

# From robotic hands to human hands: a visualization and simulation engine for grasping research

*A. Miller and P. Allen*

Department of Computer Science, Columbia University, New York, USA

*V. Santos*

Neuromuscular Biomechanics Laboratory, Cornell University, New York, USA, and

*F. Valero-Cuevas*

Neuromuscular Biomechanics Laboratory, Cornell University, New York, USA  
The Hospital for Special Surgery, New York, USA

### Abstract

**Purpose** – Robotic hands are still a long way from matching the grasping and manipulation capability of their human counterparts, but computer simulation may help us understand this disparity. We present our publicly available simulator, and describe our research projects involving the system including the development of a human hand model derived from experimental measurements.

**Design/methodology/approach** – Unlike other simulation systems, our system was built specifically to analyze grasps. It can import a wide variety of robot designs by using standard descriptions of the kinematics and link geometries. Various components support the analysis of grasps, visualization of results, dynamic simulation of grasping tasks, and grasp planning.

**Findings** – The simulator has been used in several grasping research problems and can be used to plan grasps for an actual robot. With the aid of a vision system, we have shown that these grasps can be executed by a robot.

**Research limitations/implications** – We are currently developing methods to handle deformable surfaces, tendon driven models, and non-ideal joints in order to better model human grasping.

**Practical implications** – This work is part of our current project to create a biomechanically realistic human hand model to better understand what features are most important to mimic in the designs of robotic hands. Such a model will also help clinicians better plan reconstructive hand surgeries.

**Originality/value** – We describe our publicly available grasping simulator and review experiments performed with it. The paper demonstrates the usefulness of this system as a tool for grasping research.

**Keywords** Robotics, Materials handling, Simulation

**Paper type** Research paper

## 1. Introduction

The progress in building capable robotic hands has been slow. An important factor that has affected the progress in this field is the lack of easily obtainable, low cost experimental robotic hands that can be used as test beds. The high cost and difficulty of building a robotic hand, along with the associated

electronics, control systems, and integrated sensing has led to a serious lack of experimental devices in the field. In fact, were one desirous of purchasing a robotic hand today, there appear to be few, if any, available. Custom designs exist, but usually are lacking in the higher levels of system integration that can turn a mechanical device into a full-fledged grasping system.

---

The Emerald Research Register for this journal is available at

[www.emeraldinsight.com/researchregister](http://www.emeraldinsight.com/researchregister)

The current issue and full text archive of this journal is available at

[www.emeraldinsight.com/0143-991X.htm](http://www.emeraldinsight.com/0143-991X.htm)



Industrial Robot: An International Journal  
32/1 (2005) 55–63  
© Emerald Group Publishing Limited [ISSN 0143-991X]  
[DOI 10.1108/01439910510573309]

---

This paper is based upon work supported by the National Science Foundation under Grant No. 0312271 (ITR project) from NSF's Robotics and Computer Vision Program, NSF Grant No. 0237258 (CAREER award) from NSF's Biomedical Engineering/Research to Aid Persons with Disabilities Program (to F. J. Valero-Cuevas), a Biomedical Engineering Research Grant from the Whitaker Foundation (to F. J. Valero-Cuevas), and an NSF Graduate Research Fellowship (to V. J. Santos).

This paper is a revised version of one previously published at the "Intelligent Manipulation and Grasping International Conference" (IMG'04), 1-2 July 2004, Genova, Italy

What can push this research further along? How can we design intelligent systems with the ability to grasp objects? A partial solution to this problem is to use computer modeling and simulation to effectively design and test robotic hands in typical task environments. As the computer models get better, and the hardware faster, realistic simulation can be used to learn more about robotic grasping.

At Columbia University, we have created a publicly available simulator to serve as a useful tool for grasping research. The system, called *GraspIt!* [1] can accommodate a wide variety of hand and robot designs. It includes a rapid collision detection and contact determination system that allows a user to interactively manipulate the joints of the hand and create new grasps of a target object. Each grasp is evaluated with numeric quality measures, and visualization methods allow the user to see the weak point of the grasp and create arbitrary 3D projections of the 6D grasp wrench space. The dynamics engine within *GraspIt!* Computes the motions of a group of connected robot elements, such as an arm and a hand, under the influence of controlled motor forces, joint constraint forces, contact forces and external forces. This allows the dynamic simulation of an entire grasping task, as well as the ability to test custom robot control algorithms.

