Model-based Reconstruction for Creature Animation

Maryann Simmons Jane Wilhelms Allen Van Gelder
University of California, Santa Cruz

Abstract

An semi-automatic technique for creating 3D models of creatures suitable for animation is presented. An anatomically based canonical model is deformed, given a sparse set of feature points derived from measurements describing the target animal. The layered canonical model is built on top of an articulated structure hierarchy and contains a representation of the animal's skeleton, muscles, and skin. The joint hierarchy and associated body components are transformed based on the input data. A denser set of feature points is then automatically generated from the new underlying structural components. The feature points are used to deform the attached mesh skin representation, using a segmented interpolation approach. Results are shown using measurements from a scale model and from a live horse. Our main contributions are (1) a novel approach for automatically reconstructing complete jointed creatures from an anatomically based canonical model of similar structure; and (2) an integrated application of skin interpolation for both morphing and animation. In this research, we have addressed the problem in the context of modeling and animating horses; however, the general techniques that we have developed could be applied to a wide range of creatures, at the cost of constructing a canonical model for each creature type.

Keywords: model reconstruction, animation, 3D morphing, shape interpolation

1 Introduction

Computer-animated models of non-rigid articulated creatures are at the core of many computer graphics applications including entertainment, simulation, and design. The use of such models to visualize complex 3D geometry and motion has relevance to many other fields as well. Biomechanical gait analysis studies [Back and Clayton 2000; Delp and Loan 1995], for example, can greatly benefit from the use of realistic 3D models to visualize different aspects of measured and observed locomotion.

The creation of complex models suitable for animation, however, is a difficult task. In general, content creation is becoming a limiting factor in producing complex dynamic virtual environments in computer graphics. We have sophisticated illumination and appearance models and rendering engines to produce high-quality imagery, and fast graphics hardware to deliver interactive visualization of highly complex models, but limited means of attaining such models.

Traditionally modeling has been done by procedural techniques, or painstakingly by hand by experienced designers using complicated 3D modeling packages. A more recent trend in graphics is to utilize image-based techniques, where 2D images are used instead of 3D geometry to represent an environment. Such techniques can be used to increase realism and visual complexity in virtual environments. For dynamic environments where the user wants to view, animate, and manipulate the model from possibly any vantage point under varying conditions, 3D models still hold the advantage over purely image-based techniques.

Recently, there has been a significant increase in the quality and availability of 3D capture methods, both for model and motion capture. These techniques range from devices such as laser scanners, and optical and magnetic motion capture technology, to photogrammetric techniques that derive 3D information from digital image and video input. The scanner technologies coupled with software for model reconstruction [Levoy et al. 2000] work well for capturing rigid, static objects, but often require expensive hardware. More limiting, however, is that the scanning technology requires the subject to remain motionless for a period of time (seconds to minutes): it is therefore more difficult to capture live subjects, and not possible in the case of unconstrained and/or undirectable subjects such as animals. The desired creature may not even physically exist, and therefore these approaches alone are insufficient for constructing a rich set of unique models.

This paper introduces a novel method for 3D model creation that "morphs" a canonical anatomically based representation to create novel creatures of similar type suitable for animation. The canonical model is built on top of a layered articulated segment hierarchy and contains a representation of the creature's skeleton, muscles, and skin. A new creature is built from the "inside-out": we first transform the segments, then the bones, muscles, and skin. Transformations applied to the segment hierarchy at the lowest layer are propagated out along with subsequent changes to the bones and muscles, with the final effect of deforming the creature's skin.

Most existing automated modeling approaches decouple the process of acquiring geometric models from that of producing an an-

Copyright © 2002 by the Association for Computing Machinery, Inc.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1-212-869-0481 or e-mail permissions@acm.org.
© 2002 ACM 1-58113-573-4/02/0007 $5.00

imatable representation. We demonstrate how our techniques can be used to automatically generate 3D models suitable for animation. We have implemented our approach in the context of modeling and animating horses; however, given the availability of a suitable canonical model, the general techniques that we have developed could be applied to other articulated creatures, including humans.

1.1 Related Work

A large body of work has addressed the problem of model reconstruction for human faces (see [Parke and Waters 1996; Parent 2002] for an overview). While there have been impressive advances in this area, it is far from a solved problem. Model-based techniques (e.g. [Parke 1982; Pighin et al. 1998; Kurihara and Arai 1991]), deform the geometry of a generic face model to match feature points describing a target face. If photographs are used, additional texture information can be extracted and applied, resulting in the potential for very realistic models.

The goal for many applications using facial models is animation. Novel expressions can be generated by morphing between captured expressions. Other approaches generate plausible expressions and models from a large set of input models [Blanz and Vetter 1999] or anthropometric data [DeCarlo et al. 1998]. Performance based animation approaches use motion capture to drive existing facial models [Williams 1990].

These approaches are not physically based in the way that they do reconstruction or animation. Other techniques for facial and animal modeling and animation use a layered model of skin and muscles [Chadwick et al. 1989; Wilhelms and Van Gelder 1997; Ng-Thow-Hing 2001]. In some cases the muscles are activated to produce the motions or expressions [Lee et al. 1995]. Others utilize the tissue models only to deform the skin as a result of changes in the underlying components.

Less work has been done in full-body model acquisition and animation [Kakadiaris and Metaxas 1995; Nedel and Thalmann 1998]. Physically based models have been developed for individual body parts for humans and animals [Chen and Zeltzer 1992]. While accuracy is paramount for many biomedical applications, in other applications such as entertainment, the goal can be quite different. In these domains, a plausible representation of how a creature looks and moves, and the ability to easily adjust the model and motion in not necessarily physically based ways is often key. Layered anatomically based models [Wilhelms and Van Gelder 1997; Scheepers et al. 1997] capture some of the key features of a physically based representation in an efficient and controllable manner and can be found in modeling packages designed for animation [Alias/Wavefront 2000].

Shape interpolation for morphing and animation is another related area of research. The model-based facial morphing technique of Pighin et al [Pighin et al. 1998] utilizes the feature point vertices to derive an interpolation function that maps the rest of the vertices in the canonical face model to the novel, target face. Other techniques (e.g [Sloan et al. 2001]) blend between several models to generate the final result. Interpolation is also key for animation: where the skins of articulated creatures need to deform in a plausible manner as the creature moves. Most often, the movement of vertices near the joints is calculated as a blend of the motion of the adjacent segments. This approach is subject to artifacts where the skin can crease or bulge unconvincingly.

