Physical Face Cloning

Bernd Bickel[1] Peter Kaufmann[1] Mélina Skouras[1,2] Bernhard Thomaszewski[1,2] Derek Bradley[1]
Thabo Beeler[1] Phil Jackson[3] Steve Marschner[4] Wojciech Matusik[1] Markus Gross[1,2]
[1]Disney Research Zurich [2]ETH Zurich [3]Walt Disney Imagineering [4]Cornell University

Figure 1: *Left: photograph and scanned 3D geometry of a human face. We use a physics-based optimization process to design the geometry of a synthetic skin in order to best match given target expressions. Right: final animatronic figure with fabricated skin.*

Abstract

We propose a complete process for designing, simulating, and fabricating synthetic skin for an animatronics character that mimics the face of a given subject and its expressions. The process starts with measuring the elastic properties of a material used to manufacture synthetic soft tissue. Given these measurements we use physics-based simulation to predict the behavior of a face when it is driven by the underlying robotic actuation. Next, we capture 3D facial expressions for a given target subject. As the key component of our process, we present a novel optimization scheme that determines the shape of the synthetic skin as well as the actuation parameters that provide the best match to the target expressions. We demonstrate this computational skin design by physically cloning a real human face onto an animatronics figure.

Keywords: animatronics, physics-based simulation, computational material design, facial animation, optimization.

Links: DL PDF WEB VIDEO

1 Introduction

We are naturally intrigued by the prospect of creating virtual humans in the likeness of ourselves — and it is not far-fetched to say that this is also a driving force for computer graphics research. While the latter strives to photorealistically recreate human characters on a computer screen, *animatronics* aims at creating physical robot characters that move and look like real humans. Both fields share many similarities; however, animatronics faces harder challenges since it has to deal with the physical constraints of real materials and actuators in addition to the virtual modeling component.

The human body consists of articulated rigid structures (bones) and soft tissue (e.g., flesh and skin). Therefore, it seems natural for animatronic characters to have a rigid articulated base and synthetic soft tissue. Many impressive characters have been created in this spirit, e.g., those in Disney World's *Hall of Presidents* or "Geminoids" [Nishio et al. 2007], androids that closely resemble human beings. However, creating such figures is still a difficult and labor-intensive process requiring manual work of skilled animators, material designers, and mechanical engineers. Owing to its expressive power, the human face is probably the most challenging part in this context. An animatronic character has to produce a vast range of facial expressions, each having different deformations and wrinkles. Manually designing the shape and material properties of a single skin that is able to achieve all these targets is clearly a formidable task.

The goal of this work is to automate this process, to increase the realism of the resulting character and, ultimately, to create an animatronic face that closely resembles a given human subject. In order to accomplish this task, we capitalize on recent developments from three areas in computer graphics: facial performance capture, physics-based simulation, and fabrication-oriented material design. Building on these foundations, we present a method of *physical face cloning* — a novel process for computational modeling, optimization and fabrication of synthetic skin for animatronic characters.

More specifically, our process comprises the following steps. First, we capture elastic material properties for a range of possible synthetic skin materials using a custom measurement system. We subsequently capture a collection of different expressions for a given target human face. As the central part of our pipeline, we then optimize the geometry of the skin and the actuation parameters of the underlying animatronics device to provide the best match to the target human face. Although our method applies quite generally to a broad range of possible animatronic devices, we validate the whole process in this paper by fabricating synthetic silicone skins for a specific articulated robot head. These skins are then animated, and the resulting shapes are compared both to the optimized model and to the real face.

The contributions of the research presented in this paper can be summarized as follows:

- A method to optimize for synthetic skin geometry and actuation parameters that provide the best match to a collection of given target expressions,
- A complete process for automating the physical reproduction of a human face on an animatronics device, including data acquisition, physical simulation, optimization, fabrication and validation,
- A comprehensive analysis and experimental validation of the process and its individual steps.

2 Related Work

Physical face cloning is a complex task that intersects with numerous research areas from computer graphics, applied mathematics, and material sciences. Although we are not aware of existing work to address this problem in its entirety, the individual components of our system have diverse connections to previous research.

Face Capture. Owing to the expressive power of the human face and our ability to infer a burst of information from the slightest motion, capturing, synthesizing and transferring facial expressions are important research topics in computer graphics. There exists a multitude of methods for capturing facial expressions, including traditional marker-based motion capture [Terzopoulos and Waters 1993; Bickel et al. 2007], structured light systems [Zhang et al. 2004; Wang et al. 2004], and passive markerless approaches [Beeler et al. 2010; Bradley et al. 2010; Beeler et al. 2011]. Our acquisition system uses the recent method of Beeler et al. [2011] to capture short performance sequences with high-resolution geometry and dense temporal correspondences.

Simulation. Introduced to the graphics community by Terzopoulos et al [1987], it has become common to model elastically deformable materials using continuum mechanics and finite elements. There exists a plethora of dedicated bio-mechanical models for different parts of the human body such as the face [Terzopoulos and Waters 1993; Koch et al. 1996; Sifakis et al. 2005], the hand [Gourret et al. 1989; Sueda et al. 2008], the torso [Lee and Terzopoulos 2006], various organs [Chentanez et al. 2009], or even the entire upper body [Lee et al. 2009]. Similarly, a large body of literature is concerned with modeling of biological soft tissue (see, e.g., the work by Teran et al. [2005] and references therein). We do not strive to reproduce the mechanical properties of biological tissue, but rather take an output-oriented approach: we fabricate synthetic skin using silicone rubbers, which are adequately modeled as elastic materials with isotropic and homogeneous properties. More specifically, we rely on a compressible neo-Hookean material model and, on the computational side, employ a finite element approach centered around linear tetrahedral elements.

Measuring and Fitting Material Parameters. Closely related to the choice of a constitutive law is the problem of finding a set of parameters that best approximates a given reference model. This reference data can originate either from an accurate computational model [Van Gelder 1998; Kharevych et al. 2009; Nesme et al. 2009] or, as in our case, from measurements of real-world objects. In the context of computer graphics, several methods have been described for fitting measured data to computational models, including linear-elastic materials [Becker and Teschner 2007], nonlinear viscoelastic soft tissue [Kauer et al. 2002], or nonlinear heterogeneous materials [Bickel et al. 2009]. Pai et al. [2001] describe a comprehensive system for capturing interactions with real-world objects, including deformation response, surface roughness and contact sounds. Our approach for determining parameters of different silicone samples goes along similar lines: we use a vision-based acquisition system to capture dense surface displacement fields and determine optimal material parameters by minimizing the displacement error of a corresponding finite element approximation.