In this paper we present further details regarding the various components of the system, and we also describe two projects which use *GraspIt!* as an integral part of a larger grasp planning system: one generates candidate grasps using heuristics and predefined grasp taxonomies, and the other uses a support vector machine to learn high quality grasps of parametric objects. Finally, we discuss our development of a human hand model that uses accurate geometry and kinematics derived from experimental measurements. This is part of an ongoing project to create a biomechanically realistic human hand model to better understand what features are most important to mimic in the designs of robotic hands. Such a model will also enable studies of the functional abilities of the intact and impaired human hand using the rigorous mathematical framework of robotics.

## 2. *GraspIt!*

In building *GraspIt!*, we were aware of several commercial robotics simulators available, including Delmia's IGRIP, Flow Software Technologies' Workspace5, MCS Software's ADAMS, and the Easy-Rob system, as well as past and present research projects in robot simulation, including among others Corke's Robotics Toolbox for MATLAB (Corke, 1996), Speck and Klaeren's RoboSiM (Speck and Klaeren, 1999), and Ruspini and Khatib's Simpack (Ruspini and Khatib, 1999). There are a number of important elements that set *GraspIt!* apart from this body of previous work, the most important being that it has been developed *specifically* to analyze and visualize the robotic grasping task. Accordingly, none of the simulators above have the capability to compute dynamic frictional contact forces accurately, which is an essential component of grasping problems. Further, they also lack the grasp analysis and planning algorithms that are incorporated in *GraspIt!*. Finally, *GraspIt!* is a total package that includes not only models and a simulation engine, but a powerful and advanced user interface that facilitates its use. In the following sections we provide a brief overview of the components and features of the system, as well as describe some of its applications. Further

information regarding the system and its applications can be found in the following papers (Kragić *et al.*, 2001; Miller and Allen, 1999, 2000; Miller and Christensen, 2003; Miller *et al.*, 2003; Pelossof *et al.*, 2004).

### 2.1 *GraspIt!* world elements

#### 2.1.1 Body types

A basic body consists of a pointer to its geometry, a material specification, a list of contacts, and a transform that specifies the body's pose relative to the world coordinate system. The body geometry is read from an Inventor model file that has essentially the same format as VRML 1.0. The material is one of a set of predefined material types and is used when computing the coefficient of friction between two contacting bodies.

A dynamic body inherits all of the properties of a body and defines the mass of the body, the location of its center of mass, and its inertia tensor. If the mass distribution is unknown, the system can compute the inertia tensor by assuming a uniform mass distribution. The reason for distinguishing between bodies and dynamic bodies is that some bodies are simply considered obstacles, and while they are elements of the collision detection system and can provide contacts on other bodies, they are not part of the dynamics computations and remain static. This makes it possible to create a complex world full of obstacles without making the dynamics intractable to compute.

#### 2.1.2 Robots

We have tried to make the definition of a robot as general as possible to allow a wide variety of robot designs to be imported. The system reads the kinematic parameters (specified in standard Denavit-Hartenberg notation) for each chain of links and joints from a text file, loads the geometry for each link, and constructs the robot. Our definition separates degrees of freedom (DOF) from joints and allows multiple joints to be driven by the same DOF because it is common in many hand designs to have coupled joints that are passively controlled by other joints.


A hand is a special type of robot that can form grasps of objects, and these grasps will be analyzed by the system. It also includes an auto-grasp method which closes the joints of the hand at preset velocities. Each joint stops when it has reached its limit or when a link that follows it in the kinematic chain contacts an object or another finger link.

#### 2.1.3 The robot library

The ability to easily add new robot designs is a key benefit of our system. It is a relatively simple process of specifying the parameters required by the configuration file, creating the link geometry files, and in most cases takes less than an hour or two to set up. We have already created models of a parallel jaw gripper, a Puma 560 arm, and a simplified Nomadics XR4000 mobile robot. Additionally, through collaboration with other research sites we have obtained CAD models and kinematic descriptions of four different articulated hand designs (Plate 1).

Another feature of *GraspIt!* is the ability to attach multiple robots to create robotic platforms. The system allows the definition of a tree of robots where any number of robots can be attached to the last link of a kinematic chain of another robot. This allows us to construct more complex manipulation systems, and together with a world model specifying the robot's environment, we can plan and test our entire grasping task so that we can avoid planning grasps that will conflict with these obstacles.

**Plate 1** *GraspIt!* robot models of (from left to right): a parallel jaw gripper, Barrett hand, DLR hand, NASA Robonaut hand, Rutgers hand, and Puma 560 arm



## 2.2 User interface

One of the design goals we believe we achieved was to make the user interface as intuitive and transparent as possible. When a user starts a new session, he is presented with an empty world into which he can import new obstacles, graspable bodies, or robots, and at any point the current state of the world can be saved to be loaded again later in another session. The primary element of the main window is a standard viewer, which displays a projection of a 3D world in a 2D window. The virtual camera can be rotated, panned, or zoomed, and a seek tool allows close up inspection of a particular detail in the scene.

