

Interaction Capture and Synthesis

Paul G. Kry¹²³ Dinesh K. Pai¹²

¹University of British Columbia ²Rutgers University ³EVASION / INRIA




Figure 1: Interaction capture explained. (a) Grasp on a small box captured using both motion capture markers and fingertip mounted force sensors; (b) measured contact forces; (c) interaction transferred to a similar box; (d) interaction transferred to a different shape.

Abstract

Modifying motion capture to satisfy the constraints of new animation is difficult when contact is involved, and a critical problem for animation of hands. The compliance with which a character makes contact also reveals important aspects of the movement’s purpose. We present a new technique called *interaction capture*, for capturing these contact phenomena. We capture contact forces at the same time as motion, at a high rate, and use both to estimate a nominal reference trajectory and joint compliance. Unlike traditional methods, our method estimates joint compliance without the need for motorized perturbation devices. New interactions can then be synthesized by physically based simulation. We describe a novel position-based linear complementarity problem formulation that includes friction, breaking contact, and the compliant coupling between contacts at different fingers. The technique is validated using data from previous work and our own perturbation-based estimates.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: motion capture, forces, grasping, physical modeling, contact simulation, LCP

1 Introduction

Motion capture and resynthesis has recently become an essential part of computer animation because it provides a direct and straightforward way to capture the subtleties of human motion. Even though some animations require the fine control of an artist or the flexibility provided by algorithmic motion controllers, motion capture provides the excellent combination of realism and simplicity that is well suited for a large and growing fraction of character animations today.

One of the major challenges of motion capture, however, is the modification of motion to satisfy the constraints of a target animation [Hodgins and Pollard 1997; Gleicher 1998; Popovic and Witkin 1999; Zordan and Hodgins 2002; Liu et al. 2005; Zordan et al. 2005]. Motion retargeting is relatively straightforward for free motions, but becomes very difficult when contact is involved. Contact changes the motion of both the character and the environment due to the physical interaction between the two. This is particularly crucial for animating human hands which evolved primarily for grasping and manipulation; in other words, for contact. Hands are also challenging due to their highly articulated kinematics and the numerous contacts they exhibit during interaction. We use hands to illustrate the important issues in this paper. However, similar issues also occur in many scenarios, for example, walking on uneven terrain or playing a contact sport.

From our perspective, the problem with motion capture is that it only captures the motion and not how humans move. There is considerable evidence that during physical contact humans do not control just position but also the effective compliance¹ of the contact through co-activation of antagonistic muscles [Feldman 1986; Hogan 1984; Bizzi et al. 1992]. The compliance is highly dependent on task, intent, and object geometry; it is also important for stability [Burdet et al. 2001; Rancourt and Hogan 2001].

In this paper we introduce a new technique called *Interaction Capture* which extends motion capture to explicitly include the capture of contact phenomena. We capture both motion and contact forces, and estimate an intermediate representation of the movement that we call the *interaction trajectory*. It describes both the intended motion (a reference joint angle trajectory) and passive response to contact (compliance of the joints). From this representation we can synthesize new interactions with similar objects but with different shapes, sizes, and frictional properties using physically based simulation. Figure 1 illustrates the process.

Measuring the compliance during movement can be difficult, and has traditionally required the use of complex motorized equipment for perturbing motion. We show how this difficulty can be avoided by exploiting the long latency of human reflexes, and careful monitoring of force and motion at the time of contact. This leads to a simple and practical capture method, only slightly more complicated than motion capture.

¹Compliance is the inverse of stiffness. See Section 2 for more details.

1.1 Our Contributions

- The technique of Interaction Capture, which simplifies the capture and retargeting of movements involving contact. It measures both motion and contact force and estimates *both* motion trajectories *and* joint compliances.
- A new method for estimating joint compliances at the time of contact. It is the first use of force measurements for joint compliance estimation in graphics, and it is significantly easier to use than methods developed in other fields.
- A new algorithm for retargeting movements by quasi-static simulation of compliant multipoint frictional contact between an articulated body and its dynamic environment.
- An Internet repository of interaction capture data with synchronized forces and motions captured at 500Hz, available to the community at <http://www.interactioncapture.org>. We also provide guidelines for readers to perform their own interaction captures.

2 Related Work

We review some of the related work in this section; others will be discussed in context in the rest of the paper.

Compliant joints have been used in computer animation for capturing emotion or style [Neff and Fiume 2002; Liu et al. 2005], as part of controllers for synthesizing motion [Popovic and Witkin 1999; Shapiro et al. 2003; Pollard and Zordan 2005], and for reacting to impacts [Zordan and Hodgins 2002; Yin et al. 2003; Zordan et al. 2005]. However, the compliance (or stiffness) parameters in previous work are either selected by hand, or approximated through complex optimizations that must simultaneously deal with estimating contact forces. Interaction capture can instead provide compliance estimates from capture data.

The compliance (i.e., inverse stiffness) of human movement has been investigated in the fields of haptics, biomechanics, and human motor control. In general, damping also varies systematically with stiffness [Hajian and Howe 1997]. For linear systems, the two can be combined with inertia into a complex *impedance*. Hogan [1984] and others have suggested that, for reasons of causality, the brain controls impedance rather than its inverse, *admittance*. For simulation, however, the distinction is not important; in fact, admittance is better behaved since it remains finite in the rigid limit. In this paper, we will focus on the compliance part of admittance, because damping is correlated with stiffness, and inertia can be estimated from geometry and other means (for instance, as described in Appendix A of [Liu et al. 2005]).

There exists a wealth of previous work on measuring compliance, for example, characterization of the arm during movement [Gomi and Kawato 1997; Xu and Hollerbach 1999], fingertip models [Hajian and Howe 1997; Milner and Franklin 1998; Hasser and Cutkosky 2002], and a wrist model for improved haptic interaction [Kuchenbecker et al. 2003]. Most previous work on measuring compliance uses some form of perturbation device. That is, while performing an action, a subject has an impulsive force applied to his finger, hand, or arm. System parameters are then estimated by observing the response. This is not ideal since perturbation complicates the capture process and changes the motion (i.e., the motion we are trying to measure).