Fabrication-oriented Design. As a natural step to follow acquisition and data-driven simulation, fabrication-oriented design is a recent trend that is seeing increasing attention in computer graphics. Several methods have been presented in this context, targeting the reproduction of geometric features of papercraft [Li et al. 2010] or plush toys [Mori and Igarashi 2007], general appearance properties [Weyrich et al. 2009; Hašan et al. 2010; Dong et al. 2010], and also the elastic properties of deformable media [Bickel et al. 2010]. Our approach for fabrication-oriented material design bears some similarities to this last reference, but there is also a crucial difference: whereas Bickel et al. determine spatially varying material properties for fixed geometry, we fabricate synthetic skins with uniform material properties and optimize their shape.

Shape Optimization. Engineering applications often require the shape of a structure to be optimal with respect to given conditions. Examples of such shape optimization problems include the computation of aerodynamic profiles, the design of reflectors for luminaires or the optimization of material interfaces. Such problems are typically cast as constrained optimization problems and solved numerically with finite elements (see, e.g., the textbook by Bucur and Buttazzo [2005] for an introduction). We optimize the shape of the synthetic skin by minimizing an elastic energy with respect to rest state positions, which is similar in spirit to methods for variational shape adaptation [Thoutireddy and Ortiz 2004]. But since it is generally difficult to control the adaptation of the undeformed configuration, we use a dedicated parametrization approach combined with a moving least squares interpolation technique that allows us to modify the shape only in its thickness direction.

Robotics & Animatronics. The majority of robots consist of only rigid mechanical structures. However, robots that are covered with soft-tissue materials such as rubbers or silicones have been built as well. Hara et al. [2001] was one of the first group to build an anthropomorphic head to study human-robot interaction. The face skin was made of silicone and it was driven by SMA (shape memory alloy) actuators. The head was able to reproduce six different expressions such as surprised, scared, angry, disgusted, sad and happy. Android research, which involves studying human-computer interaction, has been carried out by Intelligent Research Laboratory at Osaka University. The group has built a number of robotic characters with increasing capabilities. Their recent humanoid robot Geminoid HI-1 [Nishio et al. 2007] as well as it's successor Geminoid-DK have the realistic appearance of an adult man. These robots also have silicone-based skin driven by actuators placed underneath. KAIST and Hanson Robotics have developed a robot called Albert HUBO [Oh et al. 2006]. The robot's head resembles Albert Einstein and its underlying electro-mechanical actuation has 35 degrees of freedom. The skin is made of a sponge-like elastomer which is very soft and requires very little force to deform. In contrast to our approach, the design process for all these robotic skins is mainly manual. Shape variations are modeled by an artist, and the implications of changes to the shape have to be tested in a time consuming trial and error manner.

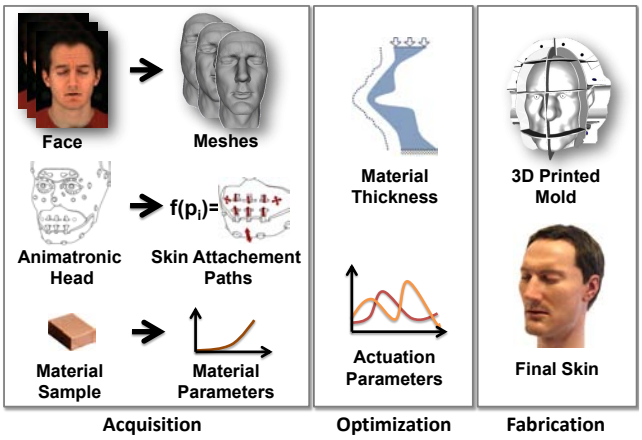

Figure 2: *Physical face cloning pipeline.*

3 Overview

Our approach to physical face cloning comprises modeling, simulation, design, and fabrication of synthetic soft tissue. The individual steps of our processing pipeline are illustrated in Fig. 2.

We start by capturing facial expressions of a human subject using an optical performance capture system that offers high-resolution 3D reconstructions including pores and wrinkles, as well as robust temporal correspondence (Sec. 4). This data provides information about the deformation behavior of the subject's skin. In addition, we also determine the operational range of the animatronic head, which consists of a set of skin attachment links actuated by motors. We attach to each moving link a marker, which allows us to sample a mapping of actuation parameters $\mathbf{p}_{\mathrm{act}}$ to 3D location and orientation of each link in space. Interpolating these samples yields a continuous function $m(\mathbf{p}_{\mathrm{act}})$ that describes the path of the skin attachment links (Sec. 7). Given the acquired example shapes of the human face and the description of the animatronic head, the goal is now to design synthetic soft tissue that matches the example shapes as closely as possible.

Our base material for the synthetic skin is silicone, which offers a wide range of stiffness that can be controlled by adjusting the concentration of plasticizer in the compound. In order to determine material parameters for the skin simulator, we numerically fit our computational model to experimentally acquired force-displacement samples of materials with different stiffnesses. We model and simulate the deformation behavior using a non-linear finite element method in combination with a neo-Hookean material (Sec. 5). The forward simulation allows accurate prediction of the deformed shape of a given synthetic skin and the resulting forces induced by the moving skin attachment links.

In order to design a synthetic skin that best matches a desired target, we introduce a novel optimization process. While the outer surface of the skin in the undeformed configuration is given by a 3D scan of the neutral target pose, we propose to vary the inner surface (thickness) of the skin to achieve spatially varying stiffness and deformation behavior. Furthermore, we also optimize the actuation parameters of the animatronic device to find the parameters that best resemble the individual target expressions. All these degrees of freedom are handled in one uniform optimization framework (Sec. 6).

We validate our simulation and optimization approach by fabricating various silicone objects using injection molding (Sec. 8). The molds are created using a rapid prototyping 3D printer. We then compare predicted simulation results and real-world behavior for different deformation scenarios (Sec. 9).

4 Measurements and Acquisition

We rely on real-world data for modeling facial geometry and expressions as well as material behavior. The following paragraphs briefly describe our techniques for data acquisition and integration.

