provide some background on modeling joint motions and visualizing such information using augmented reality. Section 3 provides the algorithm for *automatic* modeling of joint motions. Section 4 briefly describes the methods for visualization with augmented reality, and section 5 shows the results regarding both modeling and visualization.

2 Background

We shall first briefly review and illustrate the geometry and anatomy of the knee joint. Methods for representing joint motion are then reviewed, and augmented reality as a technique for visualization is introduced.

2.1 Knee Joint Anatomy

In regards to motion, the knee joint is one of the most complex anatomical joints of the human body. As such, it is of great clinical interest. We shall focus on the tibiofemoral joint (figure 2).

The tibiofemoral joint is composed of bones, muscles, tendons, menisci, ligaments, and the joint capsule. The bones are the rigid structures of the joint that are activated by the muscles attached to the bones via tendons. The menisci cushion helps stabilize and isolate the cartilage surfaces, while the ligaments limit excess motion.

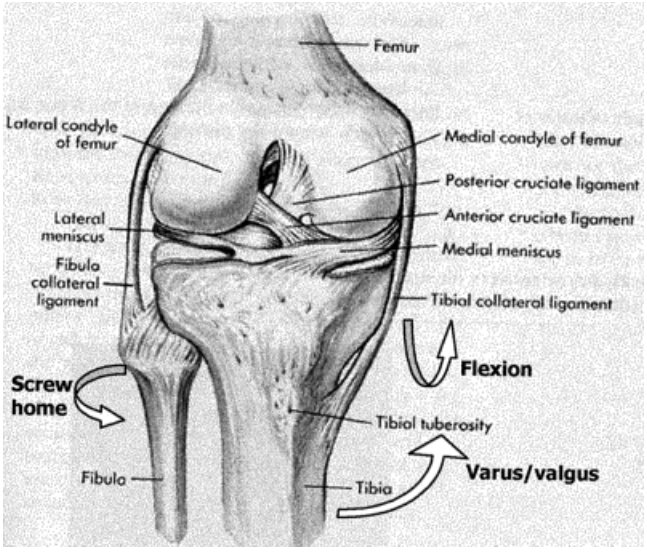

Figure 2. *Anatomy of the human knee joint (no patella).*

The bones we shall consider in the joint are the femur and the tibia. Among the other components, the patella is not considered at this time (but its motion will be relatively simple to add), and the fibula does not play a predominant role in the motion of the joint (and so will be considered fixed to the tibia). The femur has two sphere-like contact surfaces called the *condyles.* The medial condyle is closer to the center of the body and is larger than the lateral condyle. The tibia has two corresponding spherical surfaces that receive the femur condyles, the *tibial plateaus.*

The menisci are divided into the lateral and the medial meniscus. Although the meniscus is a sponge-like body, its deformation is not perceptible to the human eye (Frankel, 1997). During the motion of the joint essentially under load, the menisci slide on the tibial plateaus and wrap around the condyles to naturally fill the eventual gap between the bone surfaces (Bylski-Austrow et al., 1994).

The muscles create the motion of the joint along two degrees of freedom—flexion-extension and varus-valgus—as reported, for example, by Nordin and Frankel (1980). These degrees of freedom condition the relative orientation of the bones along two orthogonal axes. The function of the muscles is accounted for in the model proposed.

The ligaments are spring-like strings that connect the tibia and the femur at specific points and limit the motion of the bones as the joint moves (Blankevoort & Huiskes, 1988; Huiskes & Blankevoort, 1990). The contact surfaces also create constraints that limit and produce the relative positions of the bones. Furthermore, the ligaments and the contacting surfaces impose a limited range of motion for the varus-valgus. Zero degree varus-valgus corresponds to the two condyles touching the tibial plateaus.

Finally, a rotation of the tibia along its main axis known as the *screw-home effect* is produced as the knee is flexed as reported, for example, by Conlay and Long (1994). Specifically, the tibia performs an external rotation due to the longer circumference of the medial condyle. During extension of the joint, the opposite rotation occurs, locking the joint to ensure its stability. The screw-home effect is included as a natural consequence of the approach to modeling that is developed and described in this paper.

2.2 Approaches to Modeling Knee Joint Motion

As a first step to modeling and visualizing joint motion, we shall first represent bone motions while accounting for constraints that are imposed by deformable structures. Considering the relative motion of the bones, the knee joint is far from a simple cylindrical liaison as the knee flexes and extends. Instead, the center of rotation of the tibia, with respect to the femur, describes a 3-D space curve (Soudan, Van Audekercke, & Martens, 1979; Nordin & Frankel, 1980), and is thus referred as the instant center of rotation given that it is not a fixed point. The instant center of rotation results from the shape of the condyles being ellipsoid instead of spherical.

Sensing devices cannot precisely acquire the relative position of the bones down to a 0.1 mm specification due to the sensing devices sliding somewhat on the skin during motion. Therefore, a realistic model of joint motion is essential to complement the joint motion data that simply provide in this case the flexion-extension and varus-valgus angles. The precise relative position of the

anatomical components is provided by the model (Rolland, Wright, & Kancherla, 1997; Baillot, 1999).

Approaches to accurate representation of joint motion may consist of

1. Collecting dynamic 3-D medical data sets of the relative motions of individual anatomical components,
2. Collecting multiple static 3-D data sets of moving individual components that would then be interpolated to simulate a continuous dynamic data set, or
3. Modeling the motion algorithmically based on knowledge of anatomy and physical constraints.