Compliant joints and effective endpoint models (e.g., fingertips or hands) encapsulate the characteristics of tendons, muscles, and activation. A good alternative approach would be physically based musculoskeletal models [Delp and Loan 2000; Tsang et al. 2005; Teran et al. 2003], but these require a more detailed knowledge of human anatomy, biomechanics, and motor control. Our approach, in contrast, uses a kinematic structure that matches the widely used motion capture model, with the addition of compliances at the




Figure 2: One finger interaction capture.

joints. We then fit the joint compliances to interaction capture measurements to get realistic values, and let the captured motion and force guide the interaction trajectory.

We focus on hands because they are our primary tools for interacting with the world around us. Animation of hands has been important since the early years of computer graphics [Catmull 1972]. Though they have typically seen little attention in comparison to face modeling or body motion, hands have recently started attracting much more attention, as witnessed by the growing literature [ElKoura and Singh 2003; Kurihara and Miyata 2004; Tsang et al. 2005; Pollard and Zordan 2005].

3 Overview of Interaction Capture

In this section we describe the entire process of interaction capture and synthesis in a simple but real setting: an index finger gently strokes a flat surface, moving in the vertical plane (Figure 2). Our goal is to capture this interaction, and to resynthesize it in different situations, for instance, when the surface has a bump (Figure 4), or when the surface friction changes (see video). In this setting, we can describe the basic ideas without immediately delving into the gory, but unavoidable, details of articulated body kinematics and compliant frictional contact in 3D. In the rest of the paper, we show how this fully generalizes to more complex cases, using the whole human hand as an example.

3.1 Compliant Articulated Structure

The finger is modeled as a kinematic chain of three hinge (revolute) joints in the obvious way, with joint angles collected in the vector θ . The compliance C can be thought of as a collection of torsional springs that, when displaced from a reference configuration θ_r , produce joint torques τ by the relation

$$\theta = \theta_r + C\tau. \quad (1)$$

The contact force f at a point on the fingertip produces torques τ at each of the joints. They are related by the equation

$$\tau = J^T f, \quad (2)$$

where J is the Jacobian matrix, which depends on the contact point and finger configuration. How we compute this in the general case is described in Section 6.1.

3.2 The Capture Session and Data

We capture the interaction using two types of sensors. (1) The motion is measured with markers attached to the finger using a Vicon motion capture system. This part is standard, except that we capture at 500 Hz, for reasons explained below. (2) The contact force is measured using a six-axis force torque sensor, attached below the surface. The motion and force measurements are synchronized using the timing pulse provided by the Vicon data station.




Figure 3: Interaction capture for index MP flexion at the time of contact. Linear approximations of the pre- and post-contact trajectories of this contact lead to a compliance estimate of 1.04 rad/Nm.

Force sensors are readily available from several companies; we used sensors from ATI Industrial Automation and Zebra robotics (see <http://www.interactioncapture.org>). Forces can be measured by instrumenting the environment in this way in many cases, but see Section 5 for alternatives.

The motion data is processed in the usual way to estimate the joint angle trajectory. In addition, knowing the joint angles, we can compute the Jacobian matrix, and estimate joint torques τ from the contact forces using Equation 2.

Figure 3 shows plots for the index MP (metacarpo-phalangeal, or knuckle) joint angle and torque at the time of contact. Notice the clear distinction between pre- and post-contact trajectories, and the smooth, almost linear trajectories.

3.3 Joint Compliance Estimation

Traditional methods for measuring compliance use perturbation to avoid the confounding effects of spinal reflexes. The main observation is that spinal reflexes are actually quite slow due to neural conduction times, synaptic delays, and the excitation-contraction coupling of muscles. The result is that it can take around 100 ms to observe an appreciable change in muscle force due to a spinal reflex, and longer for a voluntary change [Johansson 1996].

As mentioned previously, measuring compliance has been difficult because traditional methods use extra equipment for producing perturbation. Our key insight is that natural contact itself provides such a perturbation. Therefore, by observing force changes in a small window of time immediately following contact, we can catch a glimpse at the compliance that was ‘pre-programmed’ by the central nervous system for the specific contact interaction. This is the reason for capturing forces and motion at 500 Hz, which gives us up to 50 useful samples in the time window.

Identification of the exact moment of contact is straightforward given the force measurements (see Figure 3). It would be much more complicated if only motion were captured [Ikemoto et al. 2006]. This is another immediate benefit of capturing forces.

At each contact event, we estimate the compliance from low degree polynomials fitted to the joint angle trajectory in small windows just before and just after the time of contact. These polynomials (linear or quadratic depending on data quality) serve as a smoothed version of the data from which we compute the compliance by dividing the derivative of the joint angle difference, by the derivative of the torque, at the time of contact.

In Figure 3, the compliance is approximately equivalent to a 10 degree rotation at the joint for each Newton of force at the fingertip. Note that this value depends on the task the subject wants to per-




Figure 4: Interaction resynthesis examples.

form; in Section 4 we compare the compliances for different tasks and validate our compliance estimates using independent measurements.

3.4 Reference Trajectory

The computed compliances, along with current configuration and joint torques give us a means of computing a reference configuration θ_r during contact. It can be computed simply from Equation 1. Conceptually, this is the configuration in which the finger would be if there had been no contact. The reference trajectory, along with the associated compliance, allows us to generalize the captured interaction to situations where the contacting surface is changed. In Figure 3, the reference trajectory approximately follows the extrapolated pre-contact trajectory.

Our reference trajectory is related to equilibrium-point trajectories proposed for human motor control [Feldman 1986; Bizzi et al. 1992] but very different. Our trajectory parameterizes the actual observed dynamic motion that would have been performed in the absence of contact, whereas an equilibrium-point trajectory is comprised of target points at which the system would be in equilibrium.

3.5 Interaction Synthesis

The compliance C and the reference trajectory θ_r are together called the *interaction trajectory*. It allows us to resynthesize new interactions using physical simulation. The interaction trajectory could be used with a dynamic simulation of the character, as in, for example, [Zordan and Hodgins 2002; Yin et al. 2003; Liu et al. 2005], if a suitable controller can be designed. However, the reference trajectory already encodes a lot of information about the character’s dynamics, and can also be retargeted using any previously developed motion retargeting method. We therefore focus on synthesizing the perturbations from this trajectory due to contact. For this, a quasi-static simulation is sufficient and can efficiently simulate important effects such as sticking and sliding due to static friction, and changes in contact location on the finger. This leads to a new type of linear complementarity problem, which we formulate and solve in Section 7.