Face Scanning. When cloning a human face, matching the shape of the face and how the skin deforms under different expressions is one of the main challenges. We must reconstruct various 3D expressions of the subject to be cloned, in order to seed the deformable skin optimization. Our approach is to capture a short performance of the target subject, leveraging the high resolution geometry acquisition method of Beeler et al. [2011]. The result is a detailed, compatible mesh sequence with explicit temporal correspondence, allowing us to analyze the deformation of a face under various expressions, down to the level of individual wrinkles. In our case, we chose 8 expressive poses from the sequence for optimizing the skin parameters.

Measuring & Fitting Material Parameters. Silicone rubber is an ideal base material for fabricating synthetic skin since it is easy to process and offers a vast range of stretch resistance. The animatronics device can only exert a limited force through its actuators, which imposes an upper bound on the net stiffness of the synthetic skin. Hence, we have to adjust the silicone composition such that the net stiffness of the synthetic skin complies to these constraints. To accomplish this task, we first measure the force-deformation behavior of a collection of silicone samples with different amounts of plasticizer. We then perform numerical optimization in order to determine the parameters of our computational model that best match the experimental data.

The measurement proceeds by pulling on small samples with controlled forces and capturing the resulting deformation over time using stereo reconstruction [Bradley et al. 2008]. With a set of applied loads and resulting surface displacements determined, we optimize the material parameters of our finite element solver in order to best approximate the measured stress-strain behavior. To this end, we start by defining an objective function $O(\mathbf{p}) = \sum_i |(\tilde{\mathbf{x}}_i - \mathbf{x}_i(\mathbf{p}))|^2$ that measures the difference between simulated surface positions $\mathbf{x}_i(\mathbf{p})$ and their closest corresponding point of the captured surface $\tilde{\mathbf{x}}_i$. The vector $\mathbf{p}$ holds the physical parameters of the material model, which can be identified with Young's modulus and Poisson's ratio (see Sec. 5). Given an initial guess $\mathbf{p}_0$ we iteratively compute updated parameters in a simulated annealing process. In each iteration, we compute the finite element solution $\mathbf{x}(\mathbf{p}_i)$ with the current parameters, evaluate the objective function $O(\mathbf{p}_i)$ and determine updated parameters $\mathbf{p}_{i+1}$. This process is repeated until the objective function signals convergence.

As a validation for this numerical-experimental material fitting process, we compare the stress-strain curves obtained from a standard testing procedure (ISO 37) to those obtained from a virtual counterpart of the same experiment. As can be seen from Fig. 4, we are able to match the real material behavior quite closely for deformations in the range of $[-20\%, +50\%]$. Considering the operating range of our animatronics device, this provides us with sufficiently large safety margins.

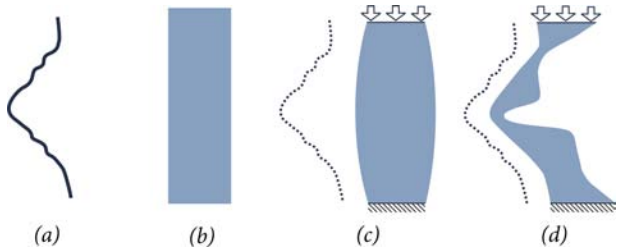

Figure 3: *Illustration of thickness optimization.* (a) *Target surface* (b) *undeformed skin patch* (c) *unoptimized skin patch deformed by vertical load* (d) *deformed skin patch with optimized thickness.*

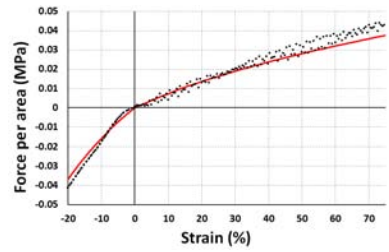

Figure 4: *Force-strain curve of real (black) and simulated (red) material.*

5 Physical Simulation of Synthetic Skin

As a central part of the face cloning pipeline, we need to have an accurate computational model for simulating deformations of a synthetic skin. Our approach is to model skin as a hyperelastic isotropic solid. Since we have to account for large rotations, stretching, and compression, we resort to finite-deformation continuum mechanics. Here we only give a brief overview and refer the reader to the textbook by Bonet and Wood [1997] for a comprehensive exposition.

Let $\mathbf{X}$ and $\mathbf{x}$ denote smooth functions describing the position of the skin in its undeformed and deformed state, respectively. Further, let $\varphi : \Omega \to \mathbb{R}^3$ denote the mapping that transforms material points from the undeformed to the deformed configuration as $\mathbf{x} = \varphi(\mathbf{X})$. The deformation of the skin at each point can be characterized by the deformation gradient $\mathbf{F}$ or, alternatively, the right Cauchy-Green tensor $\mathbf{C}$ which are defined as

$$\mathbf{F} = \frac{\partial \varphi}{\partial \mathbf{X}} \text{ , respectively } \quad \mathbf{C} = \mathbf{F}^T \mathbf{F} \text{ .}$$

For a hyperelastic material, the energy of the solid depends only on its current state of deformation, which is described in a rotation- and frame-invariant manner through the tensor $\mathbf{C}$. We denote the corresponding energy density function by $\Psi = \Psi(\mathbf{C})$. The relation between Ψ and $\mathbf{C}$ is described by material models, which are discussed next.

5.1 Material Model

Animating expressions on a virtual face generally induces both finite stretching and compression of the skin. While the simple St. Venant-Kirchhoff solid cannot be used in this setting [Irving et al. 2004], the non-linear neo-Hookean, Mooney-Rivlin, and Ogden material models are well suited for rubber-like substances such as silicone. Being the simplest representative of this class, the strain energy density of the compressible neo-Hookean material is given as

$$\Psi(\mathbf{C}) = \frac{\mu}{2}(J^{-\frac{2}{3}}\mathrm{tr}(\mathbf{C}) - 3) - \frac{\kappa}{2}(J-1)^2 \text{ ,} \tag{1}$$

where $J = \det \mathbf{F}$ and μ, κ are shear and bulk moduli, which are related to the familiar Young's modulus E and Poisson's ratio ν as

$$E = \frac{9\mu\kappa}{3\kappa + \mu} \text{ , respectively } \quad \nu = \frac{3\kappa - 2\mu}{2(3\kappa + \mu)} \text{ .} \tag{2}$$

The neo-Hookean material offers only two parameters, but for the range of deformations considered here, we found this model to be sufficiently accurate for obtaining good correspondence with the measured data. Having settled upon this material model, we determine the model parameters μ and κ for a given material sample using the numerical-experimental technique described in the previous section. Fig. 4 compares the force-strain curve for a measured and simulated material.