Recent work [Sloan et al. 2001; Lewis et al. 2000] in this area has produced techniques suitable for shape interpolation, as well as for smooth deformation of articulated skinned figures. All of the interpolation techniques listed here utilize radial basis functions for interpolation. We utilize similar interpolation techniques both for morphing of the skin and during animation. Our approach is novel in the use of a dense set of feature points, automatically derived from the anatomical components, to drive the shape interpolation to generate the positions of the remaining vertices: in this way we combine an anatomically based deformation approach with shape interpolation, producing a unified method that can be used both for morphing and animation of skinned creatures.

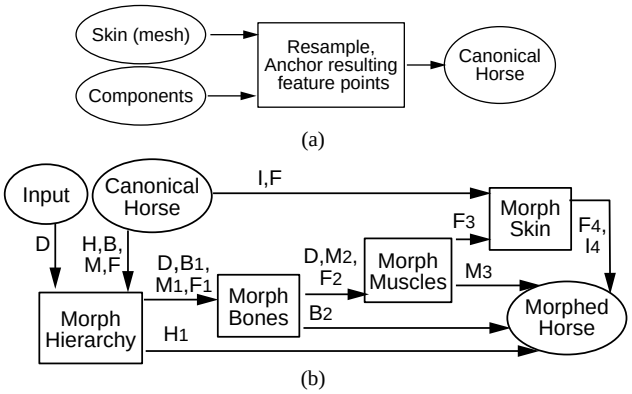

Figure 1: Overview of creature reconstruction a) Construction of canonical horse. b) Morphing steps: in the figure H, B, M, F, I stand for the structural hierarchy (H), bones (B), muscles (M), mesh feature (F) and interpolated vertices (I). The subscripts indicate the morphed version of the component after each step. D represents the input data for the target horse (e.g. marker locations).

1.2 Overview

This paper introduces a new anatomically based approach for model creation from a canonical model. For model-based approaches, three key subproblems are the acquisition of the feature points, definition of the generic canonical model, and deformation to fit the canonical model to the feature points.

One novel aspect of our approach is the use of the creature's anatomy to automatically generate a denser set of *feature points* from a sparse set of measurements. We demonstrate our algorithms with two examples: in the first we used a hand digitizer to acquire the 3D coordinates of pre-defined locations on the surface of a scale horse model, with pre-determined correspondence to vertices in our generic model; in the second example, we derived the necessary input information by taking hand measurements of a live horse.

Our canonical model utilizes a triangle mesh to model the skin of the creature. Our approach follows in the spirit of Parke's early work [Parke 1982] for facial modeling which uses a parameterized model based on conformational parameters, including positional and scale values. Our canonical model also incorporates an anatomically based layered structure hierarchy based on the techniques described in [Wilhelms and Van Gelder 1997; Schneider and Wilhelms 1998]. What is unique about our approach is that the layered representation is utilized not only to produce appropriate skin deformations during animation, but to generate the model itself. Lee *et al.* recognized not only the importance of incorporating a physically based muscle model, but also that the skull bone was responsible for the visual appearance of the skin as it is animated [Lee et al. 1995]. We incorporate the effect of structural differences in muscles and bones – both to generate and animate the model.

A subset of the skin vertices are specially designated as *feature* vertices and are used by the morphing and animation. To create a new creature, the structure hierarchy, muscles, and bones of the canonical model are deformed based on a sparse set of input measurements. A set of derived feature points is then automatically generated for the new model from the new underlying structural components. The feature points are used to deform the attached

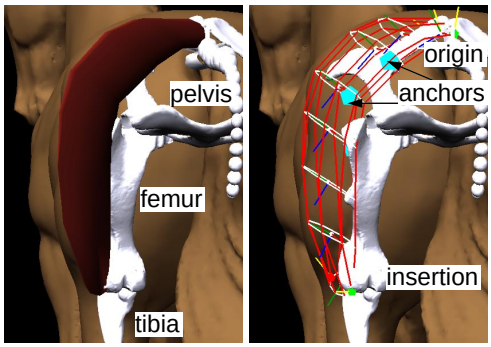

Figure 2: Muscle Model: a) left gluteal muscle, b) muscle detail.

mesh skin representation using a segmented interpolation approach.

We construct an interpolation function based on the derived feature set. Similar to previous approaches [Pighin et al. 1998; Noh and Neumann 2001; Carr et al. 2001] in model morphing and reconstruction, we use an interpolation function based on radial basis functions. We utilize a segmented interpolation scheme based on body components such that each interpolation has a computationally feasible set of input points, as well as reducing the control to more local, logical areas.

Our approach involves the following steps:

1. definition of a parameterized canonical representation,
2. morphing of the internal body components,
3. morphing of skin.

The overall approach is diagramed in Figure 1: the individual steps are described in the relevant sections below. Section 2 describes the canonical model. Section 3 describes the method for component morphing and Section 4 outlines the skin morphing method. Section 5 presents results and Section 6 concludes and discusses opportunities for future research.

2 Canonical Model

We are primarily concerned with generating visually convincing results, for creatures real or imagined. In this context, a fully accurate and complete canonical model for a particular animal would be prohibitively expensive to design and manipulate. The key to an efficient representation is to incorporate realistic components only where such detail is visually important to the model and resulting animations. To this end, we have developed a canonical horse representation from an anatomically based parameterized model composed of multiple layers [Wilhelms and Van Gelder 1997]. This model and its use for morphing and animation are described in the following sections. A more detailed description can be found in a related technical report [Simmons et al. 2002].

At the heart of the representation is the articulated structure hierarchy: where the body is represented as a collection of segments connected at joints. At the next layer, a bone representation (a general polygon mesh or ellipsoidal model) is attached to each relevant segment.

The deformation of muscles greatly affects the appearance of a creature during motion. The canonical representation includes key surface muscles attached to the appropriate bones. The muscles in our model are a very simplified approximation to true muscles, that roughly correspond to the muscle locations and shapes in horses based on information from equine anatomy books (e.g. [Goody 1983]). The goal is to produce reasonable behavior for morphing and animation.