Although current medical imaging technology does not yet allow for the first approach, the second and third approaches are discussed in the following sections.

2.2.1 Collecting Multiple Static 3-D Data Sets. The approach that consists of collecting multiple static 3-D data sets is currently pursued (Totterman et al., 1998) via magnetic resonance imaging as the medical imaging modality. The approach requires optimizing the data acquisition methods to differentiate various anatomical components. The imaging differentiation is necessary for accurate segmentation. Furthermore, the segmentation would need to be fully automated from a practical point of view, a challenge that has not yet been overcome. Finally, given a segmented data set that may serve as a geometrical 3-D model of inner anatomy for the VRDA tool, a scaling method of the geometrical model is necessary because it does not typically correspond to that of the participant who may be any student in the classroom. How scaling of the model will affect the final geometrical model and associated motions must be investigated and assessed. However, such an approach to representing joint motion is promising and may play an important role for teaching abnormal joint motions in addition to normal motions.

2.2.2 Modeling the Motion Algorithmically. This approach consists of taking a 3-D geometrical model of anatomy obtained (for example) from fine digitization of a cadaver's inner anatomy and then applying an algorithm to yield an animated model. We shall pursue and demonstrate this approach with a high-resolution geometrical model of a knee joint obtained from Viewpoint, Inc.

One of the key differences between the various algorithms developed so far arises from the required level of representation of the anatomical models and the required level of accuracy in the motion. For example, a level of representation of the bones may be as simple as sticks for bones, while the model of flexion-extension may be a simple hinge model. Those are naturally too crude for the VRDA tool application.

In reviewing algorithms for modeling joint motion, we shall distinguish between two main sets of methods: one that considers the exact motion of the joint produced by the dynamic interaction of the components, and another that employs some approximation of the knee joint motions to yield simpler models.

2.2.2.1 Method Based on Constrained Rigid-Body Dynamics. The majority of current knee models are 3-D dynamic models (Wismans et al., 1980; Garg & Walker, 1990; Huiskes & Blankevoort, 1990; Blankevoort et al., 1991). These models are dynamic because the research attempts to determine parameters that are related to forces or loads. One goal of these models may be to estimate parameters that cannot be determined in vivo. For example, the reference strain of a ligament has been determined using this method (Van Eijden et al., 1986).

The locations of the components are determined using Newton's laws combined with kinematic constraints known as *rigid-body dynamics.* The implementation of rigid-body dynamics is known to be computationally intensive, and the stability and correctness of such a model depends highly on initial conditions and the step size employed in the numerical computation.

Kinematic constraints are usually determined from measurements reported in the literature. For example, the contact point location or the orientation of some of the components can be used to reduce the degree of complexity of the model. Generally, however, these models are adapted only to the joint used and they cannot be used for geometrical models of any shape and size (Huiskes et al., 1985). To our knowledge, no model of this type has been successfully applied to the knee joint

for the complete range of motion of the flexion-extension (Hefzy & Grood, 1996). Also, none of these models include the screw-home motion.

Dynamic models may also be created to simulate dynamic movement (Hodgings et al., 1995). In this case, the models typically do not require a detailed model of the joint because the intent is to simulate realistic overall movement of a virtual human, and simple geometrical models often provide satisfactory results. If dynamic models are used at a more detailed level of the joint, they are typically nonstable and computationally intensive.

2.2.2.2 Methods Using Approximation. A simple way to model joint motion is to manually specify the position and orientation of the components for any attitude (the flexion and varus-valgus angles) of the joint in the range of motion considered. This task is tedious and must be repeated for any new geometrical model. Moreover, the complex shape of the contacting surfaces precludes a clear view of the contacting surfaces, so it is difficult to verify that the modeling yields a smooth and geometrically convincing motion.

Numerous researchers are using motion curves that represent the trajectory of the components as the ranges of motion on the degrees of freedom are spanned (Ounpuu, Gage, & Davis, 1991). Such curves are challenging to measure. Moreover, the motion measured on a specific joint may not be easily adapted to other joints because measurements found in the literature are usually not accompanied with the geometrical model of the components and the frame of reference necessary to produce the same motion.

To reduce the complexity of the computation, planar models are often considered (Yamaguchi & Zajac, 1989). The technique typically consists of mathematically approximating the contour of the extremities of the tibia and femur to establish their relative positions. A simple approximation is to define the femur condyle as an ellipse or a spiral (Rehder, 1983; Kurosawa, 1985; Loch, 1992), while the tibial plateaus are fitted to planes (Delp et al., 1990). Then, by setting the location of the contact points of the condyles on the tibia from data extracted from the literature (Nisell, 1985), the relative motion of the bones can be approximated.

A method closest to the one we shall present corrects the motion parameters extracted from experimental data by constraining the femur condyles in exact contact with the tibia. The algorithm employs the distance between control points of B-splines approximating the femur condyles and the tibia (Walker et al., 1988; Ateshian, 1993). This method uses the geometry of the joint and thus can be easily adapted to any joint. However, this method considers only two contact points, which is not realistic. Another concern with this model is that it does not yield any screw-home motion.

2.3 Augmented Reality and the VRDA Tool

One of the most promising and challenging future uses of HMDs is in applications in which virtual environments enhance rather than replace real environments. This is referred to as *augmented reality* (Bajura, Fuchs, & Ohbuchi, 1992). To obtain an enhanced view of the real environment, users wear see-through HMDs to see 3-D, computer-generated objects superimposed on their real-world view. A discussion of two technology types—optical and video see-through HMDs—is given in another paper of this special issue (Rolland & Fuchs, 2000).