5.2 Discretization

The optimization process described in the next section requires the minimization of the elastic energy subject to boundary conditions. To this end, we discretize (1) in space using tetrahedral finite elements. Let $\mathbf{x}_a \in \mathbb{R}^3, 1 \leq a \leq n$ denote deformed positions of a tetrahedral mesh. We will use the superscript e to identify quantities pertaining to a given tetrahedral element. In particuar, we let $\mathbf{X}^e \in \mathbb{R}^{12}$ and $\mathbf{x}^e \in \mathbb{R}^{12}$ denote position vectors for a given element in its undeformed and deformed state, respectively.

We proceed by approximating the configuration mapping as $\varphi \approx \sum_a \mathbf{x}_a N_a$ where N_a are piece-wise linear basis functions associated with the nodes of the mesh. The discrete deformation gradient for each element e is then expressed as

$$\mathbf{F}^e = \mathbf{F}(\mathbf{x}^e, \mathbf{X}^e) = \sum_{a=1}^{4} \mathbf{x}_a^e \left(\frac{\partial N_a^e}{\partial \mathbf{X}} \right)^T \text{ ,} \tag{3}$$

where we emphasize the fact that for each element $\mathbf{F}^e$ is a function of its position in both undeformed and deformed configuration. All other discrete quantities required for computing the elastic energy of a deformed element, i.e. $\mathrm{tr}(\mathbf{C}^e)$ and J^e, follow directly from (3). Furthermore, we note that the deformation gradient is constant within each element since N_a^e are piece-wise linear functions. With this relation established, the discrete elastic energy of a deformed element W^e can be expressed as

$$W^e(\mathbf{x}^e, \mathbf{X}^e) = \int_{\Omega^e} \Psi(\mathbf{F}^e) = \Psi(\mathbf{F}^e) V^e(\mathbf{X}^e) \text{ ,} \tag{4}$$

where Ω^e is the element's parameter domain and V^e its undeformed volume. From now on, let $\mathbf{x} = (\mathbf{x}_1^T, \ldots, \mathbf{x}_n^T)^T$ and $\mathbf{X} = (\mathbf{X}_1^T, \ldots, \mathbf{X}_n^T)^T$ denote vectors containing all deformed and undeformed positions respectively. Summing up all elemental contributions, we obtain the total energy for a deformed configuration as

$$W(\mathbf{x}, \mathbf{X}) = \sum_e W^e(\mathbf{x}^e, \mathbf{X}^e) \text{ ,} \tag{5}$$

We can now use this energy in a static equilibrium problem in order to compute the deformation of the skin in response to actuator placements, which translate into a set of position constraints. The deformed configuration of the skin is then determined as the minimum of the total energy

$$\mathbf{x} = \underset{\mathbf{x}}{\mathrm{argmin}} \ \left(W(\mathbf{x}, \mathbf{X}) + W_{\mathrm{ext}}(\mathbf{x}) \right) \text{ ,} \tag{6}$$

where W_{ext} summarizes external contributions to the potential, e.g. due to gravity. Requiring the gradient of the energy to vanish at equilibrium leads to a set of nonlinear equations, which are solved in the standard way using Newton iterations with incremental loading and line search for convergence control.

By solving (6), we can simulate skin deformation in response to arbitrary actuator placements.

6 Optimization

Recall that our goal is to fabricate a synthetic skin which, when deformed by the underlying electromechanical base, matches the target expressions as closely as possible. In order to achieve this requirement, we optimize the shape of the skin as well as the actuation parameters of the base. This process is subject to a number of constraints. We cannot change the outside surface of the skin. We also cannot modify the attachment points of the actuators. Furthermore, the motion of the actuators is physically restricted.

In order to achieve the best results given these constraints, we have to expose as many degrees of freedom as possible to the optimization process. To this end, we optimize with respect to two sets of parameters. First, we optimize the actuation parameters of the animatronic device for every target expression. Second, we change the thickness of the skin by modifying its shape on the inside – we are free to do so everywhere except at the attachment points and in their immediate vicinity. Practically, this allows the skin to produce bulges and wrinkles where they did not appear before. We start by describing the general optimization framework and then how we apply it to the individual parameter sets.

6.1 Generic Optimization Framework

Let the function $\hat{W}(\hat{\mathbf{x}}, \mathbf{p})$ denote the total energy of our physical system, including internal deformation energies and external contributions to the potential. $\mathbf{p}$ is a vector of generic parameters for which we want to optimize. $\hat{\mathbf{x}}$ contains a subset of the deformed positions $\mathbf{x}$ of the physical simulation. All other values of $\mathbf{x}$ are either constant (constrained) or computable from $\mathbf{p}$. In the most general case, both the undeformed and deformed positions of our simulation mesh are allowed to depend on $\mathbf{p}$, and we can write

$$\hat{W}(\hat{\mathbf{x}}, \mathbf{p}) = W(\mathbf{x}(\hat{\mathbf{x}}, \mathbf{p}), \mathbf{X}(\mathbf{p})) + W_{\text{ext}}(\mathbf{x}(\hat{\mathbf{x}}, \mathbf{p})). \tag{7}$$

For fixed parameter values $\mathbf{p}$, the physical system will assume the equilibrium state

$$\hat{\mathbf{x}}_{\text{eq}}(\mathbf{p}) = \underset{\hat{\mathbf{x}}}{\operatorname{argmin}} \left(\hat{W}(\hat{\mathbf{x}}, \mathbf{p}) \right) \tag{8}$$

minimizing the total energy of the system, leading to the necessary condition of a vanishing residual force

$$\left. \frac{\partial \hat{W}}{\partial \hat{\mathbf{x}}} \right|_{\hat{\mathbf{x}}_{\text{eq}}, \mathbf{p}} = \mathbf{0}. \tag{9}$$

The goal of the optimization will be to find the optimal parameter values $\mathbf{p}$ such that the positions $\mathbf{x}(\hat{\mathbf{x}}_{\text{eq}}(\mathbf{p}), \mathbf{p})$ match a desired target configuration as closely as possible. This is measured in terms of a matching energy $E_{\text{match}}(\mathbf{x})$. We additionally regularize the parameters using a regularization energy $E_{\text{reg}}(\mathbf{p})$.