Muscles are modeled as polygonal generalized cylinders. They deform in a volume-preserving manner during animation in response to movement of the bones to which they are attached. The muscle has a variable number of octagonal cross-section slices. The end slices, called the *origin* and *insertion*, are always attached to bones; the intermediate slices are optionally attached by *anchors*. The position, orientation, and size of the intermediate slices control the shape of the muscle. The origin and insertion points are constrained to move with the bones to which they are attached. The intermediate slices are automatically adjusted in orientation to smoothly interpolate between the end slices to maintain a good shape. Anchors can be specified on the intermediate slices to additionally constrain the shape, for example, to constrain the muscle to go over, rather than through, an intermediate bone. Figure 2 contains a close-up of the pelvis region and detail of the muscle structure. The muscle origin is at one of the vertebrae, the insertion at the tibia, and the muscle is anchored where it crosses over the pelvis in between.

A generalized tissue component modeled by ellipsoids is utilized to represent the remaining bulk of the animal structure where needed. The ears, for example, contain ellipsoidal components.

The skin comprises the final layer of the representation. We utilize a polygonal mesh constructed from laser scans of a scale horse model for the canonical skin. The skin is then attached to the underlying components. Other anatomically based approaches utilize the underlying components to generate the skin [Scheepers 1996; Wilhelms and Van Gelder 1997], but we believe that it is not possible with these techniques to get the high quality results that can be achieved from a detailed, carefully crafted mesh model.

The mesh is re-sampled to ensure that it contains vertices corresponding to nearby underlying components: for each bone and muscle, rays are shot from points on the component in the normal direction and intersected with the skin. A new vertex is added at the intersection point if the component is sufficiently close to the surface to make a visual impact on the skin, and the ray intersection does not occur at grazing angles. These points are designated as *feature vertices*. Each feature vertex is anchored to the associated bone or muscle. The anchor is parameterized in the local space of the component. When the underlying component is moved or deformed (in the case of a muscle) during animation, the attached skin point will move in response. Unlike previous approaches [Schneider and Wilhelms 1998; Ng-Thow-Hing 2001], we attach only the feature vertices – the remaining points are interpolated as described in Section 4. Figure 1a shows the process schematically, and 6a shows a visualization of the feature vertices.

An additional set of *marker vertices* can also be added to the mesh. These points are anchored to the underlying components by finding the closest point on the nearest muscle or bone surface. The markers indicate pre-specified locations where it is assumed positional information will be provided for the target creature, and are used for morphing as described in the following sections. Figure 3 illustrates the canonical horse model, showing a representation of the structural hierarchy, bones, muscles, and skin. The canonical model shown here contains 98 segments, 97 bones, 18 muscles (we have thus far only modeled a subset, mostly on the left side), and 87316 total vertices in the mesh, 1473 of which are feature vertices and 89 are marker vertices.

3 Morphing of Internal Components

With the canonical model in place, the goal is then to morph this generic model to another target animal, based on a sparse set of input measurements.

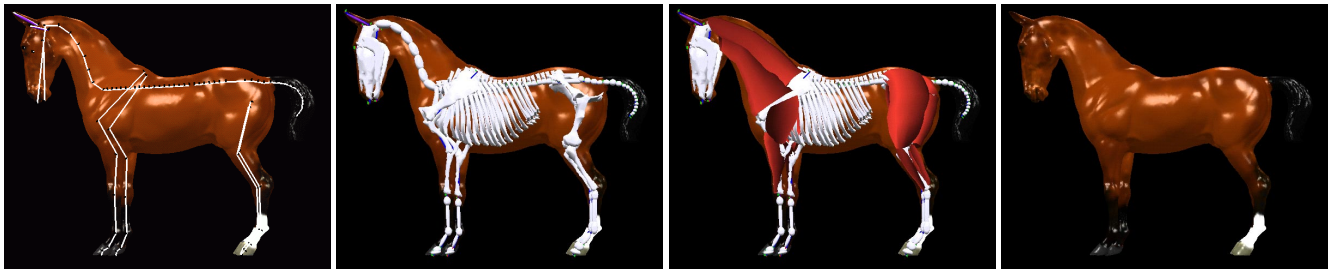

(a) (b) (c) (d)

Figure 3: Canonical Horse Model: a) segment hierarchy b) skeleton c) muscles, and d) skin.

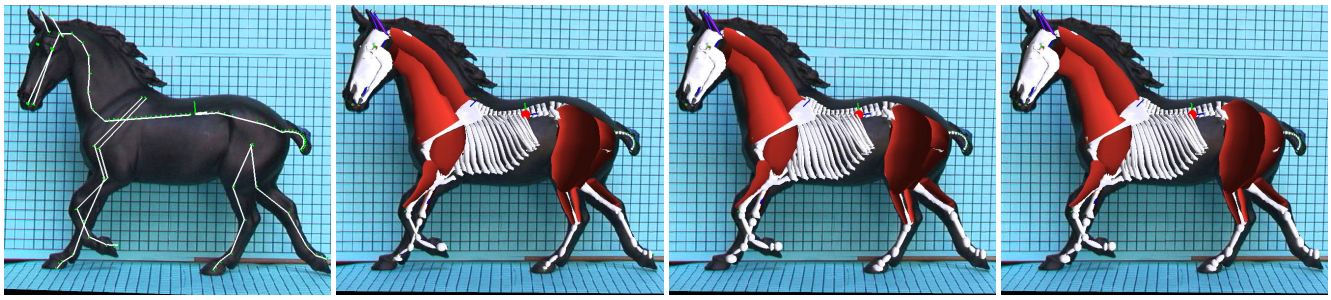

(a) (b) (c) (d)

Figure 5: Automatic Component Adjustment: a) Morphed segment hierarchy, b) resulting components (as a result of segment transformation only), c) after bone morphing, d) after muscle morphing.

3.1 Marker Points

We chose a set of 89 marker locations on the surface of the animal, roughly corresponding to structural points that are key for the purposes of shape interpolation and animation. As in the face model developed by Parke [Parke 1982], we utilize a combination of positional values and scale values to parameterize the model.

Figure 4 illustrates the marker locations on our example target animal. Points include those necessary to estimate the joint locations (e.g. *arm*) as well as points to estimate the individual structure of other components such as muscles (e.g. point *crest* gives a measure of the combined contribution of the trapezius and splenius muscles to the width of the neck). The lines indicate the proportion being measured when it is visible – a triangle indicates when the proportion is measured along the transverse axis of the animal. The total number of measured points (e.g. accounting for left and right measurements) is shown next to each marker name.