When the dynamics are not being used, obstacles, graspable bodies, and robots may be translated and rotated in 3D using an appropriate manipulator that appears upon clicking on the object. Manipulating the individual degrees of freedom of a robot is equally intuitive. Clicking on a kinematic chain of the robot brings up an interactive manipulator for each actively controllable joint in the chain. Revolute joints are controlled by dragging a disc whose axis of rotation is coincident with the joint axis (Figure 1), and prismatic joints are controlled by dragging an arrow which is aligned with the joint axis. These manipulators obey the joint limits defined in the robot configuration file and prevent the user from moving beyond them.

Another important feature is the ability to interact with portions of *GraspIt!* through MATLAB. The system can be

run as a server that accepts TCP socket connections, and then in MATLAB, compiled MEX functions are used to communicate with the server to set robot motor torques, step the dynamic simulation by one time step, and retrieve the state of the world including all of the contact forces and the current position and velocity of each body. This allows external MATLAB functions to control the simulation.


## 2.3 Contacts

To prevent bodies from passing through each other while they are being manipulated by the user, the system performs real-time collision detection using the Proximity Query Package (Larsen *et al.*, 1999). If a collision is detected, the motion of the bodies must be reversed back to the point when the contact first occurs. To find this instant, *GraspIt!* performs a binary search that ends when the bodies are separated by a distance that is less than the contact threshold (currently set at 0.1 mm). Then the system determines the convex regions of contact between the two bodies and draws a red friction cone at each contact point, which serves to visually mark the position of the contact and its normal. The width of this cone identifies how large any frictional forces can be with respect to a force applied along the contact normal.

## 2.4 Grasp analysis

When *GraspIt!* is used in the static mode, the fingers of a hand may be closed around the object without causing it to move,

**Figure 1** User interface



**Note:** Left: The user interface makes manipulating joints simple. The angle of a revolute joint can be changed by dragging a disc manipulator located at the joint. The passive distal joint moves in a fixed relationship with the medial joint. The quality of a grasp is shown as “Unstable” until a force-closure grasp is formed. Right: A force-closure grasp of a mug using a parallel jaw gripper. The worst and average case quality values are displayed in the lower left portion as the values  $e$  and  $v$ . The pair of purple indicators show the force and torque components of the worst case disturbance wrench. In the upper left is a projection of the grasp wrench space that shows the space of forces that can be applied to the mug without creating a moment, and in the lower left is a projection that shows the space of torques that can be applied without a net force acting on the object.

and every time a contact is formed or broken the system evaluates the current grasp. This evaluation is done using one or more quality metrics. The current metrics evaluate the ability of the grasp to resist disturbance forces, but future metrics will evaluate object manipulability and the grasp's robustness with respect to errors in contact placement.

To evaluate a grasp's efficiency against disturbances, the system builds the 6D grasp wrench space using a convex hull operation on the set of possible contact wrenches given a unit strength grasp (Miller and Allen, 1999). The volume of this space is used as an average case quality measure, since the larger the space is, the more efficient the grasp is. The point on the hull boundary that is closest to the origin represents the grasp's weakest point (i.e. the wrench that is most difficult for the grasp to apply). The distance to this point is used as a worst case quality measure.

These numeric results allow an objective comparison between grasps, but in some cases it is desirable to visualize more of the grasp's characteristics. To support this, the system can produce arbitrary 3D projections of the grasp wrench space which give a better sense of a particular grasp's strengths and weaknesses. It can also display the worst case disturbance wrench, which is the force-torque combination that is most difficult for the grasp to resist. Figure 1 shows an example of these results, and demonstrates why parallel jaw grippers are not good at grasping round objects. Because the flat plates only contact the mug surface in two places, the grasp cannot easily resist a torque about the axis between those two contact regions, as shown by the purple worst case disturbance wrench indicator.

### 2.5 Simulating dynamics

The system allows the user to form grasps of an object and analyze these grasps without using dynamics. To study the evolution of a grasp and to test grasping control algorithms, however, we must consider how the hand and object move over time under the influence of controlled motor forces, gravity, inertial forces, and in response to collisions. To compute the motion of each dynamic body in the world, we use a numerical integration scheme that computes the change in velocity of each body over a small finite time step given a set of external forces acting on the body and any joint constraints, non-interpenetration constraints, and friction constraints that may also be present. These constraints are linearized and formulated as a linear complementarity problem (LCP) (Anitescu and Potra, 1997), that is solved using Lemke's algorithm (Cottle *et al.*, 1992) a pivoting method similar to the simplex algorithm for linear programming. The solution provides not only the new velocity of the bodies, but also the normal and frictional impulses at the contact points. After each iteration of the dynamics is completed, the system can draw the contact forces at each contact point, and at any time the dynamics may be paused to examine the state of the system or to change the current simulation parameters.

