

Physically Based Rigging for Deformable Characters

Steve Capell, Matthew Burkhart, Brian Curless, Tom Duchamp, and Zoran Popović

University of Washington

Abstract

In this paper we introduce a framework for instrumenting (“rigging”) characters that are modeled as dynamic elastic bodies, so that their shapes can be controlled by an animator. Because the shape of such a character is determined by physical dynamics, the rigging system cannot simply dictate the shape as in traditional animation. For this reason, we introduce forces as the building blocks of rigging. Rigging forces guide the shape of the character, but are combined with other forces during simulation. Forces have other desirable features: they can be combined easily and simulated at any resolution, and since they are not tightly coupled with the surface geometry, they can be more easily transferred from one model to another. Our framework includes a new pose-dependent linearization scheme for elastic dynamics, which ensures a correspondence between forces and deformations, and at the same time produces plausible results at interactive speeds. We also introduce a novel method of handling collisions around creases.

1. Introduction

Believable computer animation requires that virtual *characters* such as humans and animals be produced with a high degree of realism. Faces should express emotions through mouth and eyebrow movements, limbs should bend at joints, muscles should bulge when in use, and soft tissue such as fat should bounce and vibrate when the character walks.

To create such realism, the traditional animation pipeline requires that each character be *rigged*, a process that is analogous to setting up a puppet to be controlled by strings. After having been rigged, a character’s shape can be controlled through a set of abstract parameters with meaningful names, like “lift left eyebrow” or “bend right knee.” For each keyframe, instead of having to position each vertex of the surface mesh, the animator need only set the values of the control parameters. To include fleshy bounces and vibrations, the animator must either create them by hand or tack a physical simulation onto the geometrically crafted deformations. In interactive settings, such as video games, these additional motions are not always possible to anticipate, necessitating some form of physical simulation.

Elastic simulation has proved to be a powerful method both for automatically creating plausible skeleton-dependent deformations and for introducing secondary motions. These simulations can be performed quickly with techniques like the finite element method for linearized elastic dynamics, al-

lowing for real-time simulation suitable for interactive video games and rapid prototyping for film production.

A significant shortcoming of these approaches is that the simulations do not provide a way for the animator to control the shape, other than by posing the skeleton. In this paper we address this limitation by introducing force-based rigging. In our system, deformations are created indirectly through the use of forces.

Section 3 describes our physical and numeric framework. Our system is built on a new method of pose-dependent linearization of elasticity, which ensures a correspondence between deformations and forces, while producing plausible results (both static and dynamic equilibria) at the speed of linear dynamics (Section 3.3). To increase realism, we introduce a new technique for handling collisions near creases (Section 3.4). Section 4 describes our rigging framework, which includes interactive optimization-based rig configuration (Section 4.2), computing rig forces from sculpted or measured surface deformations (Section 4.3), and transferring rigging between characters (Section 4.4).

2. Background

Many methods have been devised for the geometric deformation of surfaces, including free-form deformation [SP86, MJ96], shape interpolation [LCF00, SIC01, ACP02], and

wire deformation [SF98]. Such methods form the building blocks for rigging in modern animation programs such as Maya. These methods are not based on physical simulation, so the animator must hand craft physically realistic motions. The work of Kry et al. uses simulated input deformations to produce very realistic (but not dynamic) deformations [KJP02].

Dynamic deformation based on continuum elasticity was introduced to computer graphics by Terzopoulos et al. [TPBF87]. Terzopoulos and Witkin showed that the equations of motion can be linearized about a moving frame of reference, as long as the deformations are modest [TW88]. The linearized equations are more stable and can be solved quickly. We will discuss more recent quasi-linear schemes in Section 3.3.

Physical constraints have been used to make elastic models more controllable for animation [PB88, WW90, BW92], but constraints have not been shown to be useful for providing detailed shape control.

Detailed geometric deformation was combined with dynamic deformation in an anatomical modeling framework by Wilhelms and Van Gelder [WG97] (a similar approach by Scheepers et al. did not include a dynamic layer [SPCM97]). The muscle shapes are a function of the configuration of the skeleton, and a dynamic fatty layer is modeled to connect the statically deforming muscles to dynamic skin. Our framework differs in its ability to allow general shape control in unison with dynamic deformation of the entire character. An interactive system for building anatomical models was presented by Aubel and Thalmann [AT01].

Our framework is based on linear dynamics and can produce interactive animations, as in previous work on speeding up dynamic deformable models [DDCB01, JP02, MDM*02]. In particular, we build on the framework presented by Capell et al., where skeletal control for dynamic deformable models was introduced [CGC*02].

Singh et al. used forces to effect deformations [SOP95]. Our system is distinguished by the computation of such forces from desired deformations, and the overall rigging framework based on forces.

3. Deformable character formulation

Our modeling and simulation of deformable characters builds on the framework of Capell et al. [CGC*02], augmented with rigging forces, a new pose-dependent linearization scheme, and collision handling near creases.

The basic data defining a *deformable character* are: (i) an elastic domain $\Omega \subset \mathbb{R}^3$ (the interior of a triangular mesh Γ), with specified mass density, Poisson ratio, and Young's modulus, (ii) a control lattice $K \subset \mathbb{R}^3$ containing Ω , (iii) a skeleton $S \subset \Omega$ consisting of prescribed edges of K , and (iv)

a parameterized family $\mathbf{f}(\mathbf{x}, \Theta)$ of force fields (*rigs*) that can be used to deform the character into a desired shape.

The state of the character at time t is represented by the function

$$\mathbf{p}: \Omega \times \mathbb{R} \rightarrow \mathbb{R}^3 : (\mathbf{x}, t) \mapsto \mathbf{p}(\mathbf{x}, t) = \mathbf{x} + \mathbf{d}(\mathbf{x}, t), \quad (1)$$

Our animation framework gives a method for interactively determining the *displacement vector* $\mathbf{d}(\mathbf{x}, t)$ from a time-dependent family $\Theta(t)$ of control parameters that determine the pose of the skeleton and the force fields $\mathbf{f}(\mathbf{x}, \Theta(t))$ by solving a system of linear differential equations (Equation (7) below).