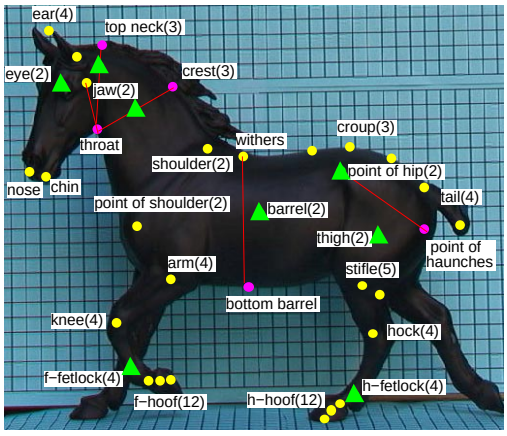

Figure 4: Marker Points used for morphing. Lines indicate measurements used for component scaling – when the measurement is along the transverse axis (i.e. into the image), it is indicated by a triangle.

It would be possible to assume symmetry to reduce the number of markers, but we chose to incorporate values from both the left and right sides. The appearance will be different not only because of asymmetries in the animal, but also because of differences in state. We want to derive as much information as possible from input data for a single position. Note that not all joints are represented: the remaining values are estimated according to knowledge about equine conformational structure. Additional values could be estimated for simplicity, or if the marker values were not available.

3.2 Segment Hierarchy Morphing

We have two methods for morphing the segment hierarchy, depending on the form that the input measurements take. One method accepts measured lengths of the subject's limbs and body dimensions. To calculate joint angles, measurements from the corresponding surface marker location to the ground and other reference planes, in combination with the limb lengths, can be used to estimate the joint angles. The second method accepts 3D coordinate values for all of the marker locations. This approach is described in more detail in what follows.

The geometric hierarchy of the canonical model is first transformed to match that of the target animal based on the estimated joint locations. The root of the hierarchy is at the lumbar vertebra whose location is estimated from the point halfway between the hip markers, which also define the local x axis. The z axis points towards the center of the top tail measurement, defining the direction of the spine at the back portion. For segments where there is an associated or derivable point from the markers, the transformation proceeds in a similar fashion down the hierarchy.

Where there are intermediate segments without associated markers, the internal segments are morphed based on the measured markers at the beginning and ending of a segment chain. The shape of the segment chain in the canonical representation is used to guide the placement of the intermediate segments. For example, there are 7 cervical vertebrae, but only marker information for the base, middle, and top of the neck. The remaining cervical locations for the top of the neck are calculated by rotating the z axis of the middle vertebra such that the line segment from the middle to the top vertebra (behind the ears) in the canonical model hierarchy lines up with the line segment between the measured locations. The difference in translation for the chain is then distributed over the rest of the vertebrae, such that the end of the chain coincides with the new desired location, while maintaining the basic shape of the neck specified in the canonical model.

Figure 5a shows the results of transforming the segment hierarchy to match the example target hierarchy. The results are shown superimposed on an image of the target animal for reference. The bones, muscles, and tissue components are parameterized with respect to the joint hierarchy, and are therefore automatically adjusted as a result of the transformation. Figure 5b shows the new bones and muscles produced as a result of the hierarchy adjustment.

3.3 Bone Morphing

At the next stage in the component morphing, the muscles and bones are further adjusted based on the input marker locations. As can be seen in Figure 5b, while the joint hierarchy matches well, it can be seen that not all of the components adequately represent the target animal (e.g. most conspicuous is the rib bones which are too small for this animal's trunk).

The bone models are differentially scaled according to the conformational scale values calculated from the markers. These values take into consideration the scale already implicitly applied due to the transformation resulting from the change in the joint hierarchy. For example, a measurement of the width of the barrel (midsection) of the horse (b_w) is used as a scale value (b_s), to adjust the rib bones by scaling outward by ($b_s = (b_{l_c} b_{w_t})/(b_{l_t} b_{w_c})$): where the b_l are the composite lengths of the canonical (c) and target (t) spine segments.

3.4 Muscle Morphing

Major muscle groups also contribute to the differences between the canonical and target model. The muscle origin, insertion, and anchor points are fixed relative to the bones, and therefore automatically reflect the changes to the bone from segment and bone morphing. The shape and volume of the muscle, however, may need to be further adjusted based on differences in the body dimensions. For example, in Figure 5c, the neck muscle is still too narrow after the composite affects of bone and component morphing.

The muscle slice dimensions are scaled based on the body dimensions in a similar manner as for the bone. For the bone, however, there is a single value each for x, y, and z applied to all the bone vertices: for the muscles the scale values are linearly interpolated across slices if there are one or more body dimension measurements that affect the width and or height of the muscle slices along the length of the muscle.

Figure 5d illustrates the final component transformation including bone and muscle adjustment.

4 Skin Morphing

The deformation of the skin proceeds by utilizing the deformed internal components to generate a set of deformed skin feature points, which are dense compared to the marker points, but quite sparse in relation to vertices of the skin mesh. These feature points are used to perform model-fitting by scattered data interpolation techniques.

As described in Section 2, the canonical skin mesh model is resampled so that it contains vertices corresponding to nearby component features. These points, and optionally the marker points, are called the *feature points*. Since the skin is anchored parametrically to the underlying components, these feature points automatically are deformed when the underlying components change as a result of the morphing. (In the context of animating one animal, the underlying components change as a result of movement, and muscles deform in response to changes in length induced by the movement, preserving their volume.)

The goal is to construct a smooth interpolating function that expresses the deformation of the non-feature skin vertices (which are not anchored in the morphed model) in terms of the changes in the feature points during morphing or animation. This problem is addressed by scattered data interpolation methods. Radial Basis Functions (RBFs) are a popular means of scattered data interpolation. They have been used in graphics for model-fitting, surface reconstruction, and for morphing [Carr et al. 2001; Lewis et al. 2000; Noh and Neumann 2001; Pighin et al. 1998; Turk and O'Brien 1999; Sloan et al. 2001]. We present a unified Anatomically-Driven-Deformation [1] and shape interpolation approach : the feature points are deformed according to the underlying structure, and these feature points are then used to drive RBF shape interpolation to generate the rest of the points.