An augmented-reality visualization tool for teaching the motion of anatomical joints (the VRDA tool) will enable a user manipulating the joint of a subject to visualize a virtual model of the inner bony anatomy superimposed on the limb as depicted in figure 1. Understanding the 3-D relationships of internal anatomical structures and the significance of body part movements is essential for the clinical examination of patients, the understanding of normal and pathological conditions, and the planning of treatment. Most students in medically related studies currently learn anatomy with a variety of limited formats including 2-D printed photographs, slides, labeled drawings, and cadaver dissection labs. Medical education, in particular, includes the clinical examination of patients and radiographic correlation with gross anatomy and pathology. Traditional methods often do not allow simultaneous visualization of both internal and external structures. Interactive videodisc, multimedia presentations, and computer dissection

simulations have been implemented and evaluated as successful. Video and computer-based demonstrations of dissections are infinitely reversible and repeatable, but they do not integrate the palpation of external anatomical landmarks. Electronic tools also do not provide the spontaneous feedback that is involved with living human models.

Because of the limitations of these traditional approaches to anatomy instruction, students may have artificial limits on their ability to quickly understand and apply the concepts. Perhaps a new teaching approach—such as the direct visualization of scaled, internal, 3-D anatomical structures in motion superimposed on the body (as done with the VRDA)—would help students form more-accurate mental models of joint motions in shorter periods of time compared to current learning processes (Wright et al., 1995). While the early versions of the VRDA tool may not simulate some of the more complex movements and elastic tissue deformations with a microscopic level of accuracy for joint movements, it certainly has a level of accuracy sufficient to provide an effective demonstration of gross joint movements and effectively demonstrate the 3-D nature of dynamic joint functions. Compared to traditional 2-D or static models, the VRDA tool offers distinct educational advantages, such as that the user interacts with the whole live model while positioning, for a more holistic approach to learning, rather than reducing the study of anatomy to one isolated and disarticulated limb at a time.

3 Algorithm Operations

3.1 Overall Approach

During the real-time use of the VRDA tool, the position of a reference point on the joint of a human subject and the orientation of the joint components are determined using a 0.1 mm accuracy and resolution optical tracker from Northern Digital. Based on the tracker data and a kinematic model of motion, a computer-generated image of a knee joint is generated and located correctly on the subject's knee. The user can see the joint image moving with the motion of the human subject. The interactive-speed rendering of the joint image is achieved by using a look-up table generated during the offline modeling procedure (Baillot, 1999).

For the kinematic model, we have developed an algorithm that allows automatic modeling of the kinematics of a 3-D geometric model of the bones of the knee joint (Baillot, Rolland, & Wright, 1998, 1999). We chose an automatic technique because a motion model can take months to construct if the bones are manually placed. For comparison, the automatic approach to modeling permits one to scale and animate a joint or some of its components within a couple of hours. We anticipate that the method we have developed can yield animations within minutes given that the algorithm is well adapted for parallelization.

The method integrates collision detection, stability detection, and a two-component algorithm which allows the automatic translation and rotation of two objects entering in collision until they reach a stable position. While ligaments and muscles are taken into account in creating the animation of the bones and all components of a joint will ultimately be represented, to simplify the model initially only the bones are demonstrated during flexion-extension of the joint. While it is beyond the scope of the current work to represent the ligaments as deformable models, their main function that consists in guiding the bones is accounted for. Finally, the cartilage is considered to be part of the rigid body to which it is attached as a consequence of the small deformations during motion (Frankel, 1997).

The physically based modeling approach we propose gives a model with 3-D motion capabilities that include flexion, screw-home, and varus/valgus angles. A strength of the method is that it yields no gap or intersection between the bones on the whole range of motion, regardless of the bone geometry we start with. Most importantly, the method enables one to animate any 3-D joint model, regardless of the size of its component parts. An inherent weakness of most current models (including ours) is that they do not account for the sliding of the menisci between the tibia and femur, which has only been recently described in the literature. Once more data describing these movements become available in the literature, they can be incorporated as necessary. The menisci are thus included in the geo-

metrical model as rigid bodies fixed to the tibia. Also, while motion of the menisci would yield a more comprehensive model, such level of accuracy is not critical for teaching radiologic positioning, the main application of the tool.

The method assumes that the muscles initiate the joint motion. The ligaments are considered to constrain and stabilize the bones in their optimal positions and orientations towards equilibrium. In the knee joint, for example, the anterior cruciate, posterior cruciate, and lateral ligaments constrain the patella, the femur, and the tibia to a stable configuration. It is then assumed that the geometry of the contact surfaces influences the degrees of freedom and the range of motion of a joint.

The relative orientation and position of the tibia with respect to the femur drives the motion of the virtual model accordingly. We consider in this model that the ligaments produce a resulting force whose direction is assumed to be constant. This direction is taken along the main axis of the tibia. We then determine for each given attitude of the joint (flexion and varus-valgus angles) the optimal relative position and orientation of the bones as if one were pushing them together along this specific direction.