Animation is controlled by a time-dependent vector of *pose parameters* $\alpha(t) \in \mathbb{R}^{N_\alpha}$, and another of *abstract parameters* $\beta(t) \in \mathbb{R}^{N_\beta}$. The pose parameters and abstract parameters are combined into a single vector of *control parameters*

$$\Theta = (\alpha_1, \dots, \alpha_{N_\alpha}, \beta_1, \dots, \beta_{N_\beta}) = (\theta_1, \dots, \theta_{N_\theta}) \in \mathbb{R}^{N_\theta}.$$

The components of β serve as general-purpose shape controllers. For example, one such parameter value could be named *raised left eyebrow* and be expected to change the shape of the character appropriately. The animator may then implement the raising of the left eyebrow via a curve $\beta(t)$ in parameter space. An animation of the entire character is created by varying the control parameters $\Theta(t)$ over time.

The parameters Θ affect the shape of the character through objects that we call *rigs* (a variation on standard animation terminology). Each rig maps the control parameters to a force field acting on the body. The collection of all rigs on a character is thus modeled by a continuous family $\mathbf{f}(\mathbf{x}, \Theta)$ of force fields on Ω . Notice that we allow all of the elements of Θ to affect the shape, not just the elements of β . This allows the pose to affect the deformation, which is useful paradigm in animation. In particular, anatomical features of the shape such as muscle bulges are often pose dependent.

3.1. Finite element approximation

To apply the finite element method, we need a suitable basis $\mathcal{B} = \{\phi_i(\mathbf{x})\}$ of functions on the domain Ω . Each basis function $\phi_i \in \mathcal{B}$ is centered at a vertex \mathbf{x}_i of K and is the restriction to Ω of a piecewise trilinear function on K , with $\phi_i(\mathbf{x}_i) = 1$ and with support contained in the union of cells of K containing \mathbf{x}_i . We write the displacement in the form

$$\mathbf{d}(\mathbf{x}, t) = \sum_i \mathbf{q}_i(t) \phi_i(\mathbf{x}), \quad \mathbf{q}_i(t) \in \mathbb{R}^3. \quad (2)$$

By construction, $\mathbf{d}(\mathbf{x}_i, t) = \mathbf{q}_i(t)$ for each vertex $\mathbf{x}_i \in K$.

By definition, the *skeleton* of the character is a set S of edges of K contained in Ω . By construction, $\mathbf{p}(\mathbf{x}, t)$ is linear on each edge of S , and the coefficients $\mathbf{q}_i(t)$ at vertices of S are constrained so that the length of each edge of S is preserved by the mapping \mathbf{p} . By requiring the set S to be a tree

that mirrors the joint hierarchy, we ensure that the skeletal parameters $(\alpha_1, \alpha_2, \dots, \alpha_{m_\alpha})$ uniquely determine the value of the control coefficients \mathbf{q}_i at the vertices of S . To simplify notation, we assemble the coefficients into two column vectors: a vector \mathbf{q}_S whose i -component is \mathbf{q}_i , for i a vertex of S , and a vector \mathbf{q} whose i -th component is \mathbf{q}_i , for i vertex of $K \setminus S$. As we already observed, \mathbf{q}_S is determined by the control parameters Θ .

3.2. Dynamics

The motion of the character is determined by specifying the control parameters $\Theta(t)$ and solving the *Euler-Lagrange* equations

$$\frac{d}{dt} \left(\frac{\partial T}{\partial \dot{\mathbf{q}}_i} \right) + \frac{\partial U}{\partial \mathbf{q}_i} = \mathbf{Q}_i \quad (3)$$

for the unconstrained coefficients \mathbf{q}_i , where $T = T(\dot{\mathbf{q}}, \mathbf{q}_S)$ is the kinetic energy, the potential energy $U = U(\mathbf{q}, \mathbf{q}_S)$ is the sum of the elastic energy (a fourth degree polynomial in \mathbf{q} and \mathbf{q}_S) and a twist energy that penalizes displacement near the skeleton and has the effect of simulating a thickened skeleton (see [CGC*02]), and $\mathbf{Q} = \mathbf{Q}(\Theta(t))$ is the vector of generalized forces determined by $\Theta(t)$. Each component

$$\mathbf{Q}_i(\Theta) = \int_{\Omega} \mathbf{f}(\mathbf{x}, \Theta) \phi_i(\mathbf{x}) dV \quad (4)$$

of \mathbf{Q} represents the force acting on one control vertex at a point in $K \setminus S$. Later in the paper we will describe the creation and placement of specific force fields.

Substituting the expansions for T , U and \mathbf{Q} in terms of the basis functions into Equation (3) yields the system of nonlinear, ordinary differential equations

$$\sum_j \mathbf{M}_{ij} \ddot{\mathbf{q}}_j = \mathbf{Q}_i^e + \mathbf{Q}_i(\Theta) - \sum_k \mathbf{N}_{ik} \ddot{\mathbf{q}}_{S,k}(t), \quad (5)$$

where $\mathbf{M}_{i,j}$ and $\mathbf{N}_{i,k}$ are constant 3×3 matrices. The function $\mathbf{Q}_i^e = -\partial U / \partial \mathbf{q}_i$ measures the internal elastic force acting on the body, and $\mathbf{Q}_i(\Theta)$ is the generalized rigging force centered at \mathbf{x}_i . The indices i and j range over vertices in $K \setminus S$ and k ranges over the vertices of S .

3.3. Skeletally warped linearization

The nonlinear system in Equation (5) is computationally expensive to solve, and potentially unstable, so various methods have been devised to linearize, quasi-linearize, or simplify the equations [TW88, CGC*02, MDM*02, ITF04, MG04]. These methods can efficiently produce plausible deformations, but each of them has shortcomings in the context of the present application. In particular, we require the following additional features that are not present in previous approaches:

- Our framework uses forces to effect a given deformation in a given pose. It is important that forces and deformations be in correspondence for a given pose.

- When computing the rig forces for a given surface deformation at a given pose, we rely on the static equilibrium solution to the equations of motion. Computing static equilibrium must be efficient.

A pose-dependent linear system satisfies these goals.