With the dynamics in place, it is possible to study the temporal formation of grasps. In the example presented in Figure 2, the Barrett hand is positioned above a wine glass which rests on its side on the table. The PD joint controllers of the Puma robot hold the wrist in place, and the desired final position of the hand joints is set to fully closed. The figure shows the initial setup and 5 different time slices during the simulation. The default time step is 2.5 ms, but smaller steps may occur due to contact events. Because there is very

little friction between the plastic and glass surfaces, and because the glass is tapered, the hand squeezes the glass out of its grasp. As the simulation continues, the wine glass slides and rolls off the table, hitting the Puma robot on its way down to the floor.

## 3. Applications

*GraspIt!* has become a platform that supports research in a variety of areas related to grasping. With its integrated grasp analysis methods, we have applied it to the problem of planning high quality grasps of objects. Below we describe two different systems we have built with *GraspIt!* that approach this challenging problem in different ways. Finally, we present some of our most recent work in creating a biomechanically realistic human hand model that will not only help clinicians better understand the mechanics of the hand and plan reconstructive surgeries, but will also help robotic hand designers build hands that come closer to matching the capabilities of our own hands.

### 3.1 Grasp planning


The grasp planning problem is extremely difficult because of the number of degrees of freedom of a robotic hand. For example, the relatively simple Barrett hand has ten degrees of freedom, six for orientation relative to the object, and four for longer manipulation. This number of DOFs creates a large search space of hand configurations. Of course, large portions of this space are worthless because the fingers are not in contact with the object, but even if the problem were reparameterized, a brute force search would still be intractable.

A variety of other approaches have been used to tackle this problem. A number of papers present contact-level grasp synthesis algorithms (Ding *et al.*, 2000; Mirtich and Canny, 1994; Ponce *et al.*, 1993). These algorithms are concerned only with finding a fixed number of contact locations without regard to hand geometry. Other systems built for use with a particular hand restrict the problem to choosing precision fingertip grasps, where there is only one contact per finger (Borst *et al.*, 1999; Hester *et al.*, 1999). These types of grasps are good for manipulating an object, but are not necessarily the most stable grasps because they do not use inner finger surfaces or the palm.

One way of limiting the large number of possible hand configurations is to use grasp preshapes. Before grasping an object, humans unconsciously simplify the task to selecting one of only a few different prehensile postures appropriate for the object and for the task to be performed. These postures have been enumerated in various grasp taxonomies (Napier, 1956).

Our first planner takes advantage of this fact, in order to reduce the search space size to a small number of grasps that are more likely to have a high quality. The system consists of two parts, one to generate a set of starting grasp locations based on a simplified object model, and one to test the feasibility and evaluate the quality of these grasps. The simplified object model is constructed from a small set of shape primitives such as spheres, cylinders, cones and boxes, and heuristic grasping strategies for these shapes allow the system to generate a set of grasp possibilities. The grasp tester moves the hand from a grasp starting position toward the object, closes the fingers around the object, and evaluates the

**Figure 2** The Barrett hand attempts a grasp of the wine glass, but due the low degree of friction at the contacts and the taper of the glass, the glass is squeezed out of the grasp



**Note:** The first frame shows the initial setup, and following from left to right, snapshots are taken at 0.5103, 0.5425, 0.6148, 0.6586, 0.6956 seconds of simulation time. The full movie is at <http://www.cs.columbia.edu/~amiller/grasplit/movies.html>.

grasp. After testing all of the generated grasp possibilities, the user is presented with the best grasps of the object in descending order of quality (Plate 2). In order to prevent infeasible grasps from being planned, the user may import a world model containing obstacles, as well as a robot arm model so that reachability constraints may be considered (Plate 3).


The drawback of this approach is that since it only considers a subset of the possible grasps, it may miss a better possibility. Humans may initially choose a sub-optimal grasp of a novel object, but with experience they will adapt their grasp to the most appropriate one for the task. Our most recent planning system applies new machine learning techniques to this problem to effectively have robots learn how to grasp arbitrary objects.