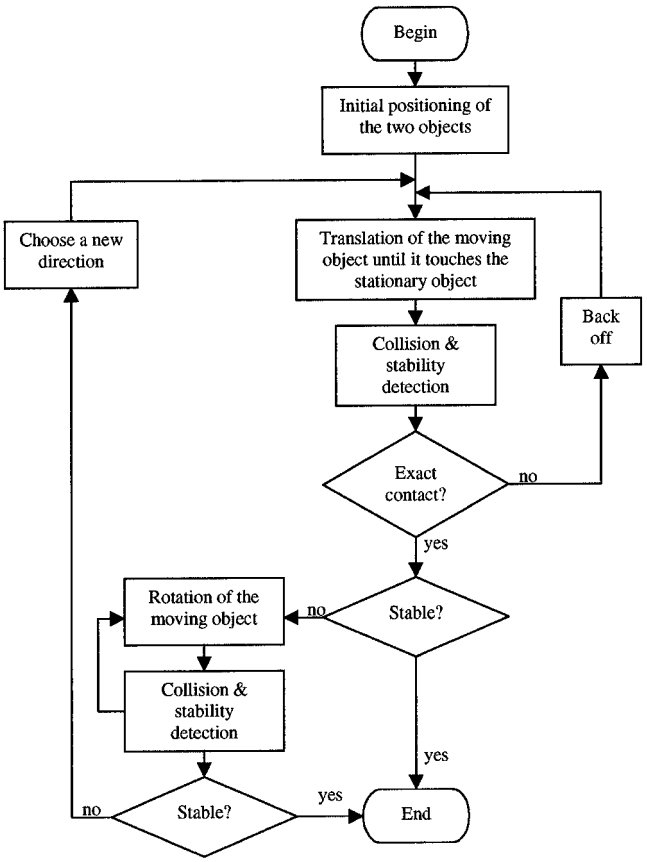

Figure 3. *Flowchart summarizing the working principle of a step of the modeling algorithm.*

3.2 Automatic Modeling Algorithm

The method we developed is general in the sense that it can be applied to all anatomical joint types. The flowchart of the working principle of the algorithm is depicted in figure 3.

The algorithm consists of two rigid-body maneuvers: translation and rotation, and collision and stability detection. Translation is defined as all points on the rigid body are moving along paths parallel to each other. Rotation is defined as there is one point on the rigid body or the extended part of the body having zero velocity. This point is called the center of rotation with all other points moving along circular paths centered around it. A collision detection algorithm is used to determine whether the rigid bodies are intersecting each other. Stability detection is used to decide if the two rigid bodies are in a stable position. A stable position is defined as neither translation nor rotation are allowed from that position.

The algorithm moves one rigid body toward the other object until they are in an exact contact and stable position. It starts with translating the rigid body of a fixed displacement along one direction. After the collision detection algorithm detects an intersection, it is backed off to establish an exact contact with the other object. Exact contact between two objects corresponds to a relative position where the object's proximity is within the contact tolerance (± 0.1 mm). Therefore, when two objects are in exact contact, they are not in collision. Then, the rigid body is rotated about a center of rotation until it is constrained, which can be decided by collision detection. The translation is resumed along another direction. The process iterates between translation and rotation until a stable position is reached. Figure 4 shows the example of a bar falling between two pyramids forming a sink. Again, three frames are shown to represent the process of settling the bar to a stable position.

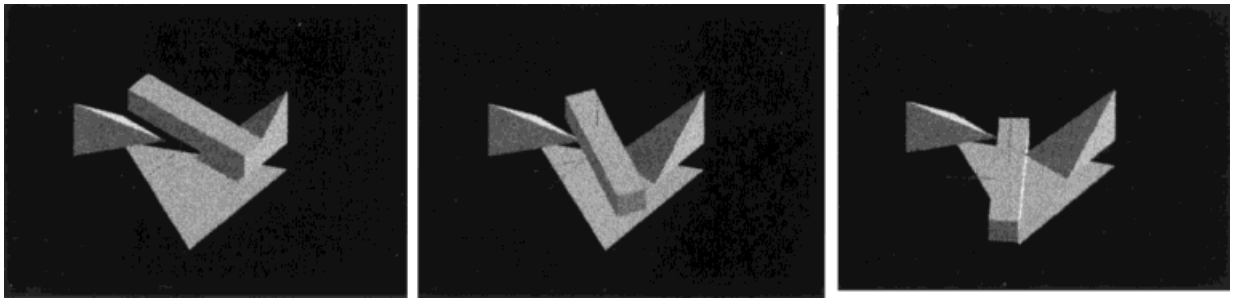

Figure 4. *Demonstration of the torque capability of the algorithm. The V-shaped object on either side of a bar falling is done with two inclined triangles. The bar is dropped on one triangle, and a torque is correctly produced on the bar as it touches the two triangles.*

3.3 Collision Detection

Exact collision detection between two objects determines which polygons forming the geometric models of the objects are intersecting in a given attitude. An original incremental modeling algorithm uses this information to make the joint elements slide against each other in a final stable position and orientation (Baillot & Rolland, 1999).

A C library—called RAPID and developed at the University of North Carolina at Chapel Hill—is used to solve the problem of finding the collisions (Gottschalk, Lin, & Manocha, 1996). This library returns the list of the intersecting triangles of the colliding geometric rigid-body models. The polygonal model was thus first transformed into triangle primitives. The normals of the geometric model are specified for each triangle vertex and averaged over each triangle to obtain the normal at each center. These vectors give the directions and application points of the reaction forces that appear when collision occurs between two rigid bodies. These vectors also indicate the direction in which the rigid bodies must move to stop the collision. The approximation of taking a center of a triangle as the collision point is valid in the model developed because the resolution of the contacting points of the surfaces (the sizes of the triangles) is small with respect to the bone sizes and of the order of magnitude of the modeling resolution.