We combine the benefits of the methods of Capell et al. [CGC*02] and Müller et al. [MDM*02], while eliminating some of their disadvantages. Like Capell et al., we use the configuration of the skeleton to predict the shape of the deformed object. But our method produces much more plausible deformations near joints. Like Müller et al., the linearization is performed independently at each vertex. But solving for static equilibrium using their method is a nonlinear problem. Our method is fast, stable, and free of gross distortions.

Similar to [MDM*02], we linearize the elastic force by estimating the local transformation at each control vertex, resulting in the following approximation:

$$\mathbf{Q}_i^e = \frac{\partial V}{\partial \mathbf{q}_i} \approx \sum_j \mathbf{R}_i \mathbf{S}_{ij} \left(\mathbf{R}_i^T (\mathbf{x}_j + \mathbf{q}_j - \mathbf{c}_i) - \mathbf{x}_j \right), \quad (6)$$

where the terms \mathbf{S}_{ij} are constant 3×3 *stiffness matrices*, and \mathbf{R}_i and \mathbf{c}_i are the estimated rotation matrix and translation vector, respectively, at node i . Note that our formulation includes an estimate of the translation in order to accommodate the translation-dependent twist energy used in [CGC*02].

Rather than measuring \mathbf{R}_i from the current state of the body, as in [MDM*02], we estimate \mathbf{R}_i and \mathbf{c}_i from the pose of the skeleton using a *skinning* methodology (e.g., [ACP02, MG03]). The pose parameters $\alpha(t)$ determine a Euclidean motion of the form $\mathbf{R}_b(\mathbf{x} + \mathbf{c}_b)$, sending each edge or “bone” $b \subset S$ to the proper position in \mathbb{R}^3 . We associate each vertex \mathbf{x}_i of K with two bones of S , denoted b_{i1} and b_{i2} and we let $\mathbf{R}_i = \omega_{i1} \mathbf{R}_{i1} \oplus \omega_{i2} \mathbf{R}_{i2}$ and $\mathbf{c}_i = \omega_{i1} \mathbf{c}_{i1} + \omega_{i2} \mathbf{c}_{i2}$. The user-defined weights ω_{ik} dictate how much each bone affects each point. The symbol \oplus in the formula for \mathbf{R}_i represents interpolation of rotations. In our current implementation we allow only two nonzero weights so that spherical linear interpolation can be used. More general methods are available for blending more than two rotations (e.g., [Ale02]), at higher computational cost.

It is also necessary to transform the force fields to be aligned with the skeleton. For instance, a force field creating a bump might change to a dimple when the body is rotated 180° , if the forces are applied in the global frame. We therefore transform the forces according to the transformation of the body. For simplicity and stability, we use the transformation predicted by the skinning procedure. Hence, given a particular skeletal configuration Θ at time t , the generalized force at vertex i is of the form $\mathbf{R}_i \mathbf{Q}_i$. We note that our approach to interpolation by warping the force fields according to the skeleton is analogous to the interpolation of displacements used by Lewis et al. [LCF00].




Figure 1: Comparison between our skeletally warped linearization method and the blended local linearization scheme from [CGC*02]. The object shown on the left, instrumented with a simple skeleton, is put into two poses. On the top row are the results using blended local linearization. Notice how the region near the bend unrealistically contracts, similar to traditional skinning methods used in animation. On the bottom row are the results using our new method, which look much more plausible.

With these changes the system (5) assumes its final form

$$\sum_j \mathbf{M}_{i,j} \ddot{\mathbf{q}}_j = \sum_j \mathbf{R}_i(\Theta) \mathbf{S}_{ij} \left(\mathbf{R}_i^T(\Theta) (\mathbf{x}_j + \mathbf{q}_j - \mathbf{c}_i(\Theta)) - \mathbf{x}_j \right) + \mathbf{R}_i(\Theta) \mathbf{Q}_i(\Theta) - \sum_k \mathbf{N}_{i,k} \ddot{\mathbf{q}}_{s,k}(\Theta). \quad (7)$$

Notice that (7) is a pose-dependent linear system. A comparison between our new method of linearization and the one presented in [CGC*02] (blended local linearization) is shown in Figure 1.

It is noteworthy that Müller et al., in [MG04], present an improvement over [MDM*02]. To eliminate ghost forces, they estimate the rotation for each cell rather than at each vertex. Since our system is highly constrained by the skeleton, ghost forces have not been a problem. However, it would not be difficult to combine our approach with that of [MG04], predicting the rotation of each cell from the nearby skeleton. The result would offer the advantages of pose-dependent linearity without ghost forces.

3.4. Self-collision detection and response

To realistically model human dynamics it is essential to take self-collisions of the surface of the character into account. When the human arm bends, for example, skin creases and




Figure 2: Without (a) and with (b) collision detection.

fold, forcing muscles and fat to bulge outward around the elbow. We model such behavior within our framework by introducing constraints that prohibit surface interpenetration.

Interactive rates are achieved through a preprocessing step that determines the locations of likely collisions for a given animation. A deterministic propagating-front collision scheme, partially inspired by [PPG04], is used to incrementally add constraints to the system as the collision surface expands. Raghupathi et al. have employed a similar technique in that they check the neighborhoods of regions in contact to find all colliding pairs in the area [RGF*04]. We expand on these ideas by introducing a preprocessing step that reduces the number of collision tests needed at runtime.

The image shown in Figure 2(b) was generated using the collision response method described in this section, absent of any rigging forces. Our solution satisfies our requirement of efficiency. As Figure 2 illustrates, the dynamical equations then produce the desired bulging effects as an artifact of the equations of elasticity.

3.4.1. Preprocessing

To efficiently detect self-collisions of the surface Γ as a simulation progresses, we determine approximate collisions using a preprocessing step. Consider a *target pose*, such as a bent elbow, in which self-collisions of the surface Γ occur. By choosing a family of control parameters $\Theta(t)$ that slowly deforms the skeleton from its rest state to the desired pose and solving the dynamic equations, we obtain a family of near equilibrium states deforming the character into the target pose. As this “slow motion” deformation evolves, we track the position of each vertex of the triangle mesh Γ .