The interpolant using RBFs is a function that returns the displacement value for each non-feature point that takes it from the original position to its position in the target form. The displacement $u_i = p_i' - p_i$ is known for each feature point p_i in the canonical model and p_i' in the target model. These displacements are utilized to construct an interpolating function $f(v)$ that returns the displacement for each skin vertex v. There are m skin vertices, of which n are feature points. We use the form:

$$f(v) = L(v) + \sum_{i=1}^{n} c_i \phi(\|v - p_i\|) \tag{1}$$

Here $L(v)$ is an affine function of $v \in R^3$, c_i are 3D-vector constants, $\|$ denotes Euclidean distance, and ϕ is a radial basis function. The points p_i are often called *centers* in the literature. Possibly $L(v)$ is zero, and for full generality it may be any polynomial. The c_i's are determined by requiring that $f(p_i) = u_i$ for $i = 1, \ldots, n$. This leads to $3n$ linear equations. When $L(v)$ is present it can be represented by a 3x4 matrix M, using the row-vector convention; M contains 12 unknowns, and 12 additional equations are derived by requiring that $\sum c_i = 0$ and $\sum c_i \otimes p_i = 0$, the latter being a 3x3 matrix formed by the tensor product. The linear system can be conveniently represented in matrix form:

$$\begin{bmatrix} \Phi & P \\ P^T & 0 \end{bmatrix} \begin{bmatrix} C \\ M^T \end{bmatrix} = \begin{bmatrix} U \\ 0 \end{bmatrix} \tag{2}$$

where the rows of P are the feature points p_i with appended homogeneous coordinate 1; the rows of C are c_i, the rows of U are u_i, and $\Phi_{ij} = \phi(\|p_j - p_i\|)$. (If $L(v)$ is omitted, omit the rows and columns for P and M.)

Nielson recommends using geodesic distance on the surface, rather than Euclidean distance as the argument to the RBF [Nielson 1993], but he only considers analytically specified surfaces. Given that we have large, irregular, triangular meshes, we adopted an approximation of this strategy based on *geodesic zones*. The model

[1] In [Lewis et al. 2000] the terminology SDD for Skeleton-Driven-Deformation is used: here the deformation includes the effects of muscles as well.

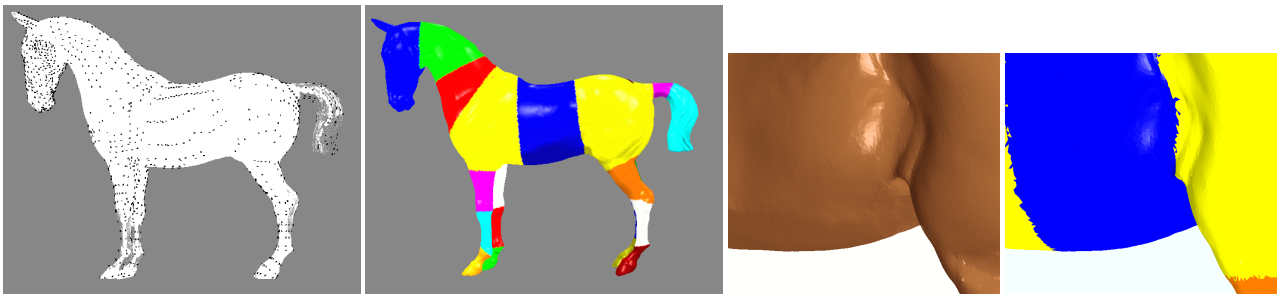

(a) (b) (c) (d)

Figure 6: Skin Morphing: a) canonical feature points, b) geodesic zones, c,d) close up of hip area shown with flat shading – illustrating that the interpolation is smooth across geodesic zone boundaries.

is segmented according to logical body parts into several geodesic zones. The idea is that Euclidean distance within one geodesic zone or neighboring zones is a fair approximation of geodesic distance.

Each geodesic zone has its own interpolation function. Its calculation is based upon the feature vertices of that zone and its immediate neighbor zones. As an example, the feature vertices for the shoulder and lower leg are used to define the interpolating function for the upper leg. This means that the skin vertices will get contributions from the displacement of "geodesically" nearby feature points, but none due to those "geodesically" distant. For example, the motion of the tail will not affect the skin vertices on the animal's head, and more importantly, one leg will not influence the other although it might be quite close in terms of Euclidean distance. See Figure 6b for a visualization of the geodesic zones.

This strategy has the effect that vertices in neighboring geodesic zones will be interpolated according to different interpolating functions, but since the functions were derived with a large overlap in the feature vertices, one does not see discontinuities, with a suitable RBF. Figure 6c,d shows a closer view of the boundary of different regions after the skin has been morphed using the interpolation, where it can be noted that the function varies smoothly across geodesic-zone boundaries.

An important aspect of using such an interpolation scheme is the selection of a radial basis function that will produce a smooth interpolation function in R^3. Several radial basis functions have been mentioned in recent literature. Some pros and cons are discussed by Carr *et al.* [Carr et al. 2001]. We evaluated several functions including the multi-quadric ($\sqrt{1+|r|^2/s^2}$), inverse-multi-quadric ($(1+|r|^2/s^2)^{-1/2}$) Gaussian ($e^{-|r|^2/2s^2}$) and soliton ($e^{-|r|/s}$), where s is a scale parameter. Because of the geodesic-zone strategy we preferred to have an RBF that died out for large values of its argument. We had better results visually with the soliton function, with s anywhere from 50 to 250 (the horse is roughly 300 units in length).

For greater than 2000 feature points, it is generally found that the matrices are too big; it becomes prohibitively expensive to solve the linear system, and matrices become ill-conditioned. Our geodesic zone strategy has the by-product that several smaller systems are solved independently, reducing the size of the matrices for each calculation.

5 Results

In this section, we present results of utilizing the morphing techniques described in this paper to automatically construct and animate a 3D horse model. To test our technique, we used a 3D digitizer [Immersion 2000] to measure the marker locations on the scale horse model shown in Figure 5a. We also took measurements off a live subject, the pony shown in Figure 7a. In this case, we hand-measured the limb lengths and body dimensions to derive the needed input data.

5.1 Model Reconstruction

Figure 7a-b shows the resulting components and mesh for the pony example, and 7c the mesh for the scale model of a percheron horse. It can be seen that the morphed horse models are not exact replications of the targets. With the sparse set of input markers, our goal was less to get a duplicate model, and more to capture the characteristics of the target horse, while maintaining a plausible horse-like appearance and behavior when animated. These criteria have been met in the morphed examples shown. To give a sense of scale, the canonical model, pony, and percheron models are shown together in Figure7d.