Figure 4 shows the result of retargeting an interaction trajectory to a similar flat surface and a surface with a bump. The finger complies realistically to the bump on the surface. The motions are easier to see in the accompanying video, which also shows retargeting to surfaces with different friction and shape.

Note that the environment around the character can be fully dynamic. Even though this simple example does not show it, it is evident in the rest of the paper and the video. A significant benefit of our approach is that the environment simulation can view the character as a time-varying object with passive compliance.

To summarize, the synthesized motion exhibits realistic dynamic behavior, in which the character’s dynamics come from the capture subject (as desired), the environment dynamics come from a dynamic simulation, and the captured compliance provides a biologically plausible interface between the two.




Figure 5: Surface exploration (left) and surface scratching (right).

Joint	Exploring	Scratching
Wrist	0.3488	0.1531
MP	0.9882	0.4788
PIP	0.6715	< 0.8426
DIP	2.0052	< 0.8432

Table 1: Compliance estimates for exploring and scratching. Units are in rad/Nm and the MP compliance is for flexion-extension.

4 Validation

Measurements of single finger interactions help validate our estimation procedures. Figure 5 shows snapshots of surface exploration and surface scratching interactions. Using a surface instrumented with a force sensor, we compute joint torques using contact force measurements that include both normal and friction forces.

4.1 Task Dependence

Table 1 shows the compliances estimated for surface exploration in comparison to those estimated for surface scratching. The compliances are clearly dependent on the task, with joints tending to be much stiffer during scratching (by about a factor of two).

4.2 Estimates from Perturbation

We have validated our compliance estimates for the surface scratching and exploring tasks with perturbation data collected using a spring-loaded platform. When the platform is unlatched (by pulling a pin via a thread, see Figure 6), the resulting sudden impulse measures the passive response of the subject, before he reflexively removes his finger from the platform. Figure 7 shows the fingertip elevation and vertical force during the perturbation trials; notice the reaction time is just under 100 ms.

In world coordinates, the platform generates a perturbation force in the vertical, or positive z , direction. From the measured z position we compute velocities and accelerations by numerical differentiation. We then fit parameters for a second order system, $f(t) = m\ddot{z}(t) + b\dot{z}(t) + kz(t) + f_0$, where f_0 is the rest force at the time of perturbation. We use least squares to find best-fit parameters using small windows of data that start at the time of perturbation and focus on the excitation of the system. This window is larger in the less stiff surface exploration case (20 ms for scratching, 25 ms for exploring since it takes longer to come to rest). Table 2 shows these estimates, but also includes the vertical component of the effective endpoint stiffness (denoted G , see Section 6.2) computed using the posture at the time of perturbation and our estimated joint compliances in Table 1. The stiffness estimated via perturbation matches surprisingly well with our joint compliance estimates.

Table 2 also compares our estimates to those of Hajian and Howe [1997], Milner and Franklin [1998], and Kuchenbecker et al. [2003]. Our mass estimates are justifiably higher than Hajian’s estimate since we are not measuring the fingertip alone, but surprisingly lower than that of Kuchenbecker’s estimated wrist model.




Figure 6: Video frames at 30 Hz showing perturbation. The pin unlatched the platform in the second frame, and by the third frame the perturbation is complete. In the fifth frame the subject starts to quickly remove his finger.




Figure 7: Fingertip perturbation data showing reaction time of about 100 ms. Fingertip elevation and force are shown for surface exploration (left) and surface scratching (right).

Parameter	Exploring	Scratching	HH97	MF98	KPN03
f (N)	0.677	1.345	2	(relaxed)	9 (grip)
m (kg)	0.017	0.018	0.006		0.142
b (Ns/m)	2.04	3.86	2		5.03
k (N/m)	89.16	257.04	137	283	538
k_G (N/m)	81.86	285.96			

Table 2: Estimated second order system parameters, and the effective endpoint stiffness in the vertical direction, k_G , computed from joint compliance estimates.




Figure 8: Extrapolation is unlikely to yield a positive compliance estimate that produces the reference trajectory shown for 1 rad/Nm.

This may be due to the size of the grip force in the latter, which also explains the larger stiffness.

4.3 Limitations

As with regular motion capture, compliance estimation can occasionally run into problems, and needs to be monitored to ensure data quality. We find the initial post contact equilibrium trajectory through extrapolation, but in some cases, joint trajectory extrapolation would suggest the compliance is negative (see Figure 8). In these cases we discard the estimate. During scratching in particular, there were few contact events from which we could compute reasonable compliance estimates for the proximal and distal interphalangeal joints (DIP and PIP).

There are several possible explanations for these difficulties. Trends in small joint motions can be hidden by motion capture noise; high joint stiffness can exacerbate this effect. Estimates can be difficult if the kinematic structure is close to singularity in the di-

rections of force measurement. This can happen when the joints are nearly straight (as in Figure 8), which can also cause inaccuracies of joint trajectory estimation from marker data. As such, we expect the actual DIP and PIP scratching values to be lower than those in Table 1. It is also plausible, however, that low degree polynomials are insufficient for trajectory interpolation, for instance, when there exists complex ‘pre-programmed’ motion around the time of contact. We believe that many of these difficult estimation situations can be avoided through a combination of careful kinematic calibration, additional motion capture markers, and smoother kinematic motion preprocessing.

Lastly, note that Equation 2 ignores the influence of inertia and gravity. This is a reasonable simplification for the finger joints in our current measurements because the motions are relatively slow and the fingers have little mass. For other joints (e.g., wrists, elbows, or knees) we could add a correction to Equation 2 using an estimated inertia and derivatives of joint angles from motion capture, though this could be tricky. Just the same, for typical graphics applications we can perhaps disregard this correction (e.g., for the slow wrist movements during scratching and exploring).