By definition, a *collision* occurs when a vertex of Γ first penetrates a triangle of Γ . When a collision occurs, we activate a constraint that prevents further penetration (see below) and continue the dynamic simulation. The *seed set* D is the sequence of vertex-vertex pairs obtained by pairing each penetrating vertex with one of the vertices of the triangle that it penetrates. We make the *a priori* assumption that this sequence of pairs will approximate the set of self-collisions that occur during any subsequent animation of the character. If this assumption holds, the result is identical to that of general collision detection.

3.4.2. Run-time behavior

During a dynamic simulation, collisions not included in the seed set (generated in the preprocessing step) may occur. To maintain interactive rates, it is essential that we maintain the list of current collisions and a list of locations at which future collisions are likely to occur.

The *collision list* is the ordered list

$$C = \{(i_1, b_1), (i_2, b_2), \dots, (i_{n_c}, b_{n_c})\} \quad (8)$$

of all pairs of the form (i, b) where i is a vertex of Γ and b are the barycentric coordinates of the point at which the vertex penetrates a triangle. The elements of C are ordered according to the time when each vertex enters the region Ω .

If a collision (i, b) occurs in the time interval $[t_0 - \Delta t, t_0]$, we impose the inequality constraint on (i, b) crafted to prevent the vertex from moving deeper into the interior of Ω

$$\chi_{i,b}(t) := (\dot{\mathbf{r}}_i(t) - \dot{\mathbf{r}}_b(t)) \cdot \mathbf{n}(t) \geq 0 \quad \text{for all } t \geq t_0, \quad (9)$$

where $\mathbf{r}_i(t)$ and $\mathbf{r}_b(t)$ denote spatial locations of i and b at time t , and $\mathbf{n}(t)$ is the normal to the triangle T at time t .

We classify elements of C into *active* and *inactive* elements. More precisely, let $n_a \leq n_c$ be the smallest non-negative integer such that if we impose the equality constraint

$$\chi_{i_k, b_k}(t) = 0 \quad \text{for all } k < n_a$$

then $\chi_{i_k, b_k} > 0$ for all $n_a \leq k \leq n_c$.

We assume that future collisions will either occur near one of the collision seeds or in a neighborhood of the current collision set. Specifically, the *neighborhood* of a pair $(i, p) \in C$ is the set $N(i, b)$ of all pairs (i', j') of vertices of Γ such that i' is in a neighborhood of i and j' is in a neighborhood of p .

At each time step we perform the following operations:

- (1) For each pair $(i, j) \in P$, check for penetration of i into each triangle in a neighborhood of j as well as for penetration of j into each triangle in a neighborhood of i . When a penetration is detected, add the corresponding pair (i, b) or (j, b) to C .
- (2) Remove any pair (i, b) from C for which the point $\mathbf{r}_i(t)$ is exterior to Ω . The inequality

$$(\mathbf{r}_i(t) - \mathbf{r}_b(t)) \cdot \mathbf{n}(t) > 0$$

provides an easy test.

- (3) Update P to be the union $P = D \cup \bigcup_{(i,b) \in C} N(i, b)$.
- (4) Find the number n_a that determines the set of active elements of C .

Note that because we detect collisions after they have occurred, small interpenetrations will occur. However, our constraints prevent any further incursion once collisions are detected. The result is visually appealing, leaving no visible gaps between regions in contact.

To maintain the constraints $\chi_{i,b} = 0$ we use the modified conjugate gradient method presented by Baraff and Witkin [BW98] and extended to FEM by Capell et al. [CGC*02].

The set of all colliding points and active constraints $\chi_{i,b} = 0$ in this formulation will generally be significantly higher than the number of degrees of freedom provided by the control lattice. We build an orthogonal basis in constraint space from the list of constraints using a QR-decomposition. This basis spanning the active constraint space is used to project out components of the solution that violate the constraints within the conjugate gradient (CG) solver. The use of such orthogonalization of constraints to project the null space of the constraint basis has been shown to still lead to guaranteed convergence of the CG [AB03].

4. Rigging – shape control using forces

Shape control is important because it enables the animator to create more realistic deformations and express the emotion and intent of a character. For example, skeletal controls will not effectively help the animator to make a character appear to breathe, produce a muscle bulge, or smile. To address this shortcoming, we introduce an additional control mechanism.

Since our framework is based on physical simulation, two natural mechanisms exist for influencing the shape: hard constraints and forces. Forces (or force fields) are better suited to our framework for a variety of reasons: (1) In contrast with hard constraints, which are rigid and therefore not a realistic model for animal tissues other than bones, forces influence the shape without destroying its dynamic nature. (2) Forces can coexist without troublesome compatibility issues—if two force fields overlap, their combined effect is obtained by vector addition. (3) Finally, force fields need not line up exactly with the character’s surface geometry. Our framework already handles the embedding of the character; if the force field extends beyond the interior of the character it can easily be restricted to the interior.

To use forces to effect deformations, we exploit an important feature of pose-dependent linearization. For a fixed skeletal pose, deformations and generalized forces are in correspondence. Consider the body in static equilibrium ($\dot{\mathbf{q}} = \mathbf{0}$ and $\ddot{\mathbf{q}} = \mathbf{0}$). Equation (7) reduces to the form

$$\mathbf{A}(\alpha)\mathbf{q} = \mathbf{b}(\alpha) + \mathbf{Q}, \quad (10)$$

where $\mathbf{A}(\alpha)$ is a matrix (containing 3×3 blocks), and $\mathbf{b}(\alpha)$ is a vector (with elements that are 3-vectors). Since $\mathbf{A}(\alpha)$ is invertible, there is a bijection between generalized forces \mathbf{Q} and displacements \mathbf{q} . This relationship allows us to think about shape in terms of forces. While our simulations are dynamic, the rigging is configured while running a static equilibrium solver (i.e., solving Equation (10)), where forces and displacements are in correspondence.