We solve this as a minimization problem with respect to $\hat{\mathbf{x}}$ and $\mathbf{p}$, with the following objective function:

$$E(\hat{\mathbf{x}}, \mathbf{p}) = \frac{1}{2}\gamma \left|\left| \left. \frac{\partial \hat{W}}{\partial \hat{\mathbf{x}}} \right|_{\hat{\mathbf{x}}, \mathbf{p}} \right|\right|^2 + E_{\text{match}}(\mathbf{x}(\hat{\mathbf{x}}, \mathbf{p})) + E_{\text{reg}}(\mathbf{p}) \tag{10}$$

The first term penalizes the violation of condition (9), ensuring with a sufficiently high penalty γ that the resulting $\hat{\mathbf{x}}$ is close to a physical solution with vanishing residual forces.

Matching Energy. The matching energy measures how well a configuration $\mathbf{x}$ is in agreement with what we want the object to look like. To this end, we embed a set of points in the simulation domain and deform them according to the mapping φ. In our implementation, we use a subset of points from high-resolution 3D scans (Sec. 4). Due to our choice of basis functions, the deformed position of a point can be computed as a linear combination of four nodal positions. We then compute the matching energy as the squared distance between the deformed points and their desired target positions $\mathbf{q}$,

$$E_{\text{match}}(\mathbf{x}) = \frac{1}{2}||\mathbf{S}\mathbf{x} - \mathbf{q}||^2 \tag{11}$$

where the matrix $\mathbf{S}$ contains the weights for the computation of the deformed points from $\mathbf{x}$.

Numerical Optimization. We minimize (10) using a Newton method. For this we need to compute its first and second derivatives with respect to $\hat{\mathbf{x}}$ and $\mathbf{p}$. While these computations are trivial for our particular choices of matching and regularization energies, we compute the second derivatives of the first term in (10) only approximately, ignoring the third order derivatives of W that would result from applying the chain rule. In each Newton step, we use a sparse direct solver [Schenk and Gärtner 2006] to solve the linear system

$$\mathbf{H} \begin{pmatrix} \Delta\hat{\mathbf{x}} \\ \Delta\mathbf{p} \end{pmatrix} = -\mathbf{f} \tag{12}$$

for the increments $\Delta\hat{\mathbf{x}}$ and $\Delta\mathbf{p}$, where $\mathbf{f}$ is the vector of first derivatives of E and $\mathbf{H}$ contains its second derivatives. We employ line search to make sure the energy decreases in each Newton step. It can happen that the matrix $\mathbf{H}$ becomes indefinite, a condition we can easily detect by the direct solver not being able to factorize $\mathbf{H}$, in which case we proceed as suggested in Nocedal and Wright [2000] by adding a multiple of the identity to $\mathbf{H}$.

6.2 Thickness Optimization

The goal of the thickness optimization is to modify the local thickness of the skin in such way that when the mechanical actuators of the robot are set to the values corresponding to a particular pose, the resulting deformation of the skin matches the pose's target positions $\mathbf{q}$ as closely as possible. In the physical simulation, the actuator settings result in hard positional constraints that can directly be applied to the corresponding deformed positions. The parameters $\mathbf{p}_{\text{thk}}$ determine the thickness distribution in the undeformed configuration without directly affecting the deformed positions, thus we can write

$$\hat{W}(\hat{\mathbf{x}}, \mathbf{p}_{\text{thk}}) = W(\hat{\mathbf{x}}, \mathbf{X}(\mathbf{p}_{\text{thk}})) + W_{\text{ext}}(\hat{\mathbf{x}}) \tag{13}$$

and minimize (10) to find the parameter values for the thickness distribution that best match the given pose.

Parameterizing the Undeformed State. Although the skin clearly exhibits a thickness direction, this information is not explicitly available from the 3D geometry such that we have to infer it first. We start by computing a low-pass filtered version of the skin's outer surface and construct a parameterization $\mathbf{r} : \mathbb{R}^2 \to \mathbb{R}^3$ using standard *uv*-mapping software. A roughly area-preserving map is desirable for obtaining unbiased sample distributions, but this is not a critical aspect in this context. The undeformed geometry of the skin can now be described as $\mathbf{X} = \mathbf{r}(u, v) + h\mathbf{n}(u, v)$, where u and v are surface coordinates and the thickness parameter h runs along

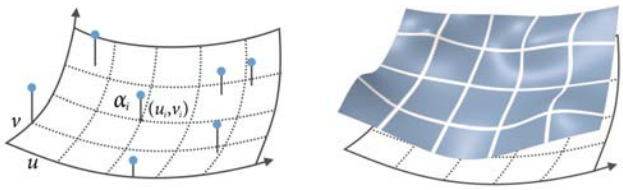

Figure 5: Left*: parametrization of the geometry and placement of MLS samples;* Right*: MLS warping field shown as surface with isoparametric lines for illustration.*

the surface's normal direction $\mathbf{n}$. Having defined the thickness in this way, we can now proceed to its optimization.

Instead of working directly on the undeformed positions, we define a *warping heightfield* on the surface in order to smoothly deform space along its normal direction. We construct a smooth warping field $w(u, v)$ from a sparse set of height samples α_i at positions $(u_i, v_i)^T$ using moving least squares (MLS) interpolation [Levin 1998]. Fig. 5 illustrates the concepts of parametrization and MLS interpolation, while Fig. 6 shows some shape variations that can be obtained for a simple one-dimensional example. To control the resolution of the warping heightfield, we adaptively place the locations of the height samples in uv space based on a user-provided density map [Clarberg et al. 2005].

Evaluating the MLS warping field at an arbitrary parameter point (u, v) amounts to fitting an affine transformation that minimizes the weighted least squares error

$$\sum_i \theta(||(u,v)^T - (u_i,v_i)^T||)||\mathbf{a}(u,v)^T(1,u_i,v_i) - \alpha_i||^2 \,, \quad (14)$$

where $\theta(x)$ is a weighting kernel and $\mathbf{a}(u, v) \in \mathbb{R}^3$ are the local coefficients of the interpolation function. The latter are determined by solving the normal equations for (14), which can be done efficiently using precomputations.