5 Force Measurement Methods

We have investigated a number of different options for measuring interaction forces. Six-axis force torque sensors provide good accuracy and temporal resolution. Objects can either be mounted on or instrumented with these sensors (see Figure 5). Other options include single axis pressure sensors such as force sensitive resistors (FSRs), or capacitance based sensors (for instance, the pressure sensitive sphere shown in Figure 13(f)). Pressure sensors are available commercially from several vendors (e.g., Pressure Profile Systems). For measuring fingertip forces on arbitrary objects, we have constructed small fingertip mounted sensors with force sensitive resistors (see Figure 9). These sensors are inexpensive, easy to measure using standard data acquisition equipment, and simple to calibrate using a force torque sensor. Details on how to build these sensors and a list of suppliers can be found at <http://www.interactioncapture.org>.




Figure 9: FSR mounted on a fingertip.

6 Interaction Capture in 3D

In this section, we describe how to extend the results of the previous sections to 3D articulated structures with branches and compliant joints, like the human hand, and set the stage for 3D interaction synthesis in Section 7.

6.1 Kinematics and Dynamics in 3D

The homogeneous coordinates of frame i with respect to another frame j is given by the 4×4 matrix jE_i . We use leading subscripts and superscripts to indicate frames. The homogeneous coordinates of a three dimensional vector x in frame i are denoted ${}^i x$. The homogeneous coordinates of this vector in frame j are given by left multiplying by jE_i .

The spatial velocity ϕ describes the relative motion of a body with respect to the fixed world frame. In coordinates of frame i , it is given by the size 6 column vector ${}^i\phi = ({}^i\omega^T, {}^i v^T)^T$, where ω is the angular velocity and v is the linear velocity of the point at the origin of frame i . Spatial forces, called wrenches, are represented

as $w = (\tau^T, f^T)^T$, where τ is the (rotational) torque and f is the (translational) force.

Spatial velocities and wrenches transform according to the adjoint transformation jAd_i . The 6×6 adjoint matrix is defined as,

$${}^jAd_i = \begin{pmatrix} \Theta & 0 \\ [p]\Theta & \Theta \end{pmatrix}, \quad \text{where } {}^jE_i = \begin{pmatrix} \Theta & p \\ 0 & 1 \end{pmatrix}.$$

Here, Θ is a 3×3 rotation matrix, p is a 3×1 displacement, and $[p]$ denotes the skew symmetric 3×3 matrix equivalent to the cross product $p \times$. Spatial velocities, being contravariant quantities, transform by left multiplying, ${}^j\phi = {}^jAd_i {}^i\phi$. Because we write spatial wrenches as column vectors, these covariant quantities are transformed by left multiplying with the inverse transpose, ${}^jw = {}^jAd_i^T {}^iw$.

We also define $\Gamma(r)$, a 3×6 matrix for computing the linear velocity of a point r given a spatial velocity. That is, $\dot{r} = \Gamma(r)\phi$, where

$$\Gamma(r) = \begin{pmatrix} [-r] & I \end{pmatrix},$$

but similarly $w = \Gamma(r)^T f$ for a linear force f acting at r .

The spatial cross product of $\phi = (\omega^T, v^T)^T$ is the linear operator with coordinate matrix

$$[\phi] = \begin{pmatrix} [\omega] & 0 \\ [v] & [\omega] \end{pmatrix}.$$

With this, the Newton-Euler equation for a rigid body can be written in body coordinates as

$$w = M\dot{\phi} - [\phi]^T M\phi, \quad (3)$$

where M is the mass inertia matrix, and all quantities are in body coordinates.

Lastly, for an articulated structure consisting of rigid components linked with joints, the Jacobian provides a linear transformation from joint angle velocities to the spatial velocity of a given link. For convenience, since our kinematic structure is a tree rather than a chain, we write the Jacobian as a matrix J , where the columns consist of *all* the twists in *world coordinates* for all degrees of freedom (i.e., each column contains the spatial velocity due to an angular velocity of 1 rad/s at the joint). Since the spatial velocity of body i is the sum of only the twists on the path between body i and the root, we select the joints affecting this body with a diagonal binary matrix S_i . Thus, $JS_i\theta$ is the spatial velocity of body i . Further, the velocity of a point r_i on body i (e.g., a contact point) is simply $\dot{r}_i = J_i\dot{\theta}$ where

$$J_i = \Gamma(r_i)JS_i. \quad (4)$$

Likewise, the Jacobian transpose maps wrenches to joint torques. For a linear force f_i at a contact point r_i , we compute $\tau = J_i^T f_i$. Note that for a small joint displacement, J_i provides a linear approximation to the displacement of r_i , that is, $\Delta u_i = J_i\Delta\theta$.

6.2 Effective Endpoint Compliance

The compliant articulated structure acts like a generalized spring, where internal torques due to forces at contact points and the angular displacements are always in equilibrium. The compliant behavior of a point on a body in the articulated structure is described by the effective compliance and is given by a 3×3 matrix. A small force Δf_i , when applied to the compliant kinematic structure at a contact point r_i , results in a displacement Δu_i at the contact point. The force at point r_i maps to internal torques via J_i^T . The torques then map to joint displacements via the compliance matrix (Equation 1), and finally, the joint displacements map to a displacement at the contact point via J_i . We write this effective compliance relationship as $\Delta u_i = J_i C J_i^T \Delta f_i$. The matrix $J_i C J_i^T$ need not be invertible,




Figure 10: The effective end point compliance for multiple contacts is coupled because of shared joints, for example, the wrist.

for example, if the joints are lined up in a singular configuration. However, we ensure the matrix has full rank by giving the skeleton’s root non-zero linear compliance in all directions; adding a small amount of compliance at the root is similar to Tikhonov regularization.