The rigging process, by which the mapping \mathbf{f} is constructed is based on two independent paradigms, based on

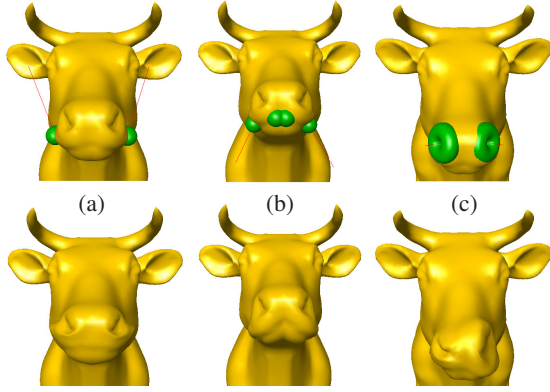


Figure 3: A cow head interactively rigged using force field templates to create simple facial expressions. The templates are visualized in the top row. The bottom row shows the resulting deformations. (a) Two bump templates create a smile. (b) Four bump templates effect a frown. (c) Two torus templates produce dilation and contraction of the nostrils.

two kinds of components: *force field templates* and *surface deformation rigs*. Force field templates, are reusable components representing parameterized force fields in a canonical coordinate frame. We will discuss them in Section 4.2. Surface deformation rigs (hereafter referred to as *surface rigs*) are rigs constructed from surface deformations. This gives the user more control by leveraging existing technologies such as surface scanners and geometric modeling software. Surface deformations are presented in Section 4.3.

4.1. The rigging process

To rig a character, individual rigs must be configured and combined to form the rigging force $\mathbf{f}(\mathbf{x}, \Theta)$. Each rig is described as a force $\mathbf{f}_a(\mathbf{x}, \eta_a)$, where η_a is a vector of *rig parameters* that determine the force produced by the a -th rig. The rigging process consists of deciding on a set of rigs and creating a mapping from Θ to η_a for each rig. For simplicity, we impose the restriction that the a -th rig depends on a single control parameter $\gamma_a \in \Theta$. The rig force is then formed by combining the force contributions from the rigs:

$$\mathbf{f}(\mathbf{x}, \Theta) = \sum_{a=1}^{N_{rigs}} \mathbf{f}_a(\mathbf{x}, \tau_a(\gamma_a)), \quad (11)$$

where the function $\tau_a : \gamma_a \mapsto \eta_a$ maps the control parameter associated with rig a to its rig parameters. It is also convenient to express the rig force purely in terms of generalized forces:

$$\mathbf{Q}_i(\Theta) = \mathbf{R}_i(\alpha) \sum_{a=1}^{N_{rigs}} \mathbf{Q}_i^a(\gamma_a), \quad (12)$$

where \mathbf{Q}_i^a is the generalized force (on the free variables) contributed by the a -th rig.

We represent the function τ_a as the linear interpolation of a set of sample points $\{(\gamma_{a,k}, \eta_{a,k}) : k = 1, 2, \dots, N_a^\tau, \gamma_{a,k} < \gamma_{a,k+1}\}$ that are created by the user during an interactive session:

$$\tau_a(\gamma_a) = \begin{cases} \eta_{a,0}, & \text{for } \gamma_a \leq \gamma_{a,0} \\ \eta_{a,k} + (\eta_{a,k+1} - \eta_{a,k}) \left(\frac{\gamma_a - \gamma_{a,k}}{\gamma_{a,k+1} - \gamma_{a,k}} \right), & \text{for } \gamma_{a,k} < \gamma_a \leq \gamma_{a,k+1}, 1 \leq k < N_a^\tau \\ \eta_{a,N_a^\tau}, & \text{for } \gamma_{a,N_a^\tau} < \gamma_a. \end{cases}$$

The process by which the user interactively rigs a character can be described algorithmically as follows:

```

a ← 1
while the character is not fully rigged do
  instantiate rig a
  choose a control parameter  $\gamma_a \in \Theta$  for the rig
  k ← 1
  while the user desires more samples do
    configure the character by setting  $\Theta$ 
    set the control parameter  $\gamma_a \leftarrow \gamma_{a,k}$ 
    configure the rig by setting  $\eta_a \leftarrow \eta_{a,k}$ 
    record the sample  $(\gamma_{a,k}, \eta_{a,k})$ 
    k ← k + 1
  end
  a ← a + 1
end

```

4.2. Rigging with force field templates

Our first approach to rigging is to rely on a library of predefined force field templates. Each template defines a parameterized force field in a canonical coordinate frame. For examples of specific force field templates and their parameters see [Cap04]. Figure 3, shows the result of rigging the head of a cow with force field templates.

Although force field templates are designed to be simple and intuitive, it is not always easy for the the user to understand which parameters to change to achieve the desired effect. We therefore support an interface wherein the user directly manipulates the surface of the character with the mouse and the system automatically optimizes the parameters of the rig so that the forces cause the surface of the character to conform to the user's input.

Using a drag-and-drop interface, the user indicates target positions for one or more points on the surface. Our goal is to choose η so that these constraints are met. This is accomplished by minimizing the cost function

$$C(\eta) = C_{points}(\mathbf{q}(\eta)) + C_{params}(\eta),$$

where the term $C_{points}(\mathbf{q}(\eta))$ enforces the constraints, and




Figure 4: In this example, a skeletal pose is computed that deforms one surface to fit the other. The spheres indicate user-selected feature points that guide the optimization.

the term $C_{params}(\eta)$ penalizes deviation of η from its preferred value η_0 . Note that $\mathbf{q} = \mathbf{q}(\eta)$ via Equation (10). Put simply, in static equilibrium, a configuration of the rig parameters η produces a particular displacement \mathbf{q} .

We optimize the cost C using the L-BFGS-B algorithm [BLNZ94], a quasi-Newtonian solver with limited memory usage. Standard optimization methods for smooth cost functions, like L-BFGS-B, require knowledge of the gradient $\frac{dC}{d\eta}$. The reference [Cap04] gives a method for efficiently computing $\frac{dC}{d\eta}$ by taking advantage of the fact that our system is linear at a fixed pose. Each evaluation of $\frac{dC}{d\eta}$ requires one integration over the character for each element of η , and a single solution of the linear Equation (10).

4.3. Deriving rigs from surface deformations

In addition to predefined force field templates our framework supports surface deformation rigs, which are constructed directly from surface deformations. A surface deformation rig uses an arbitrary generalized force to effect the desired shape change. We express the force of a surface deformation rig directly in the generalized form of Equation (12), dropping the index a since we will only be discussing a single rig:

$$\mathbf{Q}(\gamma) = \gamma \boldsymbol{\eta}. \quad (13)$$

For surface deformation rigs, the rig parameters η determine the direction of the generalized force and the rig control γ determines its magnitude. The rig is configured by computing the optimal parameters η so that a target surface is matched.