We are using supervised training to learn what is a good grasp for a robotic hand. This requires a method that allows us to try a large number of grasps of an object and report a metric on the quality of each grasp in the training set, and *GraspIt!* is perfectly suited for this. Using this training set, we


can then generate basis functions that can effectively both predict the quality of an arbitrary new set of grasping parameters and also use these basis functions to find an optimal set of grasping parameters for an object. These parameters correspond to the degrees of freedom of an actual hand, rather than the placement of point contacts on the object surface. It is also important to note that our method is not dependent on a single type of robotic hand or class of objects. It provides a robust system for testing different robotic hands and analyzing the quality space that they span.

For our tests we again used the Barrett hand, and the objects in our training set were nine different superellipsoids, which can be described with two shape parameters,  $\epsilon_1$  and  $\epsilon_2$ . For each superellipsoid we generated 1,600 grasps, which consisted of 100 random roll and spread angle combinations for each of 16 regularly sampled grasp starting positions, as shown in the left portion of Figure 3. This gave us a total of 14,400 grasps, which were evaluated by *GraspIt!* over the course of approximately 4h on a Pentium IV 2.3 GHz Windows machine.

**Plate 2** The first image shows the primitive model used for planning grasps for the coffee mug and shows the generated set of grasp positions to be tested. The other images show three of the best grasps of the mug in descending quality order



**Plate 3** The best planned grasp of the mug in the presence of obstacles and using the reachability constraints of the Puma arm



We then used an SVM regression to create a mapping between object shape, grasp parameters and grasp quality. Our learned regression mapping accepts a fixed length input vector that contains the shape and grasping parameters and returns a single scalar which estimates the grasp quality. If only provided with a subset of the input vector (i.e. the shape parameters), the regression algorithm will perform an efficient search for the optimal setting of the missing input vector values (i.e. the grasp parameters) that maximize grasp quality. The middle and right portions of Figure 3 show the results of these searches for both a shape that was in the training set and one that was not.

Our initial superquadric models are clearly not able to model the full range of possible objects to be grasped, so our next task is to extend this technique to more complex shapes, which could be composite models composed of multiple superquadric shapes. Possible methods for solving this

multi-superquadric grasping problem are discussed in Pelosoff *et al.* (2004).

### 3.2 Modeling the human hand

There is an inherent mismatch between the mechanical design and capabilities of robotic hands versus human hands. Why is this so? Robotics-based models of the human hand use idealized simple joints, torque motors and finger-pad elements. It has been shown that these simplifications do not suffice to create a realistic model of the thumb because they do not replicate the complex functional interactions among bones, ligaments and muscles (Valero-Cuevas *et al.*, 2003). If we truly desire to create hand-like functions in a robot, we need to learn from a working system such as the human hand.

Our current efforts are focused on constructing a biomechanically realistic human hand model that would allow us to determine what features of the human hand are the most important to be mimicked when designing a robotic hand. These beneficial features will be identified by creating several versions of the human hand model, each with different sets of features, and analyzing the ability of each hand to perform a set of grasping and manipulation tasks. These iterative refinements begin with developing a hand model, with rigid bodies modeling the geometry of the digits, that has a simplified version of the actual human kinematics. This will then be compared to a version of the hand that has links that deform in a manner similar to the fleshy portions of the human hand to evaluate how compliant surfaces aid stable grasping. Another version of the hand will have realistic human joints to determine the benefits of a compliant kinematic structure, and a fourth version will incorporate the network of tendons to determine what are the advantages, if any, of indirect actuation of the joints. Our simulation system will provide an arena for these comparisons, because of its ability to simulate dynamic grasping tasks and evaluate the completed grasps with a variety of quality metrics.

As a first step, we have created a human hand model that uses accurate kinematic parameters for the joint axes of the thumb. A high quality polygonal model of the skin's surface was obtained from the program Poser, which is used for animating 3D characters. For each of the fingers we use three

**Figure 3** Grasps in the training set



**Notes:** Left: For each superquadric in the training set we generate 1,600 grasp starting poses. These cover 16 positions over 1/8th of the total surface area, with 100 random combinations of different thumb orientations and finger spread angles. Here the long vector from each point denotes the grasp approach direction and the collection of short vectors shows the various thumb orientations. The spread angle is not shown. Middle: The optimal grasp of the SVM regression function for a superquadric shape previously seen in the training set ( $\epsilon_1 = 1$ ;  $\epsilon_2 = 0.3$ ). The quality of this grasp as tested in GraspIt! is 0.402. Right: The optimal grasp of the SVM regression function for novel a superquadric ( $\epsilon_1 = 0.6$ ;  $\epsilon_2 = 0.5$ ) The quality of this grasp as tested in GraspIt! is 0.315.


parallel flexion-extension (FE) joint axes and one abduction-adduction (AA) joint axis that intersects the first FE axis at a right angle. Initially, we used similar intersecting axes for the joints of the thumb (three FE axes and two orthogonal AA axes). However, while this simple orthogonal model has been considered in the past, such a model does not accurately predict maximal static thumbtip forces (Valero-Cuevas *et al.*, 2003).