Finally, we gather the parameters α_i of all sample points into $\mathbf{p}_{\text{thk}}$ and compute optimal height values by minimizing (10). Note that the MLS interpolation results in a linear mapping between undeformed positions $\mathbf{X}$ and parameters $\mathbf{p}_{\text{thk}}$, thus the matrix $\partial\mathbf{X}/\partial\mathbf{p}_{\text{thk}}$ is constant and can be precomputed. Once the optimal α_i are determined, the final rest configuration is obtained in the following way: for each point $\mathbf{X}_a$ of the original mesh we first retrieve its parameter values (u_a, v_a, h_a), evaluate the MLS warping field to obtain $w_a = w(u_a, v_a)$ and finally compute the warped positions as $\mathbf{X}'_a = \mathbf{r}(u_a, v_a) + w_a h_a \mathbf{n}(u_a, v_a)$.

Regularization. We need the undeformed mesh with nodal positions $\mathbf{X}(\mathbf{p}_{\text{thk}})$ to be a valid tetrahedral mesh where every element has positive volume. While there are different ways to formulate and enforce this constraint, we found it most practical to reuse the deformation energy function W, abusing $\mathbf{X}(\mathbf{p}_{\text{thk}})$ as the *deformed* configuration and the initial undeformed positions $\mathbf{X}(\mathbf{1})$ as the undeformed configuration:

$$E_{\text{reg}}(\mathbf{p}_{\text{thk}}) = \gamma_{\text{undefreg}} W(\mathbf{X}(\mathbf{p}_{\text{thk}}), \mathbf{X}(\mathbf{1})) \quad (15)$$

For any positive penalty value γ_{undefreg}, this energy will tend to infinity as any element's volume approaches zero, effectively preventing element inversions in the undeformed configuration.

Fabrication Constraints. Additionally, we add a quadratic energy term to E_{reg} that penalizes MLS weights exceeding a value of one, or falling below a user-defined minimum thickness. This guarantees a fabricatable skin that does not intersect the underlying mechanical structure.

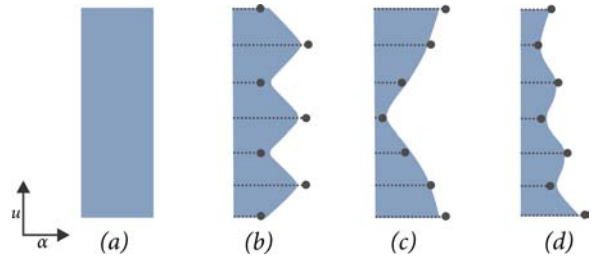

Figure 6: *Illustration of MLS interpolation using 7 sample points.* (a) *Input geometry* (b)–(d) *examples of warped geometry.*

Multiple Pose Optimization. We want to optimize the undeformed configuration $\mathbf{X}(\mathbf{p}_{\text{thk}})$ such that $m > 1$ poses can be reproduced as closely as possible: for a set of actuator settings $\mathbf{a}_i$, the given target positions $\mathbf{q}_i$ should be matched as closely as possible. In this case our optimization is slightly modified and we optimize for all deformed positions $\hat{\mathbf{x}}_1, \ldots, \hat{\mathbf{x}}_m$ and $\mathbf{p}_{\text{thk}}$ at once. We have m matching functions E^i_{match} and the objective function can be written as

$$E(\hat{\mathbf{x}}, \mathbf{p}_{\text{thk}}) = \sum_{i=1}^{m} \frac{1}{2}\gamma \left|\left| \frac{\partial \hat{W}}{\partial \hat{\mathbf{x}}}\bigg|_{\hat{\mathbf{x}}_i, \mathbf{p}_{\text{thk}}} \right|\right|^2 \quad (16)$$

$$+ \sum_{i=1}^{m} E^i_{\text{match}}(\mathbf{x}(\hat{\mathbf{x}}_i, \mathbf{p}_{\text{thk}})) \quad (17)$$

$$+ E_{\text{reg}}(\mathbf{p}_{\text{thk}}). \quad (18)$$

We initialize the optimization with values $\hat{\mathbf{x}}_i$ computed from a per-pose forward simulation with incremental loading, and use $\mathbf{p}_{\text{thk}} = \mathbf{1}$ for the MLS parameters.

6.3 Actuation Parameter Optimization

The goal of the actuation parameter optimization is to find the best values for controlling the mechanical actuators such that the resulting deformed skin matches a given target pose as closely as possible. In this setting, the parameters $\mathbf{p}_{\text{act}}$ of the optimization define the positional constraints to apply to the nodal positions $\mathbf{x}$. In the objective function (10) we use

$$\hat{W}(\hat{\mathbf{x}}, \mathbf{p}_{\text{act}}) = W(\mathbf{x}(\hat{\mathbf{x}}, \mathbf{p}_{\text{act}}), \mathbf{X}) + W_{\text{ext}}(\mathbf{x}(\hat{\mathbf{x}}, \mathbf{p}_{\text{act}})). \quad (19)$$

Parameterization. The mapping from unconstrained nodal positions $\hat{\mathbf{x}}$ and parameters $\mathbf{p}_{\text{act}}$ to deformed positions $\mathbf{x}$ is

$$\mathbf{x}_i = \begin{cases} \hat{\mathbf{x}}_i & \text{if node } i \text{ unconstrained} \\ \mathbf{M}_a(\mathbf{p}_{\text{act}})\mathbf{X}_i & \text{if node } i \text{ constrained by actuator } a. \end{cases} \quad (20)$$

For the sake of simplicity, we linearize the transformation matrix $\mathbf{M}_a(\mathbf{p}_{\text{act}})$ of actuator a around the current parameter values $\mathbf{p}_{\text{act}}$ in each step of the optimization.

Regularization. Similar to the thickness optimization, we use the regularization energy E_{reg} to keep the actuator parameter values within the range of physically feasible values.

7 Animatronic Head

To actuate the silicone skin we use a proprietary animatronic head developed by Walt Disney Imagineering. It is driven by electric motors and features $a = 13$ parameters $\mathbf{p}_{\text{act}} \in \mathbb{R}^a$ to control the movement of the attachment links indicated in red in the figure on the left. The length of the red arrows illustrates the range of potential motion. Furthermore, there are several rigid attachment links that constrain the movement of the skin. We refer to the video for a visualization of the actuation level of each actuation parameter for a facial animation sequence, providing an intuition of which expressions are inside or at the border of the gamut of reproducible poses. We treat the underlying mechanics and controlling mechanism as a black box, so in theory our optimization algorithm can be used with any animatronics device.