Given a target surface $\bar{\Gamma}$ represented as a triangle mesh, we want to compute a generalized force such that $\bar{\Gamma} \approx \Gamma$, where Γ is the surface of the simulated body at static equilibrium. This computation replaces the user interaction in the rigging process. Because Γ also depends on the skeletal configuration parameters α , which are unknown for the target surface, we must also compute an optimal skeletal pose. Our approach to optimizing the skeletal pose is based on the work of Allen et al. [ACP02].

Our procedure for pose optimization is based on matching a set of user-selected feature points $\{\mathbf{p}_i\}$ on Γ and $\{\bar{\mathbf{p}}_i\}$ on $\bar{\Gamma}$ using the following cost function:

$$C_{features}(\alpha) = \frac{1}{N_{features}} \sum_{k=1}^{N_{features}} |\mathbf{p}_k - \bar{\mathbf{p}}_k|^2,$$




Figure 5: A surface deformation rig for a bent arm. (a) A scanned arm in its rest configuration. The arm has been instrumented with a skeleton and control lattice for elastic simulation. (b) The target surface, which was also acquired by scanning. We want to compute the forces that approximately produce the target surface when applied to the simulated arm shown in (a). (c) The static equilibrium shape of the arm in (a) posed to match (b) as well as possible. (d) A comparison between (b) and (c). (e) The bent arm after applying an optimized surface deformation rig that uses forces to produce the target shape. (f) A comparison between (b) and (e).

where $N_{features}$ is the number of feature points. Although it requires user interaction, the selection of feature points by the user has the advantage that points can be chosen that are representative of the configuration of the underlying skeleton, such as places where the bone is close to the surface. The L-BFGS-B algorithm is used to solve for the optimal pose parameters α . Figure 4 shows the results of optimizing the pose to fit a target surface.

After finding the optimal skeletal pose associated with $\bar{\Gamma}$, and choosing the value of γ , we compute the optimal generalized force (via η) that aligns the deformed surface Γ with the target surface $\bar{\Gamma}$. Our goal is then to find a vector η that minimizes the following cost function that penalizes mismatch between Γ and $\bar{\Gamma}$:

$$C(\eta) = C_{surf}(\eta) + \omega C_{smooth}(\eta), \quad (14)$$

where C_{surf} penalizes the distance between the deformed surface and the goal surface, C_{smooth} enforces smoothness, and ω is a user-specified smoothing parameter. For the precise definition of the cost terms see [Cap04].

Figure 5 shows the results of creating a surface deformation rig to match a bent arm. Because the deformation is associated with the arm in a bent pose, we chose the angle of the elbow joint as the control parameter γ associated with the

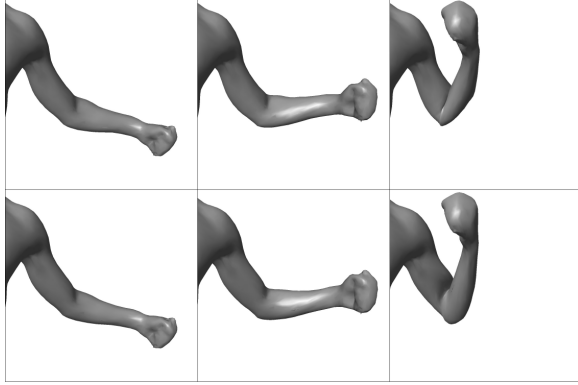


Figure 6: The rig from Figure 5 used at different poses. The top row shows the simulation without rigging. The bottom row shows the same poses simulated with rig forces in effect. Notice that the biceps and the area around the crease of the elbow are much more realistic looking when the rig is used.

rig. As γ varies from 0 to 1 the force that effects the deformation is gradually introduced. In this manner, intermediate poses of the elbow produce intermediate deformations. Figure 6 shows the effect of the rig when the pose of the arm differs from the pose of the target data.

4.4. Retargetting surface deformation rigs

Deformed surfaces can be difficult to construct, and sometimes a deformation is available for one model but not for another similar model. For these reasons we would like to transfer rigs between characters. In this section we discuss the transfer of surface deformation rigs from one character to another. We accomplish this by creating a mapping between the two characters and then using the map to transfer forces.

Suppose that we have a character whose domain is Ω with surface $\Gamma = \partial\Omega$ and skeleton S , and that the character has been parameterized by a cell complex K in the manner described in Section 3.2. Our goal is to create a mapping from Ω to a new character $\bar{\Omega}$ with surface $\bar{\Gamma} = \partial\bar{\Omega}$ and skeleton \bar{S} . We approximate such a mapping by repositioning the vertices of the control lattice to form a homeomorphism $\mathbf{h} : K \rightarrow \mathbb{R}^3$ using the trilinear basis. Our goal is to choose \mathbf{h} so that $\mathbf{h}(\Omega) \approx \bar{\Omega}$.

We reposition the vertices of K using an optimization procedure that minimizes a cost function of the form

$$C(\mathbf{h}) = \omega_{skel}C_{skel}(\mathbf{h}) + \omega_{surf}C_{surf}(\mathbf{h}) + C_{distort}(\mathbf{h}). \quad (15)$$

The first term encourages the skeletons to match by penalizing the difference $\bar{S} - \mathbf{h}(S)$. The second term (defined in Section 4.3) matches the surfaces by penalizing the difference $\bar{\Gamma} - \mathbf{h}(\Gamma)$. The last term penalizes distortion in the mapping. The user defined weights ω_{skel} and ω_{surf} determine




Figure 7: (a) The unrigged Ganesh arm (in a bent pose). (b) The arm with rigging transferred from the arm scan in Figure 5. On Ganesh the rig causes similar effects as on the original arm model: the elbow extends, becoming less rounded, and the biceps bulges).

how much each term contributes to the total cost. We refer the reader to [Cap04] for details.