Through the use of Markov Chain Monte Carlo simulations (Gilks *et al.*, 1996) and the virtual five-link thumb model (Giurintano *et al.*, 1995) we have arrived at a set of kinematic parameters that can best fit, in a least squares sense, experimental thumbtip forces and electromyography data (Santos and Valero-Cuevas, 2004). The difference in these two sets of joint axes can be seen in Figure 4.

Since these models rely on ideal revolute joints, they will not be able to truly mimic all of the possible motions of the digits. This limitation can be seen in Figure 5, where the hand grasps a mug. Initially the side of the thumb is in contact with the mug, but in a real grasp, the thumb would twist passively as pressure is increased at the contact. The current kinematic model prevents this from happening in simulation. What is needed is a model to predict the kinematics of a joint as a function of the shape of the bone articular surfaces in contact, the geometry and material properties of the ligaments and the applied load.


In addition, as the pressure increases at a contact, the finger surface deforms and begins to approach the local shape of the

**Figure 4** The difference in two set of joint axes



**Notes:** Left: The five joint axes of the thumb are arranged in a way commonly seen in mechanical designs with orthogonal and intersecting pairs of axes. Right: When using ideal revolute joints to model the kinematics of the hand, this set of skewed axes more accurately predicted experimental thumbtip forces and muscle coordination patterns

**Figure 5** The thumb contacts the mug on its ulnar side, but because of the constraints of using idealized joints, it will not twist and make contact with the thumbtip pad as pressure is increased at the contact



grasped object at the contact point. This results in an expanded contact area that reduces shear stresses for a given fingertip force and or torque magnitude, improves tactile information about the contact, and increases the safety margin to microslips, thus improving the stability of the grasp. Currently in *GraspIt!* all of the links are considered individual rigid bodies; however, we are beginning to develop a physically-based deformable finger model to predict the geometric change in the contact regions and accurately compute contact forces between nonrigid bodies during grasping. We are investigating the use of adaptive refinement finite element methods (Grinspun *et al.*, 2002) to compute the skin deformations in both contact and joint movement situations, and we plan to use an accurate model of the skeletal structure to constrain the positions of the innermost elements, which contact the bone. CT and MRI scans currently being collected will provide the necessary geometric detail to model the bones and soft tissue of the hand. We will also use published experimentally derived models of viscoelastic properties of the finger pads to accurately simulate these deformations (Pawluk and Howe, 1999; Serina *et al.*, 1998).


Finally, we are interested in modeling the indirect actuation of the tendon networks within each finger. Using an existing computer-controlled system for actuating multiple muscles in cadaveric hands (Plate 4), we can measure finger forces and torques or movement for given input tendon forces or excursions. Then with this data we will use model-based estimation techniques Bar-Shalom *et al.*, 2001 to determine the statistically best suited representation of the finger extensor mechanism (Figure 6). This will involve estimating the number and location of the connections (i.e. points of bifurcation or confluence) among the tendons.

We have used an idealized extensor mechanism description dating from the 17th century (Figure 7(a)) and have simulated its behavior as a 3D floating net. The distribution of tension among the different branches depends on their geometric configuration (Valero-Cuevas *et al.*, 1998) which is known to change with finger posture (Garcia-Elias *et al.*, 1991a) (Figure 7(b)). In contrast, previous descriptions of the extensor mechanism assumed a fixed distribution of tension among its elements. More recently, other studies have used this floating net approach to study finger movement (Sancho-Bru *et al.*, 2001) and thumb force production (Valero-Cuevas *et al.*, 2003). However, these models of the extensor

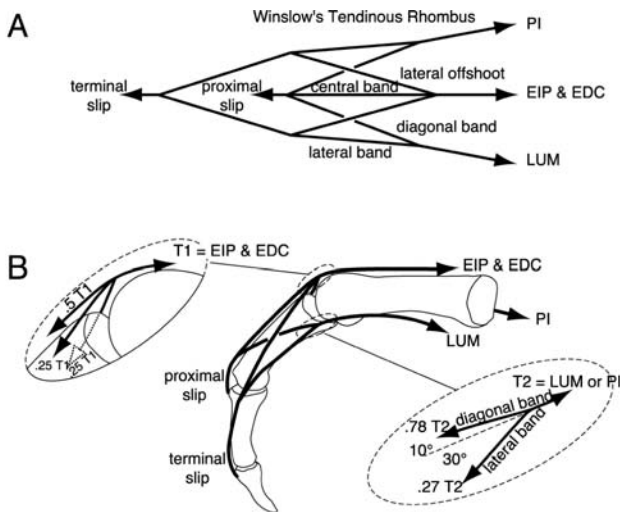
**Plate 4** A computer-controlled system to deliver known tensions and displacements to cadaveric hands (Pearlman *et al.*, 2004; Valero-Cuevas *et al.*, 2000)