In Figure 7e, a portion of an interpolated skin for the percheron is shown flat-shaded in the bottom image. Despite the fact that the relative position and orientation of the animal changed substantially in non-uniform-ways (e.g the body of the target horse is much more substantial, but the legs are quite short), the resulting mesh is fairly smooth and captures the shape of the animal. An additional smoothing step could also be applied after the interpolation. The top image shows the same model if every vertex in the mesh is treated as a feature point (i.e. is anchored to the closest component). The rigid connections imposed by this approach lead to artifacts in the mesh, especially around joint locations. Our combined anchoring/interpolation approach provides a solution in which the skin follows the underlying components closely at selected locations where the bone or muscle is close to the surface, but can smoothly interpolate to avoid most of the artifacts that can occur when the surface "buckles" at joints.

In our approach, the marker points are optionally included in the feature point set for interpolation. We have found that since the marker points are constrained to map to a specific vertex in the mesh, the accuracy of the marker placement greatly influences the quality of the results. If the marker points are accurate, their inclusion improves the reconstruction, but can cause significant artifacts if they are very misaligned. In the examples shown, the markers were not included in the feature set. We calculated the difference in the interpolated marker locations from that of the input values. The r.m.s error was less than 2% of the length of the horse for the marker vertices and the difference in the overall interpolation with and without using the markers was r.m.s. error less than 1% (this corresponds to approx. .5 inches). This indicates that the feature point/interpolation scheme is performing well.

The running time of the algorithm is dominated by the RBF calculation. For the examples shown, the component morphing completed in approx. one second, while the morph of the skin took

closer to one minute for a mesh of 87K vertices. We have not as yet made any efforts at optimizing this stage: our current implementation links in Matlab to do the matrix calculations. Our method could benefit from recently introduced, faster approximate methods that can process larger numbers of points with high-quality results as described by Carr *et al.* [Carr et al. 2001].

The crux of the algorithm is the design of the canonical model. Slight changes in the location of the automatically generated feature points can greatly affect the results. Our current muscle implementation can be unstable for large relative changes in the origin and attachment bones during morphing and animation. The benefit of our approach is that a good quality canonical model will generate good quality novel models, and therefore the effort to construct such a model is well-spent.

5.2 Animation

Figures 7f,g show the results of animating both the canonical model and the automatically generated models. The interpolation techniques described in Section 4 were applied to estimate the locations of the non-feature mesh vertices based on the movement of the anchor points when the horse is animated. The images in this paper were generated using $s = 50$ for the RBF soliton function.

The animations shown required that the horse model be in a specified starting position. This was done semi-automatically by a routine that adjusts the joint angles to approximately match that of the starting position. Extracting the correct rest state automatically from measured values is a more difficult problem – requiring the identification of differences in muscle shape and joint angle that are inherent to the animal's structure from those that occur from a change in state. The animation was also adjusted to adapt for the differences in the structure of the different horses. A more automatic approach for retargeting the animation could be applied at this stage [Gleicher 1998].

The constructed model could also be utilized with other animation systems that make use of an underlying component hierarchy, with different muscle deformation and skinning solutions [Alias/Wavefront 2000].

6 Conclusion and Future Work

This paper has presented a new technique for automatic model reconstruction and demonstrated its use on capturing novel horse models by deforming a canonical horse model.
The contributions of this method are the following:

- This is a novel approach for automatically reconstructing anatomically based models of complete jointed creatures from a canonical model.
- Interpolation is successfully used both for morphing and for animation. Our combined anchoring/interpolation approach causes the skin to follow underlying components at important locations, and the smooth interpolation between those points avoids the artifacts (such as buckling) that occur on detailed skin meshes when all points are attached to underlying parts.
- The model is very flexible. It can be morphed and animated completely automatically, or the component hierarchy can be easily adapted both locally and globally to produce novel, perhaps fanciful creatures. New base models can be generated with different skin and muscle models.

Our goal is to extract not only the geometry and animation hierarchy automatically, but appearance and motion as well. While the current prototype system provides a step towards that goal, there are many aspects that need to be incorporated or improved on:

Kinematic Model The current model is not physically based, and therefore not suitable for quantitative analysis, such as might be required for rigorous biomechanical studies. Such a tool, however, could still provide an important educational and analysis tool for that area.

Canonical Model The canonical model could be made more realistic, if desired. We have utilized scans of scale models for the skin and bone representations. These could be replaced with scans of real bones, and a more realistic horse skin model. The current framework could be used with more physically based muscle, skin, and tissue models [Lee et al. 1995].

It could also be possible to automatically add in different skins to the canonical model and use the same types of techniques presented here, in the inverse direction: using the skin to automatically morph the underlying components. Schneider et al [Schneider and Wilhelms 1998] took a step in this direction, but their semi-automatic technique only addresses the anchoring of a new skin to an existing component hierarchy. It would be very useful to automatically produce underlying components that better represent the attached skin. We inserted the muscles into the canonical model by hand-and it is difficult to get them to fit the skin in a robust and effective way. To do the full inverse problem, it is necessary to have a consistent parameterization– here the mesh parameterization approach of Praun *et al.* [Praun et al. 2001] is applicable.

Muscle Model The muscle model based on deformable cylinders is somewhat restrictive and requires considerable human interaction during creation of the canonical model (muscles automatically adjust for morphed models, however).

Component Deformation The current heuristic for joint location estimation could benefit from results from biomechanical research [Back and Clayton 2000]. Anthropometric measurements could be used [DeCarlo et al. 1998] to improve the heuristics for component deformation based on the feature points, but such data is less readily available for non-human subjects.

Automatic Feature Point Extraction In this work we do not focus on the process of acquiring the initial marker points. It would be interesting to investigate approaches to automatically derive marker points from, for example, video.

Appearance Modeling For realistic appearance modeling, an additional "layer" must also be added to the representation to describe the appearance of the animal (e.g. the fur), and how the appearance changes under varying illumination and motion.

State Extraction For animation, it is important to derive a symmetrical rest state for the animal's components, given that the model was captured in a non-standard pose. It would then be possible to differentiate between feature properties that arise from the animal's base conformation or from state changes (e.g. the slope of the shoulder, or width of the neck muscle).

Acknowledgments

The authors would like to thank Marc Levoy and his group for the use of Stanford's scanning hardware and software – many thanks to David Koller and Steve Marschner for assisting with the scanning process itself. Thanks to Mark Slater for helping process the resulting meshes. This research was funded by NSF CSR 9972464.

References

ALIAS/WAVEFRONT, 2000. Maya: Character set-up. Alias/Wavefront.

BACK, W., AND CLAYTON, H. 2000. *Equine Locomotion*. W.B. Saunders, London.