With N contacts on multiple different links in the kinematic structure (see Figure 10), the effective compliance becomes coupled because of shared joints on the path between the contacts and the root. This relationship is described by the matrix,

$$G = J C J^T, \quad (5)$$

where J^T is the block row matrix (J_1^T, \dots, J_N^T) . Grouping linear forces and displacements into block column vectors, Δf and Δu respectively, the local linear model of the compliance for the current equilibrium configuration is written as

$$G \Delta f = \Delta u. \quad (6)$$

7 Interaction Synthesis

We now describe how the captured interaction trajectory can be used to interact with objects using dynamics simulation. This requires computing the contact forces that the character would impose on a rigid object, accounting for friction and intermittent contact. This is a hard problem in general, because of the pose-dependent compliant coupling between different fingers. To address this, we formulate a new linear complementarity problem (LCP) which is related to, but significantly extends, the quasi-rigid contact model of Pauly et al. [2004]. Unlike previous work, we incorporate friction into a quasi-rigid LCP, and account for contact with a moving compliant articulated structure using a locally linear approximation. The result is a novel formulation that handles multiple contacts, breaking contact, and a quasi-static compliant articulated structure.

7.1 Algorithm

The state of the system at time k is given by C^k , θ_r^k , ϕ^k , q^k , θ^k , f_i^k , r_i^k , $i = 1 \dots N$. The values C^k and θ^k come from the captured interaction, ϕ^k and q^k describe the spatial velocity and configuration of the environment, f_i^k are the forces, and r_i^k are the contact locations. The quasi-static compliant structure is *approximately* maintained in equilibrium, i.e.,

$$\theta^k - \theta_r^k + \theta_{err} = \sum_i C^k J_i^k T f_i^k \quad (7)$$

will hold for a small error in joint angles θ_{err} , where J_i is the Jacobian for contact r_i and configuration θ^k (see Equation 4). This small error exists because of updates in the contact location and changes in the configuration (computed via locally linear approximation).

Ultimately, we want to find the forces f_i^{k+1} that evolve the state of the rigid body forward (providing ϕ^{k+1} and q^{k+1} upon integration in time) and let us find the θ^{k+1} necessary to keep the compliant structure in equilibrium given new reference angles θ_r^{k+1} .

7.1.1 Friction and Breaking Contact

We first define vectors at each contact point that let us describe the friction cone. This construction is similar to those used in complementarity formulations of rigid body friction (see, for example, [Lotstedt 1981; Baraff 1994; Stewart and Trinkle 1996; Anitescu and Potra 2002; Miller and Christensen 2003]). We define d_{i0} as the outward unit normal at contact point i and a set of paired unit tangent vectors d_{ij} , $j = 1..2m$, such that $\|d_{ij}\| = 1$, $d_{i(2j-1)} = -d_{i(2j)}$, $j = 1..m$ (see Figure 11). We assemble these vectors into a matrix $D_i = (d_{i0}, \dots, d_{i(2m)})$, as well as a matrix consisting of just the tangent vectors, $D_{i*} = (d_{i1}, \dots, d_{i(2m)})$.

The Coulomb friction cone is defined as the set of possible forces that can be supported by the frictional surface. We build a polyhedral approximation to the cone using the vectors in D_i . Let β_{i0} be the normal force magnitude at contact i , that is, $\beta_{i0} = d_{i0} f_i$. Given a coefficient of friction μ_i , the tangential friction force allowed by Coulomb friction lies inside a circle in the tangent plane of radius $\mu_i \beta_{i0}$. We approximate this set of allowable friction forces by the convex hull of the unit length tangent vectors, d_{ij} $j = 1..2m$, scaled by the normal force times the coefficient of friction,

$$\left\{ \sum_{j=1}^{2m} d_{ij} \beta_{ij} \mid \beta_{ij} \geq 0, \sum_{j=1}^{2m} \beta_{ij} \leq \mu_i \beta_{i0} \right\}.$$

Therefore, we approximate Coulomb friction at contact point i by $f_i = D_i \beta_i$, $\beta_{i*} \geq 0$, $F_i \beta_i \geq 0$, where $\beta_i = (\beta_{i0}, \dots, \beta_{i(2m)})^T$, and $F_i = (\mu_i, -1, \dots, -1)$. Note that we always use the \star subscript to denote “all but the first element”. The force at contact i , written $D_i \beta_i$, is actually f_i^{k+1} , the force at the next step for which we are solving. The Coulomb friction approximation gives us a constraint on the force, but we must also constrain the motion to satisfy the principle of maximum dissipation. The actual force maximizes the instantaneous energy dissipation due to frictional work, i.e., a minimization of the *negative* frictional work,

$$\min_{\beta_{i*}} \Delta u_i^T D_{i*} \beta_{i*}, \quad \beta_{i*} \geq 0, \quad F_i \beta_i \geq 0.$$

Here, $D_{i*} \beta_{i*}$ is the friction force, and Δu_i is the motion at the contact point (a fixed quantity for the purpose of minimization). We first rewrite this constrained minimization problem using Lagrange multipliers (see [Murty 1988]). The Lagrangian for this system is

$$\mathcal{L} = \Delta u_i^T D_{i*} \beta_{i*} - v_{i*}^T \beta_{i*} - \lambda_i F_i \beta_i,$$

with the KKT optimality condition coming from differentiation with respect to β_{i*} . We write it as $v_{i*} = D_{i*}^T \Delta u_i - E_{i*}^T \lambda_i$, where $E_i = (0, -1, \dots, -1)$. Defining $\sigma_i = F_i \beta_i$, we can then write the KKT conditions with complementarity notation ($a \perp b$ meaning $a \geq 0, b \geq 0, a^T b = 0$) as

$$\begin{aligned} \beta_{i*} \perp v_{i*} &= D_{i*}^T \Delta u_i - E_{i*}^T \lambda_i, \\ \lambda_i \perp \sigma_i &= F_i \beta_i. \end{aligned} \quad (8)$$

Addressing non-interpenetration, we define s_{i0} as the current linear distance between closest points (or contact points) along the contact normal. This distance is provided by proximity detection and will be negative when there is interpenetration (in which case the contact points are selected as extremal points). We define the separation, v_{i0} , of contact point i as

$$v_{i0} = d_{i0}^T \Delta u_i + s_{i0}.$$




Figure 11: Discretized friction cone for $m = 3$.




Figure 12: Left, contact motion due to changing forces at the contacts as specified by the effective endpoint compliance G . Middle, contact motion due to a small change in the reference trajectory, $\Delta\theta_r$. Right, contact motion due to rigid body motion.