3.4 Optimal Positioning Paradigm

Generally stated, one of the rigid bodies, the *reference,* is fixed in space. The other rigid bodies are translated towards, or slide or rotate against, the reference in order to find their optimal position and orientation, given that the motion is constrained within a few degrees of freedom. In the case of the femur-tibia relationship, the femur is the reference, and the tibia is moved with respect to the femur during the modeling.

The motion of the moving rigid body is initially set to the estimated direction of the resulting force that would be produced by the ligaments in a real knee under various manipulations of the joint. A first translation allows placing the rigid bodies in preliminary contact. Then, at each step, the procedure verifies if there were collisions during the last motion (translation or rotation). If no collision occurred, the original direction of motion is reestablished. If there were a collision, exact contact is made along the last performed motion (translation or rotation).

At this point, one of three scenarios could occur in the physical world: the moving solid could rotate (due to a torque), translate, or adopt an equilibrium position. The moving solid can translate along only a set direction imposed by ligament forces and in the plane orthogonal to that direction, and can rotate around only that direction axis, the other orientations being fixed for a specific given attitude. The translation along the direction is performed when no collision occurred at the last step. The translation in the plane orthogonal to the direction is performed if no rotation around the direction axis is possible as a result of a torque.

3.5 Algorithm Convergence

A cycle detection procedure has been implemented to avoid problems of convergence when large modeling

steps are employed or computational errors of the result occur in the case of the translation. A cycle here refers to the fact that the position output by the algorithm oscillates between two values. Therefore, the procedure first verifies if the current computed attitude of the solid yields an attitude equal to a previous step. If the difference in angle is less than the angular resolution and the difference in position is less than the translation resolution, a cycle is detected and the modeling step size corresponding to the last motion for either orientation or translation is divided in half. The step size is reset to its original value if either the original direction of motion is reestablished following a rotation or if a torque is produced after a translation. When collision occurs during the last computed motion, reducing the step size allows the resulting motion to become finally smaller than the required resolution. This allows the algorithm to converge effectively to a stable position.

The cycle detection procedure also solves a problem that could occur when a torque should be produced, but does not occur as a consequence of the solid moving stepwise. Consider the bar falling between two surfaces shown in figure 4. When the bar touches both surfaces, it rotates as a consequence of a torque. The original position of the bar can be such that the bar never touches both surfaces at the same time, because there is no collision when the exact contact occurs and because the bar is moved in steps that are too large. However, in such a case, a cycle would be produced, and the step size would be regularly decreased, bringing the bar slowly towards a position where both surfaces touch and consequently produce a torque on the bar. This method allows motion using adaptive modeling steps within the size of an object for fast convergence without unstable behavior.

3.6 Implementation

The algorithm was implemented on an SGI Onyx Deskside with two processors running at 150 MHz. We used a high-resolution geometric model (from Viewpoint Datalabs, Inc.) that includes geometric models of the bones, the meniscus, the ligaments, the tendons, and the muscles. Each polygonal model was transformed into triangle primitives for compatibility with the collision detection engine. The model was described in the extension position, vertically, with the patella in front. In this configuration, all the bones had their origin at the same location and in the same orientation. We considered two parameters in our model (the flexion and the varus/valgus angles). For both orientation and translation, we used a modeling step ten times larger than the associated resolution. The resolution of the human eye is one arc minute, and the viewing distance of the model is typically 0.5 m in our application. The maximum resolution in translation resolvable by the human eye was thus 0.15 mm. We set 0.1 mm as the translation resolution and 1 mm as the translation step during modeling. To obtain the corresponding resolution for the rotation, we accounted for the tibia being enclosed in a circle of radius 30 mm. Then an arclength of 0.145 mm must be produced by an angle of 0.27 deg. (0.145 mm/0.03 rad). We set the orientation resolution to 0.1 deg., and the modeling rotation to 1 deg.

4 Visualization With Augmented Reality: Methods

A first visualization of the superimposition of a graphical knee joint on a real counterpart was performed using the bench prototype, optical see-through HMD designed and calibrated for perception experiments (Rolland, Ariely, & Gibson, 1995; Rolland et al., 1997). The bench prototype has the capability of adjusting interpupillary distance and setting the optical virtual images at any distance of accommodation (0.8 m in this experiment), and is equipped with a linear head-motion parallax to simulate side-to-side head motion. The tracking of the user's head is done using optical encoders, and the anatomical part is also tracked optically using the OPTOTRAK 3020 from Northern Digital. A high-performance graphics C library called Performer was employed for interactive-speed rendering.