BLANZ, V., AND VETTER, T. 1999. A morphable model for the synthesis of 3d faces. In *Proceedings of SIGGRAPH 1999*, Computer Graphics Proceedings, Annual Conference Series, ACM, 187–194.

CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, ACM, 67–76.

CHADWICK, J., HAUMANN, D., AND PARENT, R. 1989. Layered construction for deformable animated characters. In *Proceedings of SIGGRAPH 1989*, Computer Graphics Proceedings, Annual Conference Series, ACM, 243–252.

CHEN, D., AND ZELTZER, D. 1992. Pump it up: Computer animation of a biomechanically based model of muscle using the finite element method. In *Proceedings of SIGGRAPH 1992*, Computer Graphics Proceedings, Annual Conference Series, ACM, 89–98.

DECARLO, D., METAXAS, D., AND STONE, M. 1998. An anthropometric face model using variational techniques. In *Proceedings of SIGGRAPH 1998*, Computer Graphics Proceedings, Annual Conference Series, ACM, 67–74.

DELP, S. L., AND LOAN, J. P. 1995. A graphics-based software system to develop and analyze models of musculoskeletal structures. *Computers in Biology and Medicine 25(1)* (January), 21–34.

GLEICHER, M. 1998. Retargeting motion to new characters. In *Proceedings of SIGGRAPH 1998*, Computer Graphics Proceedings, Annual Conference Series, ACM, 33–42.

GOODY, P. C. 1983. *Horse Anatomy*. J. A. Allen and Co. Limited, London.

IMMERSION, 2000. MicroSribe-3D user's manual. Immersion 3D Capture.

KAKADIARIS, I., AND METAXAS, D. 1995. 3D human body model acquisition from multiple views. In *Proceedings of Fifth International Conference on Computer Vision*, 20–23.

KURIHARA, T., AND ARAI, K. 1991. A transformation method for modeling and animation of the human face from photographs. In *Proceedings of Computer Animation 1991*, 45–58.

LEE, Y., TERZOPOULOS, D., AND WATERS, K. 1995. Realistic modeling for facial animation. In *Proceedings of SIGGRAPH 1995*, Computer Graphics Proceedings, Annual Conference Series, ACM, 55–62.

LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINZTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., SHADE, J., AND FULK, D. 2000. The Digital Michelangelo Project: 3D scanning of large statues. In *Proceedings of SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, ACM, 131–144.

LEWIS, J., CORDNER, M., AND FONG, N. 2000. Pose space deformations: A unified approach to shape interpolation and skeleton-driven deformation. In *Proceedings of SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, ACM, 165–172.

NEDEL, L., AND THALMANN, D. 1998. Modeling and deformation of the human body using an anatomically based approach. In *Proceedings of Computer Animation 1998*.

NG-THOW-HING, V. 2001. *Anatomically-based models for physical and geometric reconstruction of humans and other animals*. PhD thesis, University of Toronto.

NIELSON, G. M. 1993. Scattered data modeling. *IEEE Computer Graphics and Applications 13*, 1, 60–70.

NOH, J.-Y., AND NEUMANN, U. 2001. Expression cloning. In *Proceedings of SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, ACM, 277–288.

PARENT, R. 2002. *Computer Animation: Algorithms and Techniques*. Morgan Kaufmann, San Francisco, CA.

PARKE, F., AND WATERS, K. 1996. *Computer Facial Animation*. A.K. Peters, Wellesley, MA.

PARKE, F. I. 1982. Parameterized models for facial animation. *IEEE Computer Graphics and Applications 2*, 9 (Nov.), 61–64, 66–68.

PIGHIN, F., HECKER, J., LISCHINSKI, D., SZELISKI, R., AND SALESIN, D. H. 1998. Synthesizing realistic facial expressions from photographs. In *Proceedings of SIGGRAPH 1998*, Computer Graphics Proceedings, Annual Conference Series, ACM, 75–84.

PRAUN, E., SWELDENS, W., AND SCHROEDER, P. 2001. Consistent mesh parameterizations. In *Proceedings of SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, ACM, 179–184.

SCHEEPERS, F., PARENT, R., CARLSON, W., AND MAY, S. 1997. Anatomy-based modeling of the human musculature. In *Proceedings of SIGGRAPH 1997*, Computer Graphics Proceedings, Annual Conference Series, ACM, 163–172.

SCHEEPERS, F. 1996. *Anatomy-based Surface Generation for Articulated Models of Human Figures*. PhD thesis, Ohio State University.

SCHNEIDER, P. J., AND WILHELMS, J. 1998. Hybrid anatomically based modeling of animals. In *Proceedings of Computer Animation 1998*, 161–169.

SIMMONS, M., WILHELMS, J., AND VAN GELDER, A. 2002. A canonical horse model for animation and morphing. Technical Report UCSC-CRL-02-24, University of California, Santa Cruz, Jack Baskin School of Engineering.

SLOAN, P.-P., ROSE, C. F., AND COHEN, M. F. 2001. Shape by example. In *Proceedings of 2001 Symposium on Interactive 3D Graphics*, ACM.

TURK, G., AND O'BRIEN, J. F. 1999. Shape transformation using variational implicit functions. In *Proceedings of SIGGRAPH 1999*, Computer Graphics Proceedings, Annual Conference Series, ACM, 335–342.

WILHELMS, J., AND VAN GELDER, A. 1997. Anatomically based modeling. In *Proceedings of SIGGRAPH 1997*, Computer Graphics Proceedings, Annual Conference Series, ACM, 173–180.

WILLIAMS, L. 1990. Performance-driven facial animation. In *Proceedings of SIGGRAPH 1990*, Computer Graphics Proceedings, Annual Conference Series, ACM, 235–242.