For non-penetration, v_{i0} must be positive, but if we want “non-velcro” forces then we also require the magnitude of the normal force to be non-negative, $\beta_{i0} = d_{i0}^T f_i \geq 0$. These values are complementary to each other as we can only have contact forces when the separation is zero. We combine this with Equation 8 by defining $v_i = (v_{i0}, v_{i\kappa}^T)^T$, and $s_i = (s_{i0}, 0, \dots, 0)^T$, which lets us write the complementarity conditions at point i as

$$\begin{aligned} \beta_i \perp v_i &= D_i^T \Delta u_i - E_i^T \lambda_i + s_i, \\ \lambda_i \perp \sigma_i &= F_i \beta_i. \end{aligned} \quad (9)$$

7.1.2 Contact Point Motion

We define Δu_i in an implicit fashion, combining the different sources of motion for contact points. The block vector Δu has three terms (see Figure 12),

$$\Delta u = G\Delta f + J\Delta\theta_r - (U_\phi f^{k+1} + u_\phi). \quad (10)$$

The first term accounts for movement of the contacts due to the change in contact forces, $\Delta f = f^{k+1} - f^k$. This is the motion necessary to maintain balance in the equilibrium equation for the new contact forces; the matrix G gives us the locally linear approximation of this motion for small changes in force.

The second term of Equation 10 is also related to maintaining equilibrium, as it gives the motion of the contact points necessary to preserve equilibrium given the current forces and the change in the reference angles, $\Delta\theta_r = \theta_r^{k+1} - \theta_r^k$.

The third term of Equation 10 is the motion of the contact point on the rigid body due to rigid body motion. This motion has two parts. The first part, $U_\phi f^{k+1}$, describes the motion due to contact forces, while the second part, u_ϕ , is due to the body’s current spatial velocity and external forces (e.g., gravity). We use a discretization of the rigid body dynamics equation that is implicit in the contact forces to improve stability (the contact forces can change quickly in comparison to the other variables). Let w^{k+1} be the wrench on the body, in body coordinates, due to the sum of forces f_i^{k+1} at contact points r_i^k . This wrench is written

$$w^{k+1} = -{}^w_b \text{Ad}^T \Gamma^T f^{k+1},$$

where Γ is a block column matrix of size $3N \times 6$ (for N contacts) with block i equal to $\Gamma(r_i^k)$. This sums the forces in the block column vector f as wrenches in world coordinates. Notice we use the inverse transpose of the adjoint to map the summed wrench from world coordinates to body coordinates (recall the Newton-Euler equation is in body coordinates). We first rearrange Equation 3 to solve for accelerations,

$$\dot{\phi} = M^{-1}(w^{k+1} + [\phi^k]^T M \phi^k + w_{ext}),$$

where w_{ext} is the external wrench and includes forces such as gravity. Taking one Euler step of size h we can write the new velocity as

$$\phi^{k+1} = \phi^k + hM^{-1}(w^{k+1} + [\phi^k]^T M \phi^k + w_{ext}).$$

Taking an additional Euler step, we can write the linear approximation of the motion of the contact points due to rigid body motion by multiplying by $h \Gamma {}^w_b \text{Ad}$. Here, the adjoint ${}^w_b \text{Ad}$ converts the body spatial velocity back to the world coordinates used in our previous equations, and the contact point velocities are computed via multiplication by Γ . Hence, we have the implicit linear approximation $U_\phi f^{k+1} + u_\phi$, where,

$$\begin{aligned} U_\phi &= -h^2 \Gamma {}^w_b \text{Ad} M^{-1} {}^w_b \text{Ad}^T \Gamma^T, \\ u_\phi &= h \Gamma {}^w_b \text{Ad} (\phi^k + hM^{-1}([\phi^k]^T M \phi^k + w_{ext})). \end{aligned}$$

7.1.3 Linear Complementarity Problem

Finally, we use the pieces we have derived above to build a Linear Complementarity Problem (LCP). The complementarity conditions at each point (Equation 9) can be combined into a single system, where the forces and displacements are coupled using Equation 10. Letting $D = \text{diag}(D_1, \dots, D_N)$, $E = \text{diag}(E_1, \dots, E_N)$, $F = \text{diag}(F_1, \dots, F_N)$, $s = (s_1^T, \dots, s_N^T)^T$, $\beta = (\beta_1^T, \dots, \beta_N^T)^T$, $v = (v_1^T, \dots, v_N^T)^T$, $\lambda = (\lambda_1, \dots, \lambda_N)^T$, and $\sigma = (\sigma_1, \dots, \sigma_N)^T$, we now write

$$\begin{pmatrix} D^T(G - U_\phi)D & -E^T \\ F & 0 \end{pmatrix} \begin{pmatrix} \beta \\ \lambda \end{pmatrix} + \begin{pmatrix} s' \\ 0 \end{pmatrix} = \begin{pmatrix} v \\ \sigma \end{pmatrix} \perp \begin{pmatrix} \beta \\ \lambda \end{pmatrix},$$

where $s' = s + D^T(J\Delta\theta_r - Gf^k - u_\phi)$. This LCP can be solved with Lemke’s method (see [Murty 1988]). The solution provides the contact forces, $f^{k+1} = D\beta$, with which we advance the rigid body system by integrating Equation 3 twice. We also use these forces to compute the new joint angles from the equilibrium equation,

$$\theta^{k+1} = \theta_r^{k+1} + C^{k+1} J^T f^{k+1}.$$

This is effectively a single linear approximation step towards the equilibrium, and is one of the sources of error in Equation 7.

The last step in advancing the system is to identify new contact points for the new configuration θ^{k+1} and body configuration q^{k+1} . If existing contacts are replaced with the new points, additional error is introduced to the equilibrium equation due to translation of the forces. This can be accommodated at the next time step by adding the resulting joint displacement to $\Delta\theta_r$ when computing s' .

8 Results

Our retargeting simulator is implemented in Java. Complex models and multi-finger contact can make the simulation run slightly slower than real time due to collision detection and larger LCPs.