Once the mapping \mathbf{h} from Ω to $\bar{\Omega}$ has been established, we can use it to transfer a generalized force \mathbf{Q} (or η , recall Equation (13)) from Ω to $\bar{\Omega}$. Transferring the generalized forces directly between different physical systems produced unintuitive results, so instead we convert the generalized forces to displacements, transfer the displacements, and then convert back to generalized forces. In doing so we take advantage of the fact that our pose-dependent linear system draws a correspondence between generalized forces and displacements.

Consider a generalized force \mathbf{Q} acting on the character whose domain is Ω . We can convert it to a displacement \mathbf{d} by first applying Equation (10) to compute the generalized coordinates \mathbf{q} and then applying the basis expansion in Equation (2). The laws of vector transformation dictate that the transformed displacement $\bar{\mathbf{d}}$ satisfies the equation $\bar{\mathbf{d}}(\bar{\mathbf{x}}) = \mathbf{J}(\mathbf{x})\mathbf{d}(\mathbf{x})$, where $\mathbf{J}(\mathbf{x}) = \frac{\partial \mathbf{h}}{\partial \mathbf{x}}(\mathbf{x})$. We approximate this transformation by transforming the coefficients: $\bar{\mathbf{q}}_i = \mathbf{J}(\mathbf{x}_i)\mathbf{q}_i$.

Although \mathbf{h} is not differentiable at \mathbf{x}_i , we can estimate \mathbf{J}_i at the vertex \mathbf{x}_i of the lattice by computing the affine map that minimizes the energy function

$$E(\mathbf{J}_i) = \sum_{j=1}^N (\mathbf{J}_i \mathbf{e}_{ij} - \bar{\mathbf{e}}_{ij})^2, \quad (16)$$

where $\mathbf{e}_{ij} = \mathbf{r}_j - \mathbf{r}_i$ and $\bar{\mathbf{e}}_{ij} = \bar{\mathbf{r}}_j - \bar{\mathbf{r}}_i$ are edge vectors surrounding vertex i in K and \bar{K} , respectively. We minimize E by the Procrustes algorithm [BM92]. Once we have computed $\bar{\mathbf{q}}$ as above, Equation (10) can be applied again to produce the generalized force $\bar{\mathbf{Q}}$ for the target character.

5. Results

Some of the results of our rigging system have already been shown in Figures 2, 3, 5, and 6. We tested the system on two

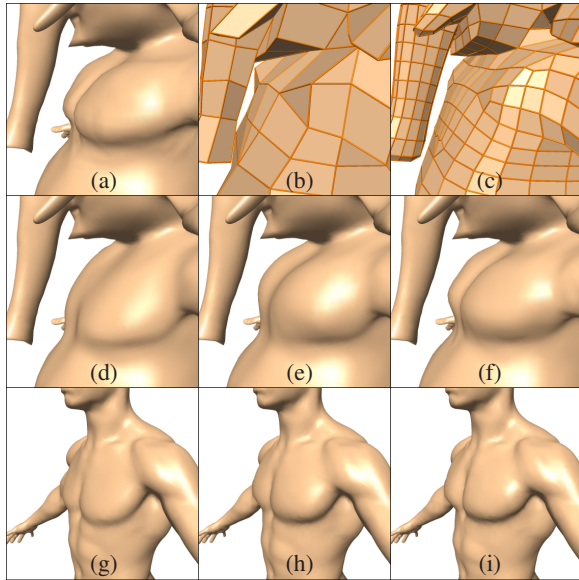


Figure 8: A surface deformation rig effecting a chest flex. (a) A deformation representing Ganesh with a flexed chest, created using geometric modeling software. The second row shows the use of a force field rig to effect the same shape change. (d) The model at rest. (e) A coarse simulation (the surface of the volumetric lattice is shown in (b) above). (f) A finer simulation (the lattice is shown in (c) above). The bottom row shows the chest flex rig transferred to the human model: (g) at rest, (h) coarse, and (i) fine.

additional input surfaces: “Ganesh” (a rotund humanoid figure with an elephant’s head) designed using geometric modeling software, and a scan of a human. The Ganesh character was instrumented with three rigs which were then transferred to the human character. One of the rigs was derived from a third model, the arm scan demonstrated in Figures 5 and 6. Another rig, representing the flexing of the chest, was derived from a surface deformation designed by an artist. A third rig, effecting a breathing motion of the torso, was created using a force field template.

Figures 7 and 8 demonstrate the transfer of rigging between characters. In Figure 7, the rig derived from the bent arm scan (Figure 5) has been transferred to the Ganesh character. It is noteworthy that the thin human arm on which the rig is based and the Ganesh arm behave quite differently when using the unrigged model. The Ganesh arm is short, fat, and cartoon-like, so the deficiencies of the simplified physical model are not as apparent as for the thin human arm as in Figure 5 (c) and (d). Despite their significant differences, the rig produces a similar effect on Ganesh as on the scanned arm. The biceps bulges in an appropriate place and the elbow extends, counteracting an unrealistic contraction.




Figure 9: Animations using a chest flex rig. (a) Animation with no force-based rigging. (b) Animation using a chest flex rig. (c) Animation using the chest flex and a belt constraint.

The transfer of a chest flex rig from ganesh to the human model is shown in Figure 8. The chest flex rig was created using a target deformation created using geometric modeling software. The transferred rig produces a plausible deformation despite the fact that the characters are very different.

Figure 9 demonstrates frames of an animation of the rigged Ganesh character. In row (c), a belt constraint has been applied, demonstrating that our system supports constraints along with rigging. Although we have not done extensive timing studies, the cost of simulation is about equally split between overhead (e.g., bookkeeping, rendering) and solving the sparse linear system at each time step.

6. Conclusion

Our system gives animators control over the shapes of elastic deformable characters by introducing force fields as the building blocks of rigging. The resulting simulations combine the guidance of the animator with other influences such as gravity, physical constraints, and user interaction.

A disadvantage of using forces to effect deformations is that in order to create finely detailed deformations, the dynamic simulation must be computed at a high resolution. To reduce computation in such cases, adding displacements to the dynamic shape might be preferable to using forces. It would be interesting to further explore this tradeoff.

Another possible extension to this work would be the interpolation of physical properties as a part of the rigging. In addition to forces, it would fit nicely into our framework to adjust the stiffness of part of the character as a function of rig parameters. It would also be interesting to try to extract such physical properties from captured data.