**Figure 6** Stylized anatomy of the middle finger showing the tendinous network that makes up the finger extensor mechanism (Netter, 1997)



**Figure 7** Valero-Cuevas' model of the extensor mechanism of the index finger (Valero-Cuevas *et al.*, 1998)



mechanisms are not entirely validated for all finger postures (as would be desirable for a general-purpose model of the hand), nor is there a mechanics-based description of the extensor mechanism that can be efficiently encoded to be part of a computer simulator of the human hand.

#### 4. Conclusion

We have presented a versatile simulation system that can aid grasping research by facilitating the testing of grasp planning and grasping control algorithms without the need for costly robotic hands. Of course, if one has a manipulation system, this software can be used effectively as a planning environment to devise a grasping strategy for that hand (Kragić *et al.*, 2001). Synthesizing appropriate grasps for a given task is not a trivial problem, however, given the large number of possibilities, and we have demonstrated two different grasp planning systems that are inspired by the way humans approach the problem. The first attempts to limit this

search space by using taxonomies of prehensile grasping postures and rules on how to choose appropriate postures for given primitive shapes within the object model, and the second attempts to learn the relationship between grasp quality and the combination of object shape and grasping parameters.

*GraspIt!* can also be used to evaluate new robotic hand designs. Simulation is often used in the design of many other mechanical devices because it allows a designer to investigate how changes to various parameters affect the performance of the device without the need to build time-consuming and costly prototypes. Many current robotic hand designs attempt to emulate the human hand because of its proven manipulation capabilities (Wilkinson *et al.*, 2003). Still however, we do not fully understand all the intricacies of this complex biological mechanism, nor do we have a clear sense of which elements are most instrumental in making the hand such an effective manipulator. Thus we have embarked on a project to create a simulated biomechanical model of the human hand that we hope will give us insight into how to build a better robotic hand.

#### Note

- 1 The source code for *GraspIt!* is available for download from <http://www.cs.columbia.edu/~amiller/graspit>

#### References

Anitescu, M. and Potra, F.A. (1997), "Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems", *Nonlinear Dynamics*, Vol. 14, pp. 231-47.

Bar-Shalom, Y., Li, X. and Kirubarajan, T. (2001), *Estimation with Applications to Tracking and Navigation*, Wiley, New York, NY.

Borst, C., Fischer, M. and Hirzinger, G. (1999), "A fast and robust grasp planner for arbitrary 3D objects", *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, pp. 1890-6.

Corke, P. (1996), "A robotics toolbox for MATLAB", *IEEE Robotics and Automation Magazine*, Vol. 3 No. 1, pp. 24-32.

Cottle, R.W., Pang, J.S. and Stone, R.E. (1992), *The Linear Complementarity Problem*, Academic Press, p. 1992.

Ding, D., Liu, Y.-H. and Wang, S. (2000), "Computing 3-D optimal formclosure grasps", *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, pp. 3573-8.

Garcia-Elias, M., An, K.N., Berglund, L., Linscheid, R.L., Cooney, W.P. and Chao, E.Y. (1991a), "Extensor mechanism of the fingers: I A quantitative geometric study", *J. Hand Surgery (American)*, Vol. 16, pp. 1130-40.

Gilks, W.R., Richardson, S. and Spiegelhalter, D.J. (1996), *Markov Chain Monte Carlo in Practice*, Chapman & Hall/CRC, Boca Raton, FL.

Giurintano, D.J., Hollister, A.M., Buford, W.L., Thompson, D.E. and Myers, L.M. (1995), "A virtual five-link model of the thumb", *Medical Engineering and Physics*, Vol. 17 No. 4, pp. 297-303.

Grinspun, E., Krysl, P. and Schröder, P. (2002), "CHARMS: a simple framework for adaptive simulation", *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, ACM Press, pp. 281-90.