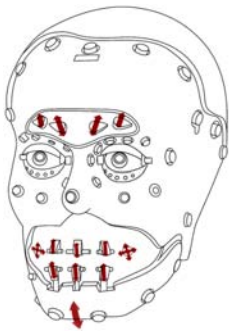

To determine the operating range and control parameters of the head we attach a marker to each moving link and track its position under a number of sample locations and orientations. Tracking is performed using the same capture system as used for face scanning (Sec. 4). The mapping $m(\mathbf{p}_{\text{act}}) : \mathbb{R}^a \rightarrow \mathbb{R}^{l \cdot 6}$ from animation parameters $\mathbf{p}_{\text{act}}$ to 3D locations and orientations of attachment points is then specified by quadratically interpolating the samples. These are then used as hard boundary constraints in our forward simulation. To obtain the specific animation parameters for a given facial expression, we solve the inverse problem as described in Sec. 6.3.

8 Fabrication

Our choice to use silicone is based on the fact that the stiffness can be controlled accurately, the color can be adjusted, it is robust, and the fabrication process is easy and safe. Silicone belongs to the class of elastically deformable polymers with a high yield strength. After having optimized the ideal skin geometry, we use liquid injection molding to fabricate the physical skin as shown in the accompanying video. Our mold consists of a core that resembles the inner surface of the skin and six outer parts. Note that multiple outer parts are required so that the skin can be removed without destroying the mold. The mold itself is fabricated using a rapid prototyping 3D printer. For our experiments, we used the two-part silicone GI-245 by Silicones Inc. After mixing it with the corresponding catalyst, it has a liquid consistency and can be injected into the mold. Approximately seven days of room-temperature curing is required until it is ready to be used.

9 Results

In this section we present our results, starting with validation of our pipeline and optimization process, and then showcasing our method with an example of real physical face cloning for an animatronic figure.

Validation Simulation. We validated our pipeline for measuring, modeling, and simulating silicone materials. Fig. 4 shows a comparison of the force-strain curve of a simulated material to a measured sample. The measurements are acquired by running an uniaxial stretch and compression test on a rectangular ($55 \times 5 \times 2$ mm — stretch) and cylindrical (50 mm diameter, 12 mm height — compression) object. Furthermore, we validate the simulation by comparing simulated deformations of our silicone skin with 3D scanned poses of the actual animatronic figure in Fig. 9, showing an average error of less than 1.35 mm.

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

Figure 7: *Volumetric shape optimization. Starting with a volumetric object (a) and a target surface (b), our optimization process computes the thickness variation (c) such that the surface matches the deformation of the target under actuation (d). Without optimization (e), the target is not well approximated.*

Validation Optimization. Our optimization process computes the volumetric shape of soft-tissue material under actuation such that the surface matches the deformation of given target surfaces. We validated the optimization process in simulation. We start by generating ground truth data by randomly varying the thickness of a uniform block using the moving least squares warp as described in Sec. 6.2. From the forward simulation of these objects we extract target surfaces. We then used these target surfaces and the undeformed block as input into our optimization — see e.g. Fig. 7. In all experiments, our algorithm was able to exactly reproduce the spatially-varying thickness of the ground truth objects that were used to generate the inputs. As shown in Fig. 8, we also fabricated a subset of these examples, and compared the deformation of the predicted optimized shape to the real object. For these 95mm long bars, the average errors between the deformed fabricated surface and the goal shape are 0.44mm (top) and 0.38mm (bottom).

Face Cloning. We showcase our complete pipeline by physically cloning a human face for an animatronic figure. We started by acquiring more than 100 frames of an expressive performance of our subject. We then aligned the neutral pose to a scan of the animatronic head. Because the animatronic head has very pronounced cheek bones, we slightly warp the acquired 3D geometry of the human face to prevent intersections. Consequently, we apply the same warping transformation to the captured geometry before doing comparisons (such as the error plots in Fig. 10) with our results. The gap between the surface of the neutral pose and the underlying animatronics defines the initial unoptimized skin, which we represent as a tetrahedral mesh consisting of ≈ 63k elements. To better represent important facial deformation features on the forehead, we additionally use a high-resolution mesh of ≈ 157k elements to optimize this area. We then use our optimization process to obtain animation parameters for the entire captured face performance sequence. As matching criteria we select small surface regions close to the attachment links. In those regions, the surface thickness is constrained and cannot be changed because the skin is connected to the moving and static links. For all other regions we use our optimization process to achieve the desired deformation behavior by varying the skin thickness. In our experiment, we used a subset of 8 expressive poses for thickness optimization. Fig. 10 illustrates the unoptimized and optimized geometry and provides a side-by-side comparison of the deformation behavior.

Finally, we also fabricated the optimized skin, using material with Young's modulus of $E = 78$kPa and Poisson's ratio of $\nu = 0.47$ in its cured state. Figure finishing included industry-standard techniques for adding hair and facial stubbles and painting a texture by a professional makeup artist. Fig. 9 shows our figure in action and compares it to the human face and the simulation under vari-

ous facial expressions. We also scanned these poses of the animatronic figure and provide error visualizations showing the distance between the simulated surface and the scanned surface. As can be seen from the error plots, the agreement between the shape predicted by our simulation and the actual observed shape of the actuated skin is generally good, with less than 4mm separating the surfaces everywhere. This shows that our simulation is a good model for the real synthetic face. The difference between the scans of the real face and the simulation is larger, representing the limitations in the range of motion of the robot head and the range of deformation of the attached skin. Comparing the human face to the cloned synthetic face, the large scale deformations and general expressions are captured well, and the differences show where improvements should be made in the physical system to enable the fitting methods to achieve a better match. For further comparisons we refer to the accompanied video.

10 Conclusion

In this paper we have described a process for computationally guided design of an animatronic character composed of a soft-tissue material and an electromechanical base. We compute a volumetric shape of the soft-tissue material and a set of actuation forces that can reproduce target deformations of an exterior character surface. As an example, this process allows us to construct an animatronic head that approximates the shape and behavior of a real person whose expressions are acquired using a 3D scanning system. The core requirement of this process is a physically based simulation framework that lets us accurately predict the behavior of the soft tissue when subject to external forces. Our experimental validation shows that using this process provides a principled way to design soft-tissue animatronic characters.