Acknowledgments Thanks to Brett Allen, Keith Grochow, Daichi Sasaki, Supriyo, Ernest Wu, and Yeuhi Abe. This research was supported by NSF grants EIA-0121326, CCR-0092970, IIS-0113007, and CCR-0098005, and by UW Animation Research Labs, an Alfred P. Sloan Fellowship, Electronic Arts, Sony, Microsoft Research, Alias, and Washington Research Foundation.

References

- [AB03] ASCHER U. M., BOXERMAN E.: On the modified conjugate gradient method in cloth simulation. *Visual Computer* 19, 7-8 (2003), 526–531.
- [ACP02] ALLEN B., CURLESS B., POPOVIĆ Z.: Articulated body deformation from range scan data. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2002)* 21, 3 (2002), 612–619.
- [Ale02] ALEXA M.: Linear combination of transformations. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2002)* 21, 3 (2002), 380–387.
- [AT01] AUBEL A., THALMANN D.: Interactive modeling of the human musculature. In *Proceedings of Computer Animation 2001* (2001).
- [BLNZ94] BYRD R. H., LU P., NOCEDAL J., ZHU C.: A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16, 6 (1994), 1190–1208.
- [BM92] BESL P. J., MCKAY N. D.: A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 2 (Feb. 1992), 239–258.
- [BW92] BARAFF D., WITKIN A.: Dynamic simulation of non-penetrating flexible bodies. *Computer Graphics (Proceedings of SIGGRAPH 92)* 26, 2 (July 1992), 303–308.
- [BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. In *Proceedings of SIGGRAPH 98* (July 1998), pp. 43–54.
- [Cap04] CAPELL S.: *Interactive Character Animation Using Dynamic Elastic Simulation*. PhD thesis, University of Washington, Department of Computer Science and Engineering, 2004.
- [CGC*02] CAPELL S., GREEN S., CURLESS B., DUCHAMP T., POPOVIĆ Z.: Interactive skeleton-driven dynamic deformations. *ACM Transaction on Graphics (Proceedings of ACM SIGGRAPH 2002)* 21, 3 (2002).
- [DDCB01] DEBUNNE G., DESBRUN M., CANI M.-P., BARR A. H.: Dynamic real-time deformations using space & time adaptive sampling. *Proceedings of SIGGRAPH 2001* (2001).
- [ITF04] IRVING G., TERAN J., FEDKIW R.: Invertible finite elements for robust simulation of large deformation. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on computer animation* (2004), pp. 131–140.
- [JP02] JAMES D. L., PAI D. K.: Dyrat: Dynamic response textures for real time deformation simulation with graphics hardware. *Proceedings of SIGGRAPH 2002* (2002), 582–585.
- [KJP02] KRY P. G., JAMES D. L., PAI D. K.: Eigenskin: Real time large deformation character skinning in hardware. In *Proceedings of the 2002 ACM SIGGRAPH Symposium on Computer Animation* (2002), pp. 153–159.
- [LCF00] LEWIS J. P., CORDNER M., FONG N.: Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. In *Proceedings of SIGGRAPH 2000* (2000), pp. 165–172.
- [MDM*02] MÜLLER M., DORSEY J., MCMILLAN L., JAGNOW R., CUTLER B.: Stable real-time deformations. In *Proceedings of the 2002 ACM SIGGRAPH Symposium on Computer Animation* (2002), pp. 49–54,189.
- [MG03] MOHR A., GLEICHER M.: Building efficient, accurate character skins from examples. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2003)* (2003), 562–568.
- [MG04] MÜLLER M., GROSS M.: Interactive virtual materials. In *Proceedings of Graphics Interface 2004* (2004), pp. 239–246.
- [MJ96] MACCRACKEN R., JOY K. I.: Free-form deformations with lattices of arbitrary topology. *Computer Graphics (Proceedings of SIGGRAPH 96)* (1996), 181–188.
- [PB88] PLATT J. C., BARR A. H.: Constraint methods for flexible models. *Computer Graphics (Proceedings of SIGGRAPH 88)* 22, 4 (August 1988), 279–288.
- [PPG04] PAULY M., PAI D. K., GUIBAS L. J.: Quasi-rigid objects in contact. In *SCA '04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2004), ACM Press, pp. 109–119.
- [RGF*04] RAGHUPATHI L., GRISONI L., FAURE F., MARCHALL D., CANI M.-P., CHAILLOU C.: An intestine surgery simulator: Real-time collision processing and visualization. *IEEE Transactions on Visualization and Computer Graphics* 10, 6 (2004), 708–718.
- [SF98] SINGH K., FIUME E.: Wires: A geometric deformation technique. In *Proceedings of ACM SIGGRAPH 98* (1998), pp. 405–414.
- [SIC01] SLOAN P.-P. J., III C. F. R., COHEN M. F.: Shape by example. In *Proceedings of 2001 Symposium on Interactive 3D Graphics* (2001), pp. 135–143.
- [SOP95] SINGH K., OHYA J., PARENT R.: Human figure synthesis and animation for virtual space teleconferencing. In *Proceedings of the Virtual Reality Annual International Symposium* (1995), pp. 118–126.
- [SP86] SEDERBERG T. W., PARRY S. R.: Free-form deformation of solid geometric models. *Computer Graphics* 20, 4 (Aug. 1986), 151–160.
- [SPCM97] SCHEEPERS F., PARENT R. E., CARLSON W. E., MAY S. F.: Anatomy-based modeling of the human musculature. In *Proceedings of SIGGRAPH 97* (1997), pp. 163–172.
- [TPBF87] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. *Computer Graphics (Proceedings of SIGGRAPH 87)* 21, 4 (July 1987), 205–214.
- [TW88] TERZOPOULOS D., WITKIN A.: Physically based models with rigid and deformable components. *IEEE Computer Graphics and Applications* 8, 6 (Nov. 1988), 41–51.
- [WG97] WILHELMS J., GELDER A. V.: Anatomically based modeling. In *Proceedings of SIGGRAPH 97* (1997), pp. 173–180.
- [WW90] WITKIN A., WELCH W.: Fast animation and control of nonrigid structures. *Computer Graphics (Proceedings of SIGGRAPH 90)* 24, 4 (August 1990), 243–252.