- Hester, R.D., Cetin, M., Kapoor, C. and Tesar, D. (1999), "A criteria-based approach to grasp synthesis", *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, pp. 1255-60.
- Kragić, D., Miller, A. and Allen, P. (2001), "Real-time tracking meets online grasp planning", *Proceedings of the 2001 IEEE International Conference on Robotics and Automation*, pp. 2460-5.
- Larsen, E., Gottschalk, S., Lin, M. and Manocha, D. (1999), "Fast proximity queries with swept sphere volumes", Technical Report TR99-018, Department of Computer Science, University of North Carolina, Chapel Hill, NC.
- Miller, A. and Allen, P. (1999), "Examples of 3D grasp quality computations", *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, pp. 1240-6.
- Miller, A. and Allen, P. (2000), "GraspIt!: A versatile simulator for grasping analysis", *Proceedings of the ASME Dynamic Systems and Control Division*, Vol. 2, pp. 1251-8.
- Miller, A. and Christensen, H. (2003), "Implementation of multi-rigid-body dynamics within a robotic grasping simulator", *Proceedings of the 2003 IEEE International Conference on Robotics and Automation*, pp. 2262-8.
- Miller, A., Knoop, S., Christensen, H. and Allen, P. (2003), "Automatic grasp planning using shape primitives", *Proceedings of the 2003 IEEE International Conference on Robotics and Automation*, pp. 1824-9.
- Mirtich, B. and Canny, J. (1994), "Easily computable optimum grasps in 2-D and 3-D", *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pp. 739-47.
- Napier, J. (1956), "The prehensile movements of the human hand", *J. Bone and Joint Surgery*, Vol. 38B No. 4, pp. 902-13.
- Netter, F.H. (1997), *Atlas of Human Anatomy*, 2nd edn., Icon Learning Systems.
- Pawluk, D. and Howe, R. (1999), "Dynamic lumped element response of the human fingerpad", *J. Biomech. Eng.*, Vol. 121 No. 2, pp. 178-83.
- Pearlman, J.L., Roach, S.S. and Valero-Cuevas, F.J. (2004), "The fundamental thumb-tip force vectors produced by the muscles of the thumb", *J. Orthopaedic Research*, Vol. 22 No. 2, pp. 306-12.
- Pelossos, R., Miller, A., Allen, P. and Jebara, T. (2004), "An SVM learning approach to robotic grasping", *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, pp. 3215-8.
- Ponce, J., Sullivan, S., Boissonnat, J.-D. and Merlet, J.-P. (1993), "On characterizing and computing three- and four-finger force-closure grasps of polyhedral objects", *Proceedings of the 1993 IEEE International Conference on Robotics and Automation*, pp. 821-7.
- Ruspini, D. and Khatib, O. (1999), "Collision/contact models for the dynamic simulation and haptic interaction", *Proceedings of the Ninth International Symposium of Robotics Research, ISRR'99*, pp. 185-94.
- Sancho-Bru, J.L., Perez-Gonzalez, A., Vergara-Monedero, M. and Giurintano, D. (2001), "A 3-d dynamic model of human finger for studying free movements", *J. Biomech.*, Vol. 34, pp. 1491-500.
- Santos, V. and Valero-Cuevas, F. (2004), "Investigating the interaction between variability in both musculoskeletal structure and muscle coordination for maximal voluntary static thumb forces", *Proceedings of the Neural Control of Movement Annual Meeting*, pp. 27-8.
- Serina, E.R., Mockensturm, E., Mote, C.D. and Rempel, D. (1998), "A structural model of the forced compression of the fingertip pulp", *J. Biomech.*, Vol. 31 No. 7, pp. 639-46.
- Speck, A. and Klaeren, H. (1999), "RoboSiM: Java 3D robot visualization", *IECON '99*, pp. 821-6.
- Valero-Cuevas, F., Johanson, M.E. and Towles, J.D. (2003), "Towards a realistic biomechanical model of the thumb: The choice of kinematic description is more critical than the solution method or the variability/uncertainty of musculoskeletal parameters", *J. Biomech.*, Vol. 36 No. 7, pp. 1019-30.
- Valero-Cuevas, F.J., Towles, J.D. and Hentz, V.R. (2000), "Quantification of fingertip force reduction in the forefinger following simulated paralysis of extensor and intrinsic muscles", *J. Biomech.*, Vol. 33, pp. 1601-9.
- Valero-Cuevas, F.J., Zajac, F.E. and Burgar, C.G. (1998), "Large index-fingertip forces are produced by subject-independent patterns of muscle excitation", *J. Biomech.*, Vol. 31, pp. 693-703.
- Wilkinson, D.D., Weghe, M.V. and Matsuoka, Y. (2003), "An extensor mechanism for an anatomical robotic hand", *Proceedings of the 2003 IEEE International Conference on Robotics and Automation*, pp. 238-43.