In this first implementation, we used a mannequin leg to perform the superimposition. Some infrared markers were placed around the top of the tibia and around the bottom of the femur to capture the orientation of the leg. We accounted for the fact that the infrared markers

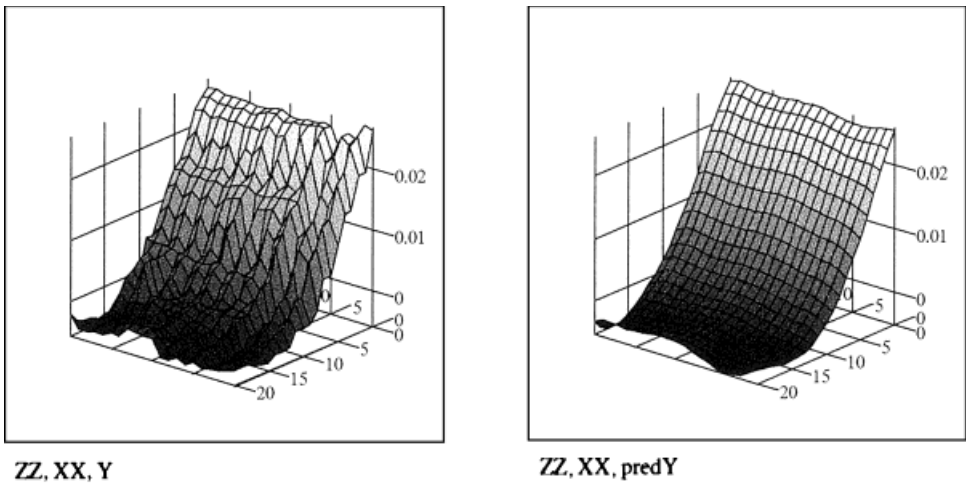

Figure 5. *An example of the smoothing of an ensemble of motion curves. In this case, the z coordinate of the tibia referential as a function of the flexion-extension angle on one axis and the varus/valgus angle on the other. (a) The original curve surface. (b) The same surface obtained after least-square fit using a polynomial of degree 6.*

are located on the surface of the leg instead of on the bones in computing the location of the virtual bones for superimposition. More-advanced methods were also developed for a real leg and will be reported elsewhere (Outters, Argotti, & Rolland, 1999). The graphical bones were used at their original scale; however, the mannequin was female while the geometric model came from a male cadaver, which will yield discrepancies in registration. Although the mismatch does not provide an ideal setting, it is sufficient for the goal of this experiment as we aim to demonstrate subjective transparency of the opaque leg with optical superimposition of the graphical joint when it is properly positioned and oriented with respect to the leg model.

5 Results and Discussion

5.1 Modeling Flexion-Extension of the Knee Joint

Let's define the anterior-posterior drawer effect as the excessive motion front to back of the tibia with respect to the femur when the knee is flexed. The curve of the translation produced by the AP-drawer effect, as a function of the flexion angle, showed some discontinuities in two places that were also observed during simulation. We understand that these discontinuities are due to the discrete nature of the model composed of either polygons or triangles. The discontinuity is a consequence of the polygon roughness that prevents the motion of the tibia during a range of flexion angles, yet at a certain angle the motion is possible because one stopping edge is favorably oriented. We performed a least-square polynomial surface regression to smooth the discontinuous motion curves. By using a polynomial of degree 6 for the two translations and the rotation curves, a smooth-motion model was rendered without perceptible collision. The smoothing of the surface curve of the z coordinate (the A-P drawer effect) as a function of the flexion and varus/valgus angles is shown in figure 5. The final rendered and smoothed models of the flexion-extension motion of a knee joint are shown in figure 6. Various snapshots acquired during the full cycle of flexion-extension from various viewpoints are shown. The lookup table used to animate the joint was constructed such that the condyles touch for every value of the flexion-extension angle.

5.2 Visualization with Augmented Reality

We focused our first experiments on assessing the subjective feeling of inclusion of the graphics within the

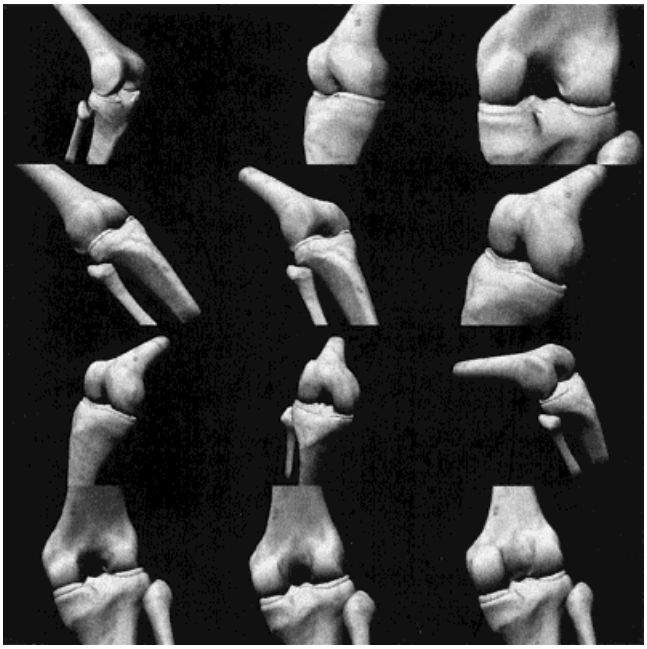

Figure 6. *The animated textured joint bones used in the VRDA tool. The look-up tables have been constructed to obtain a varus/valgus of zero (the two condyles touch) for all flexion-extension angles. These snapshots of the automatically animated knee joint during flexion-extension show multiple views of the joint with both condyles touching.*