Figure 1(a) shows a typical interaction capture trial with a small box measuring $3 \times 5 \times 7$ cm. The box has a mass of approximately 12 g. Figure 13 columns (a) and (b) show a similar interaction applied to a virtual version of this box, but with different masses. Because our technique preserves the idea, or intent, of the original interaction, the virtual hand does not automatically succeed in picking up a heavier object. However, we can emulate human responses to such changes, which are of two kinds: unanticipated, and anticipated.

8.1 Grip Adjustment

When picking up an object that has been unexpectedly changed, with no visual cues to identify the difference, we can adapt with a relatively slow grip tightening response [Johansson 1998]. The reaction time for this unanticipated change is about 100 ms. Figure 13 columns (c) and (d) show an example of a simulated delayed grip tightening response due to slip at the fingertips.




Figure 13: Interaction synthesis examples. (a) Grasp applied to a small box with same mass; (b) on a heavy box; (c) without grip tightening response to unanticipated changes; (d) with grip tightening on slip with 100 ms delay; (e) grip adjustment based on imminent slip for anticipated changes; (f) interaction capture with a pressure sensitive ball; (g) manipulation applied to a similar object.




Figure 14: Grip adjustment based on imminent slip, shown for contact at the thumb. Left, when the ratio equals the coefficient of friction the object slips out of grasp. Right, when the ratio exceeds $\mu_{\text{threshold}}$, the compliance is reduced and the grip force increases.

In contrast to an unanticipated change, if the difference is evident (e.g., from visual cues such as shape variations) we do not wait until we feel our fingers slipping. Instead, we directly apply the appropriate grip force. One possible view is that we are actually adapting the same motor program to handle a family of different shapes. A similar strategy may be used when skin mechanoreceptors provide advanced warning of possible slip, by measuring vibration at the periphery of the contact patch. In this vein, we can modify the interaction trajectory to accommodate new objects by monitoring imminent slip.

We check for imminent slip during simulation by comparing the length ratio of friction and normal forces with the coefficient of friction. When the ratio exceeds a threshold, we deem the grip to be on the verge of slipping and decrease the compliances by a small factor, which tightens the grip at the next time step. Decreasing compliance to increase the grip strength raises the joint torques, but also provides the higher stiffness we would expect to see with larger contact forces.

Figure 13(e) shows an example for the box grasp applied to a wedge with grip adjustment. Due to its shape, the wedge requires larger grip forces to overcome gravity. Figure 14 shows the corresponding synthesized forces and ratio, with and without adjustment. With a discretized friction cone, slip can occur for a ratio less than μ . Here we use a threshold of 0.9μ , and we only compute the ratio when contact forces exceed 0.2 N , which allows slip during the very light touch at the initial time of contact.

Note that grip tightening may not always improve grasp quality; in some cases, shape variations are better suited to different interaction trajectories (i.e., different motor programs), or trajectories with reference angles modified to control the direction of grasp forces.

8.2 Pressure Sensitive Surfaces

Fingertip mounted sensors (shown in Figures 1(a) and 9) are useful because they allow the capture of interaction with arbitrary objects. For certain interactions, however, the attached sensors affect the subject's motion. Alternatively, we can capture interaction with objects that measure forces on their surface. For instance, Figure 13(f) shows a manipulation task captured using a pressure sensitive ball. We did not estimate compliances from the captured data because the pressure measurement rate and accuracy for this sensor are inadequate. Nevertheless, we can use the data to construct an interaction trajectory using previously estimated compliance values. Figure 13(g) shows the result of resynthesizing the captured finger-gaiting motion with a virtual light bulb.

8.3 Limitations

Although we do not address joint limits, this is not a problem for the examples in this paper because we focus on resynthesis under relatively similar conditions (e.g., we do not warp the interaction trajectory to accommodate large variations in object size). Nevertheless, our LCP should be easily modifiable to handle the inequality constraints imposed by joint limits (i.e., consider adding additional inequality constraints to the minimization problem in Section 7.1.1). In this paper we do not address the existence or uniqueness of the LCP in Section 7.1.3. For the cases we considered, we observe that our implementation always finds a solution. Furthermore, we believe our formulation is better behaved than dynamic LCPs as it is more closely related to the elastic LCPs used by Pauly et al. [2004].

Note that our implementation does not have any path or grasp planning. As such, in the grasping examples, we manually place the object in approximately the same location as it was during interaction capture. We then push the object to an improved initial configuration by simulating the grasping motion until all fingers are in contact followed by the same joint motion in reverse.

9 Conclusions

We have introduced the technique of Interaction Capture to address the difficulties of dealing with contact during motion capture. Interaction capture and synthesis promises to improve how motion capture retargeting can be done when contact is involved. It also has potential applications in variety of areas, such as robot programming, evaluation of ergonomics, and design of prosthetic devices. Additionally, though our focus is on hands and grasping, the work we present here could be applied to arms, interaction through the use of tools, and compliant articulated structures in general.

Our compliance estimation approach has the important advantage that we do not perturb the subject during the capture. Although the use of fingertip mounted sensors may alter the subject's motion, the effect of this is less disruptive than the perturbation in traditional methods. The impedance used by a subject during an interaction depends on many factors (user, task, intent, geometry, etc.). Because of this, our approach is desirable in the same way that motion capture is desirable in traditional settings; it provides compliance estimates for a given captured interaction of interest.

Finally, our interaction resynthesis technique allows a quasi-static compliant kinematic model to be incorporated into a dynamic environment with friction. This method consists of a position based linear complementarity problem that incorporates friction, breaking contact, and the compliant coupling between contacts at different fingers.

Acknowledgements: This work was supported in part by NSF grants IIS-0308157, ITR-0205671, EIA-0215887, Rutgers University, and Alias. Special thanks to Timothy Edmunds, Danny Kaufman, Shinjiro Sueda, Qi Wei, and François Faure.