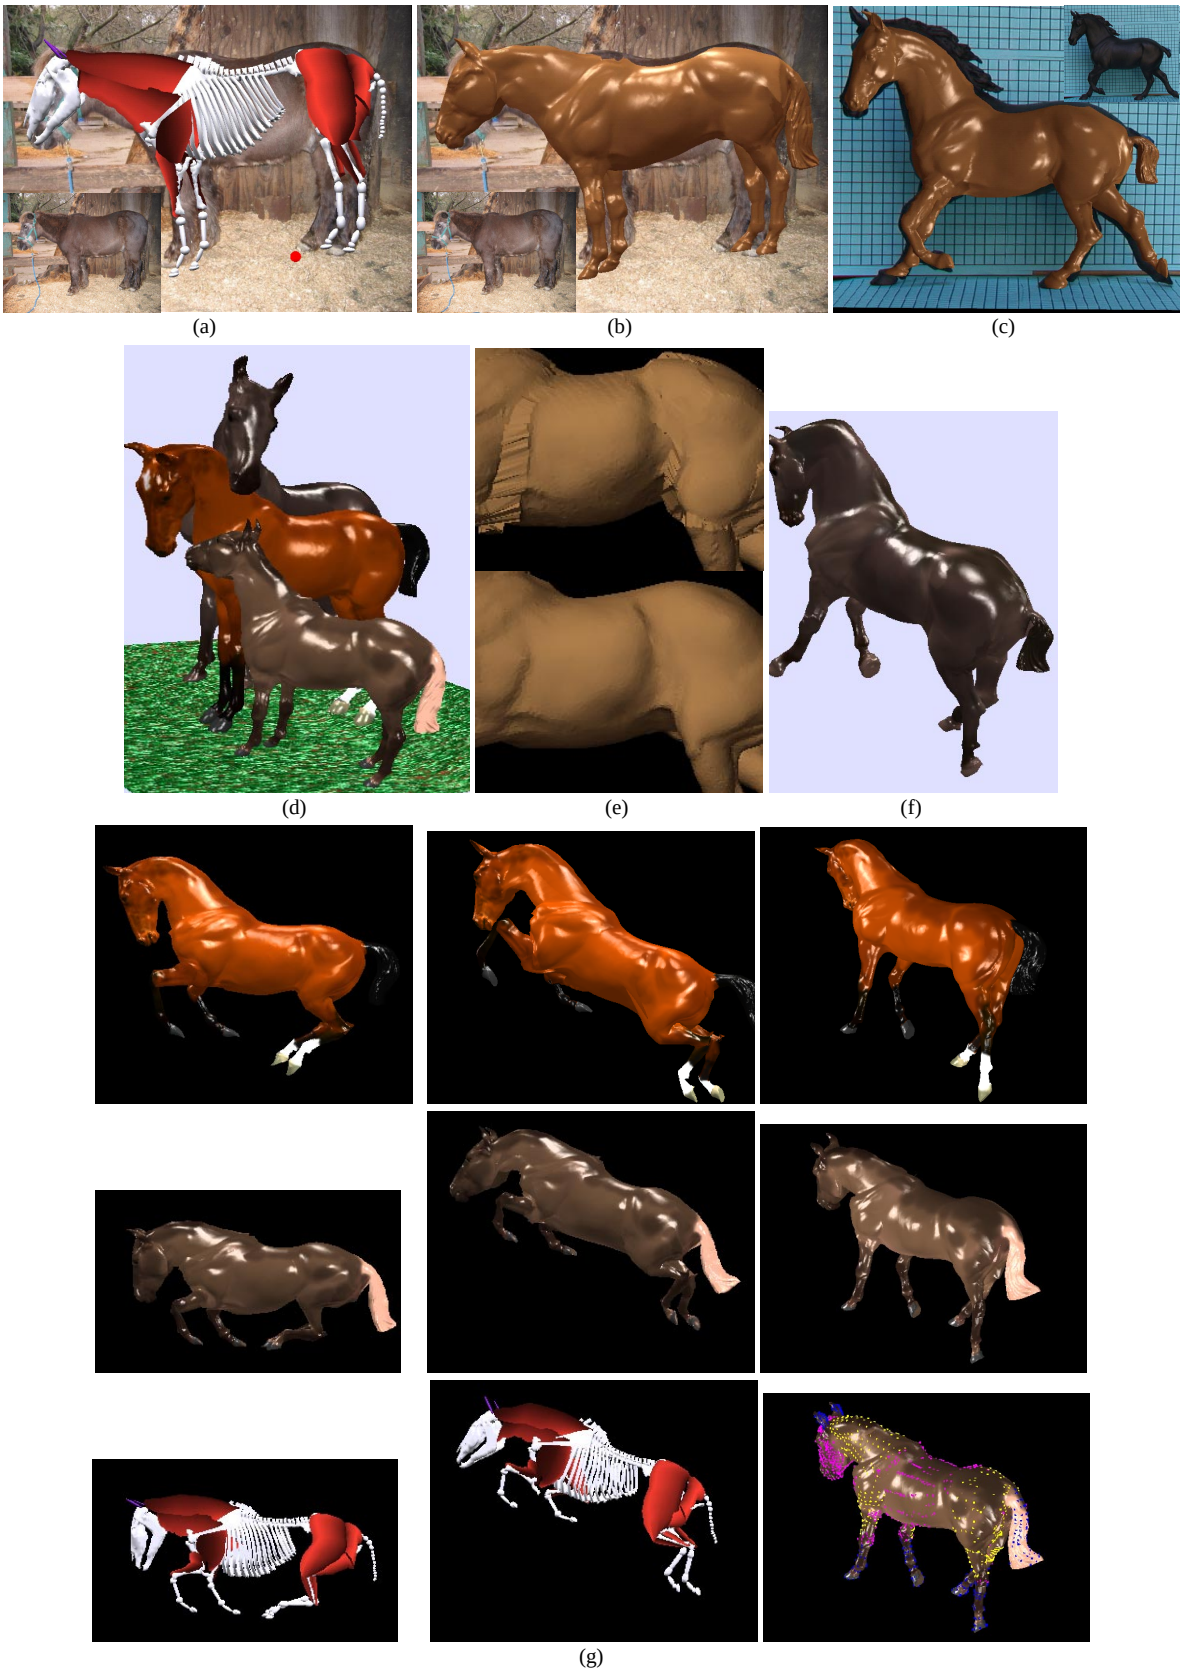

Figure 7: Results: a,b) Resulting components (a) and skin (b) for the pony. c) Resulting skin for the percheron (see Figure 5d for the components). d) Canonical horse and the two morphed examples shown together to give a sense of scale. (Note: the color of the horses' skin was done by hand-not automatically extracted). e) Interpolation: the top image shows the skin (flat-shaded) after morphing if interpolation is not performed. The bottom image shows the skin (flat-shaded) with interpolation. f,g) The rest of the images are examples from animating the models. f) shows the percheron at a walk. The top row of g) shows the canonical model, the middle row the pony: from the same steps in an animation sequence (the view is different). The final row shows the pony with components visualized for the first two examples, and the feature points for the third.