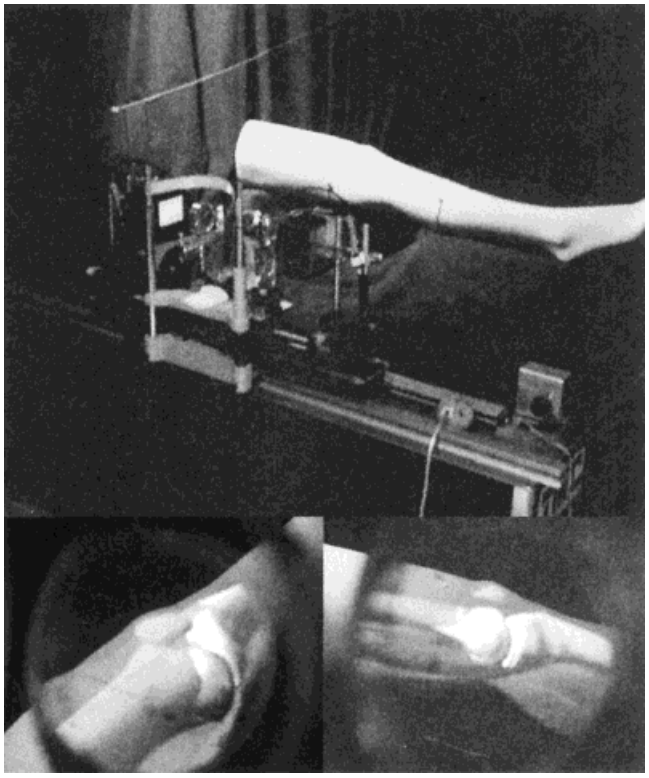

Figure 7. *First demonstration of the superimposition of a graphical knee joint superimposed on a leg model for use in the VRDA tool. (a) A picture of the bench prototype set-up and the leg model where the superimposition was conducted; (b) and (c) a snapshot of the superimposition through one lens of the set-up in (b) a diagonal view, and (c) a side view. Active infrared markers were placed on the leg model to capture the orientation of the leg. This information, as well as the calibration of the overall system (the set-up and the optical tracker) were used to achieve the superimposition.*

leg model from different viewpoints of the leg and the head. Two views taken behind one of the lenses of the stereoscopic display are shown in figure 7. In this visualization (in which only the head of the user moves), we included the bones, ligaments, and muscles which will be included in the full development of the VRDA tool. The subjective judgment of having "X-ray" vision through this leg model in the stereoscopic bench prototype display was powerful and convincing. This subjective judgment is also communicated quite vividly through the snapshots included in figure 7.

6 Conclusion and Future Work

Two milestones in the development of the VRDA tool were reported. First, a method was described for the automatic modeling of joint motions and the application to the flexion-extension of a knee joint. The method is based on collision detection and biomedical knowledge. The promise of the modeling technique lies in the precomputation of motion curves for the joint and the use of a look-up table generated from the modeling to render kinematic motion of a joint at interactive speed. Secondly, the first optical superimposition of the graphical knee joint including bones, ligaments, and muscles on a leg model is demonstrated. Results show promise for the success of the VRDA tool.

Acknowledgments

This research is supported by a First Award from the National Institute of Health, grant 1-R29-LM06322-01A1. We thank

Michael Moshell at the University of Central Florida and Saara Totterman at the University of Rochester for stimulating discussions about this research. We also thank Peter Loan from Musculographics, Inc., for discussing freely with us the basic motions involved in the knee joint and how current models were approximating them.

References

Ateshian, G. A. (1993). B-spline least-square surface fitting for articular surface of diathroidal joints. *J. Biomech. Eng. 115*(4), 366–373.

Baillot, Y. (1999). *First implementation of the Virtual Reality Dynamic Anatomy (VRDA) tool.* Unpublished master's thesis. University of Central Florida, Orlando, FL.

Baillot, Y., & Rolland, J. P. (1999). Novel algorithms for automatic modeling of joint motions. U.S. patent application.

Baillot, Y., Rolland, J. P., & Wright, D. L. (1998). Modeling of a knee joint for the VRDA tool. *Proc. of Medicine Meets Virtual Reality 98,* 366–367.

———. (1999). Kinematic modeling of knee-joint motion for a virtual reality tool. *Proc. of Medicine Meets Virtual Reality 99,* 30–35.

Bajura, M., Fuchs, H., & Ohbuchi, R. (1992). Merging virtual objects with the real world. *Computer Graphics, 26,* 203–210.

Blankevoort, L., & Huiskes, R. (1988). Ligament-bone interaction in a three-dimensional model of the knee. *J. Biomech. Eng., 113,* 263–269.

Blankevoort, L., Kuiper, J. H., Huiskes, R., & Grootenboer, H. J. (1991). Articular contact in a three-dimensional model of the knee. *J. Biomech., 24*(11), 1019–1031.

Bylski-Austrow, D. I., Ciarelli, M. J., Kayner, D. C., Mattews, L. S., & Goldstein, S. A. (1994). Displacements of the menisci under joint load: An in-vitro study of human knees. *J. Biomech., 27*(4), 421–431.

Conlay, M. D., & Long, G. L. (1994). Ligament forces, condyle reactions, line geometry, screw theory and human knee stability. In *Advances in Bioengineering American Society of Mechanical Engineers, Bioengineering Division (publication) BED* 28 (pp. 169–170). New York: AMSE.

Davis, L., Peuchot, B., & Rolland, J. P. (1999). The calibration of optical tracking probes and measurement of precision and accuracy [Abstract]. *Optical Society of America Annual Meeting.*

Delp, S. L., & Loan, J. P. (1995). A graphics-based software system to develop and analyse models of musculoskeletal structures. *Comput. Biol. Med., 25*(1), 21–34.