References

- ANITESCU, M., AND POTRA, F. A. 2002. A time-stepping method for stiff multibody dynamics with contact and friction. *International J. Numer. Methods Engineering* 55, 753–784.
- BARAFF, D. 1994. Fast contact force computation for nonpenetrating rigid bodies. In *Proceedings of SIGGRAPH '94*, ACM Press, A. Glassner, Ed., Computer Graphics Proceedings, Annual Conference Series, 23–34.
- BIZZI, E., HOGAN, N., MUSSA-IVALDI, F. A., AND GISZTER, A. 1992. Does the nervous system use equilibrium-point control to guide single and multiple joint movements? *Behavioral and Brain Sciences* 15, 603–613.
- BURDET, E., OSU, R., FRANKLIN, D. W., MILNER, T. E., AND KAWATO, M. 2001. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414 (November), 446–449.
- CATMULL, E. 1972. A system for computer generated movies. In *Proceedings of the ACM Annual Conference*, 422–431.
- DELP, S., AND LOAN, J. 2000. A computational framework for simulating and analyzing human and animal movement. *Computing in Science and Engineering* (September).
- ELKOURA, G., AND SINGH, K. 2003. Handrix: Animating the human hand. In *ACM SIGGRAPH Symposium on Computer Animation*, 110–119.
- FELDMAN, A. G. 1986. Once more on the equilibrium-point hypothesis (lambda model) for motor control. *Journal of Motor Behavior* 18, 1 (March), 17–54.
- GLEICHER, M. 1998. Retargeting motion to new characters. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, 33–42.
- GOMI, H., AND KAWATO, M. 1997. Human arm stiffness and equilibrium-point trajectory during multi-joint movement. *Biological Cybernetics* 76, 163–171.
- HAIJAN, A. Z., AND HOWE, R. D. 1997. Identification of the mechanical impedance at the human finger tip. *Journal of biomechanical engineering* 119, 1, 109–114.
- HASSER, C. J., AND CUTKOSKY, M. R. 2002. System identification of the human hand grasping a haptic knob. In *Proceedings of the 10th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems (HAPTICS'02)*.
- HODGINS, J. K., AND POLLARD, N. S. 1997. Adapting simulated behaviors for new characters. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, 153–162.
- HOGAN, N. 1984. Adaptive control of mechanical impedance by coactivation of antagonist muscles. *IEEE Transactions on Automatic Control* AC-29, 8, 681–690.
- IKEMOTO, L., ARIKAN, O., AND FORSYTH, D. A. 2006. Knowing when to put your foot down. In *SI3D '06: Proceedings of the 2006 symposium on Interactive 3D graphics and games*, ACM Press, New York, NY, USA, 49–53.
- JOHANSSON, R. S. 1996. *Sensory and memory information in the control of dextrous manipulation*. Kluwer Academic, 205–260.
- JOHANSSON, R. S. 1998. Sensory input and control of grip. In *Novartis Foundation Symposium*, vol. 218, 45–63.
- KUCHENBECKER, K. J., PARK, J. G., AND NIEMEYER, G. 2003. Characterizing the human wrist for improved haptic interaction. In *Proc. ASME Int. Mechanical Engineering Congress and Exposition*, vol. 2.
- KURIHARA, T., AND MIYATA, N. 2004. Modeling deformable human hands from medical images. In *ACM SIGGRAPH Symposium on Computer Animation*, 357–365.
- LIU, C. K., HERTZMANN, A., AND POPOVIC, Z. 2005. Learning physics-based motion style with nonlinear inverse optimization. *ACM Trans. Graph.* 24, 3, 1071–1081.
- LOTSTEDT, P. 1981. Coulomb friction in two-dimensional rigid-body systems. In *Zeitschrift für Angewandte Mathematik und Mechanik*, 605–615.
- MILLER, A. T., AND CHRISTENSEN, H. I. 2003. Implementation of multi-rigid-body dynamics within a robotic grasping simulator. In *IEEE International Conference on Robotics and Automation*, vol. 2, 2262–2268.
- MILNER, T. E., AND FRANKLIN, D. W. 1998. Characterization of multijoint finger stiffness: dependence on finger posture and force direction. *IEEE Transactions on Biomedical Engineering* 45, 11 (November), 1363–1375.
- MURTY, K. G. 1988. *Linear Complementarity, Linear and Nonlinear Programming*. Heldermann Verlag, Berlin.
- NEFF, M., AND FIUME, E. 2002. Modeling tension and relaxation for computer animation. In *SCA '02: Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 81–88.
- PAULY, M., PAI, D. K., AND GUIBAS, L. J. 2004. Quasi-rigid objects in contact. In *SCA '04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ACM Press, New York, NY, USA, 109–119.
- POLLARD, N. S., AND ZORDAN, V. B. 2005. Physically based grasping control from example. In *SCA '05: Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ACM Press, New York, NY, USA, 311–318.
- POPOVIC, Z., AND WITKIN, A. 1999. Physically based motion transformation. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, ACM Press/Addison-Wesley Publishing Co., 11–20.
- RANCOURT, D., AND HOGAN, N. 2001. Stability in force-production tasks. *Journal of Motor Behavior* 33, 2, 193–204.
- SHAPIRO, A., PIGHIN, F., AND FALOUTSOS, P. 2003. Hybrid control for interactive character animation. In *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, IEEE Computer Society, 455.
- STEWART, D. E., AND TRINKLE, J. C. 1996. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *International J. Numer. Methods Engineering* 39, 2673–2691.
- TERAN, J., BLEMKER, S., HING, V. N. T., AND FEDKIW, R. 2003. Finite volume methods for the simulation of skeletal muscle. In *SCA '03: Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, 68–74.
- TSANG, W., SINGH, K., AND FIUME, E. 2005. Helping hand: an anatomically accurate inverse dynamics solution for unconstrained hand motion. In *SCA '05: Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ACM Press, New York, NY, USA, 319–328.
- XU, Y., AND HOLLERBACH, J. M. 1999. A robust ensemble data method for identification of human joint mechanical properties during movement. *IEEE Transactions on Biomedical Engineering* 46, 4 (April), 409–419.
- YIN, K., CLINE, M. B., AND PAI, D. K. 2003. Motion perturbation based on simple neuromotor control models. In *Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, 445–449.
- ZORDAN, V. B., AND HODGINS, J. K. 2002. Motion capture-driven simulations that hit and react. In *SCA '02: Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 89–96.
- ZORDAN, V. B., MAJKOWSKA, A., CHIU, B., AND FAST, M. 2005. Dynamic response for motion capture animation. *ACM Trans. Graph.* 24, 3, 697–701.