Delp, S. L., Loan, J. P., Hoy, M. G., Zajac, F. E., Topp, E. L., & Rosen, J. M. (1990). An interactive graphics-based model of the lower extremity to study orthopaedic surgical procedures. *IEEE Transactions on Biomedical Eng, 37*(8), 757–767.

Frankel, V. H. (1997). Personal communications.

Frankel, V. H., & Nordin, M. (1980). *Basic Biomechanics of the Skeletal System.* Philadelphia: Lea and Feibiger.

Garg, A., & Walker, P. S. (1990). Prediction of total knee motion using a three-dimensional computer-graphics model. *J. Biomech., 23*(1), 45–58.

Gottschalk, S., Lin, M. C., & Manocha, D. (1996). OBBTree: A hierarchical structure for rapid interfence detection (TR96-013). Department of Computer Science, University of North Carolina, Chapel Hill.

Hefzy, M., & Grood, E. (1996). Review of knee models: 1996 update. *App. Mech. Rev., 41,* 1–13.

Hodgings, J. K., Wooten, W. L., Brogan, D. C., & O'Brien, J. F. (1995). Animating human athletics. *Proc. of SIGGRAPH,* 71–78.

Huiskes, R., & Blankevoort, L. (1990). The relationship between knee joint motion and articular surface geometry. In *Biomechanics of Diarthrodial Joints* (pp. 269–286). New York: Springer-Verlag.

Huiskes, R., Kremers, J., Lange, A., Woltring, H. J., Selvik, G., & Rens, Th. J. G. van (1985). Analytical stereophotogrammetric determination of the three-dimensional knee-joint geometry. *J. Biomech., 18,* 559–570.

Kurosawa H., Walker, P. S., Abe, S., Garg, A., & Hunter, T. (1985). Geometry and motion of the knee for implant and orthopaedic design. *J. Biomech., 18,* 487–499.

Loch, D. A., Luo, Z., Lewis, J. L., & Stewart, N. J. (1992). A theoretical model of the knee and ACL: Theory and experimental verification. *J. Biomech., 25*(1), 81–90.

Nisell, R. (1985). Mechanics of the knee: A study of joint and muscle load with clinical applications. *Acta Orthop. Scand., 56*(216), 5–41.

Nordin, M., & Frankel, V. H. (1980). *Basic Biomechanics of the Skeletal System* (ch. 4). Philadelphia: Lea and Febiger.

Ounpuu, S., Gage, J. R., & Davis, R. B. (1991). Three-dimensional lower extremity joint kinetics in normal pediatric gait. *J. Pediat. Orthop., 11,* 341–349.

Outters, V., Argotti, Y., & Rolland, J. P. (1999). *Knee motion capture and representation in augmented reality* (Technical Report TR99-006). University of Central Florida, Orlando, FL.

Rehder, U. (1983). Morphometrical studies on the symmetry of the human knee joint: femoral condyles. *J. Biomech., 16,* 351–361.

Rolland, J. P. (1998). Mounted displays. *Optics and Photonics News, 9*(11), 26–30.

Rolland, J. P. (2000). Wide-angle, off-axis, see-through, head-mounted display. *Optical Engineering—Special Issue on Pushing the Envelope in Optical Design Software, 39*(7), (in press).

Rolland, J. P., Ariely, D., & Gibson, W. (1995). Towards quantifying depth and size perception in virtual environments. *Presence: Teleoperators and Virtual Environments, 4*(1), 24–49.

Rolland, J. P., & Fuchs, H. (2000). Optical versus video see-through head-mounted displays in medical visualization. *Presence: Teleoperators and Virtual Environments, 9*(3): 287–309.

Rolland, J. P., Wright, D. L., & Kancherla, A. R. (1997). Towards a novel augmented-reality tool to visualize dynamic 3-D anatomy. *Proceedings of the Medicine Meets Virtual Reality, 5,* 337–348.

Rolland, J. P., Yoshida, A., Davis, L., & Reif, J. H. (1998). High-resolution inset head-mounted display. *Applied Optics, 37*(19), 4183–4193.

Soudan, K., Van Audekercke, R., & Martens, M. (1979). Methods, difficulties and inaccuracies in the study of human joint kinematics and pathokinematics by the instant axis concept. Example: the knee joint. *J. Biomech., 12,* 27–33.

Totterman, S., Tamez-Pena, J., Kwok, E., Strang, J., Smith, J., Rubens, D., & Parker, D. (1998). 3D visual presentation of shoulder joint motion. *Proc. of MMVR* 98, IOS Press, 27–33.

Van Eijden, T. M. G. J., Kouwenhoven, E., Verburg, J., & Weijs, W. A. (1986). A mathematical model of the patello-femoral joint. *J. Biomech., 19,* 219–229.

Walker, P. S., Rovick, J. S., Robertson, D. D., & Schrtager, R. J. (1988). The effects of knee brace hinge design and placement on joint mechanics. *J. Biomech., 21,* 965.

Wismans, J., Veldpaus, F., Janseen, J., Huson, A., & Struben, P. (1980). A three-dimensional mathematical model of the knee-joint. *J. Biomech., 13,* 677–685.

Wright, D. L., Rolland, J. P., & Kancherla, A. R. (1995). Using virtual reality to teach radiographic positioning. *Radiologic Technology, 66*(4), 233–238.

Yamaguchi, G. T., & Zajac, F. E. (1989). A planar model of the knee joint to characterize the knee extensor mechanism. *J. Biomech., 22*(1), 1–10.