Limitations and Future Work. We believe that our computationally guided process for designing and fabricating animatronic characters will serve as a blueprint for how future robots with soft tissue should be built. Because this work is the first in the area, it has a number of limitations that should be addressed in future research. First, in our current framework the attachment of the actuators is fixed to specific places underneath the character's skin. Relaxing this constraint would likely give us more degrees of freedom, e.g., the rest shape of the soft tissue would not need to be the same as the rest shape of the electromechanical base. Second, we currently only use a single soft tissue material to produce a synthetic skin. We believe we can extend the range of possible deformations by using multi-layered materials similarly to Bickel et al. [2010]. Moreover, it is important to note that the motions of current animatronic characters are still limited as they are less expressive than most humans. Thus, it is important to develop retargeting algorithms that can map the expressions of a target human to the physical constraints of the electromechanical base and materials used. In this paper, we have not addressed the appearance aspect of the synthetic skin tissue including reflectance, subsurface scattering, and hair. This is also a very interesting direction for future work and we envision that similar computationally guided processes could be developed to solve these problems. Finally, while in this work we have only shown examples of replicating human faces, we predict that similar frameworks will be used to design realistic full-body characters.

Acknowledgements

We would like to thank the anonymous reviewers for their comments, Brian Tye, Frank Mezzatesta, and Walt Disney Imagineering for their support, and Lee Romaire for figure finishing. This work was supported in part by the NCCR Co-Me of the Swiss NSF.

References

BECKER, M., AND TESCHNER, M. 2007. Robust and efficient estimation of elasticity parameters using the linear finite element method. In *SimVis*, 15–28.

BEELER, T., BICKEL, B., BEARDSLEY, P., SUMNER, B., AND GROSS, M. 2010. High-quality single-shot capture of facial geometry. *ACM Trans. Graph. 29*, 4 (July), 40:1–40:9.

BEELER, T., HAHN, F., BRADLEY, D., BICKEL, B., BEARDSLEY, P., GOTSMAN, C., SUMNER, R. W., AND GROSS, M. 2011. High-quality passive facial performance capture using anchor frames. *ACM Trans. Graph. 30*, 4 (August), 75:1–75:10.

BICKEL, B., BOTSCH, M., ANGST, R., MATUSIK, W., OTADUY, M., PFISTER, H., AND GROSS, M. 2007. Multi-scale capture of facial geometry and motion. *ACM Trans. Graph. 26*, 3 (July), 33:1–33:10.

BICKEL, B., BÄCHER, M., OTADUY, M. A., MATUSIK, W., PFISTER, H., AND GROSS, M. 2009. Capture and modeling of non-linear heterogeneous soft tissue. *ACM Trans. Graph. 28*, 3 (July), 89:1–89:9.

BICKEL, B., BÄCHER, M., OTADUY, M. A., LEE, H. R., PFISTER, H., GROSS, M., AND MATUSIK, W. 2010. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph. 29*, 4 (July), 63:1–63:10.

BONET, J., AND WOOD, R. D. 1997. *Nonlinear Continuum Mechanics for Finite Element Analysis*. Cambridge Uni. Press.

BRADLEY, D., BOUBEKEUR, T., AND HEIDRICH, W. 2008. Accurate multi-view reconstruction using robust binocular stereo and surface meshing. In *Proc. CVPR*.

BRADLEY, D., HEIDRICH, W., POPA, T., AND SHEFFER, A. 2010. High resolution passive facial performance capture. *ACM Trans. Graph. 29*, 4 (July), 41:1–41:10.

BUCUR, D., AND BUTTAZZO, G. 2005. *Variational Methods in Shape Optimization Problems*. Birkhuser Mathematics.

CHENTANEZ, N., ALTEROVITZ, R., RITCHIE, D., CHO, L., HAUSER, K. K., GOLDBERG, K., SHEWCHUK, J. R., AND O'BRIEN, J. F. 2009. Interactive simulation of surgical needle insertion and steering. *ACM Trans. Graph. 28*, 3 (July), 88:1–88:10.

CLARBERG, P., JAROSZ, W., AKENINE-MÖLLER, T., AND JENSEN, H. W. 2005. Wavelet importance sampling: efficiently evaluating products of complex functions. *ACM Trans. Graph. 24*, 3 (Aug.), 1166–1175.

DONG, Y., WANG, J., PELLACINI, F., TONG, X., AND GUO, B. 2010. Fabricating spatially-varying subsurface scattering. *ACM Trans. Graph. 29* (July), 62:1–62:10.

GOURRET, J.-P., THALMANN, N. M., AND THALMANN, D. 1989. Simulation of object and human skin deformations in a grasping task. In *Comp. Graph. (Proc. SIGGRAPH)*, 21–30.

HARA, F., AKAZAWA, H., AND KOBAYASHI, H. 2001. Realistic facial expressions by sma driven face robot. In *Proc. of IEEE International Workshop on Robot and Human Interactive Communication*, 504 –511.

HAŠAN, M., FUCHS, M., MATUSIK, W., PFISTER, H., AND RUSINKIEWICZ, S. 2010. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. Graph. 29* (July), 61:1–61:10.

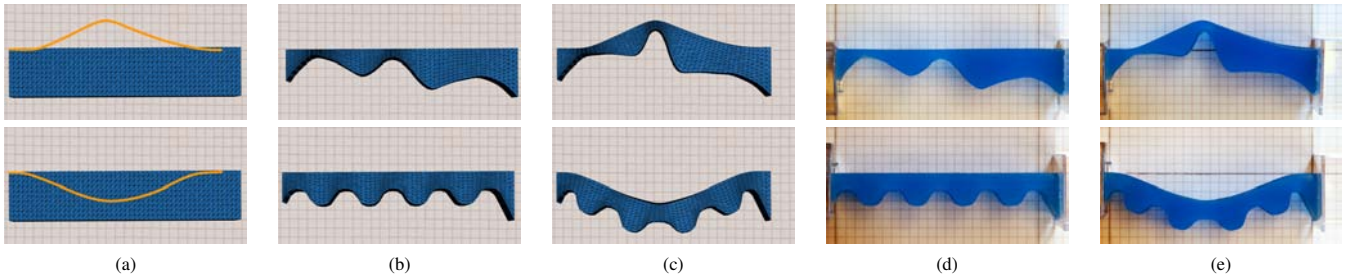

Figure 8: *Illustrating the volumetric shape optimization in 1D. Given a volumetric object and a target surface (indicated in orange) as shown in column (a), our optimization process computes the volumetric shape (b) such that the surface matches the deformation of the target under actuation (c). Columns (d) and (e) show the undeformed and deformed state of the corresponding fabricated objects against a 4mm grid.*

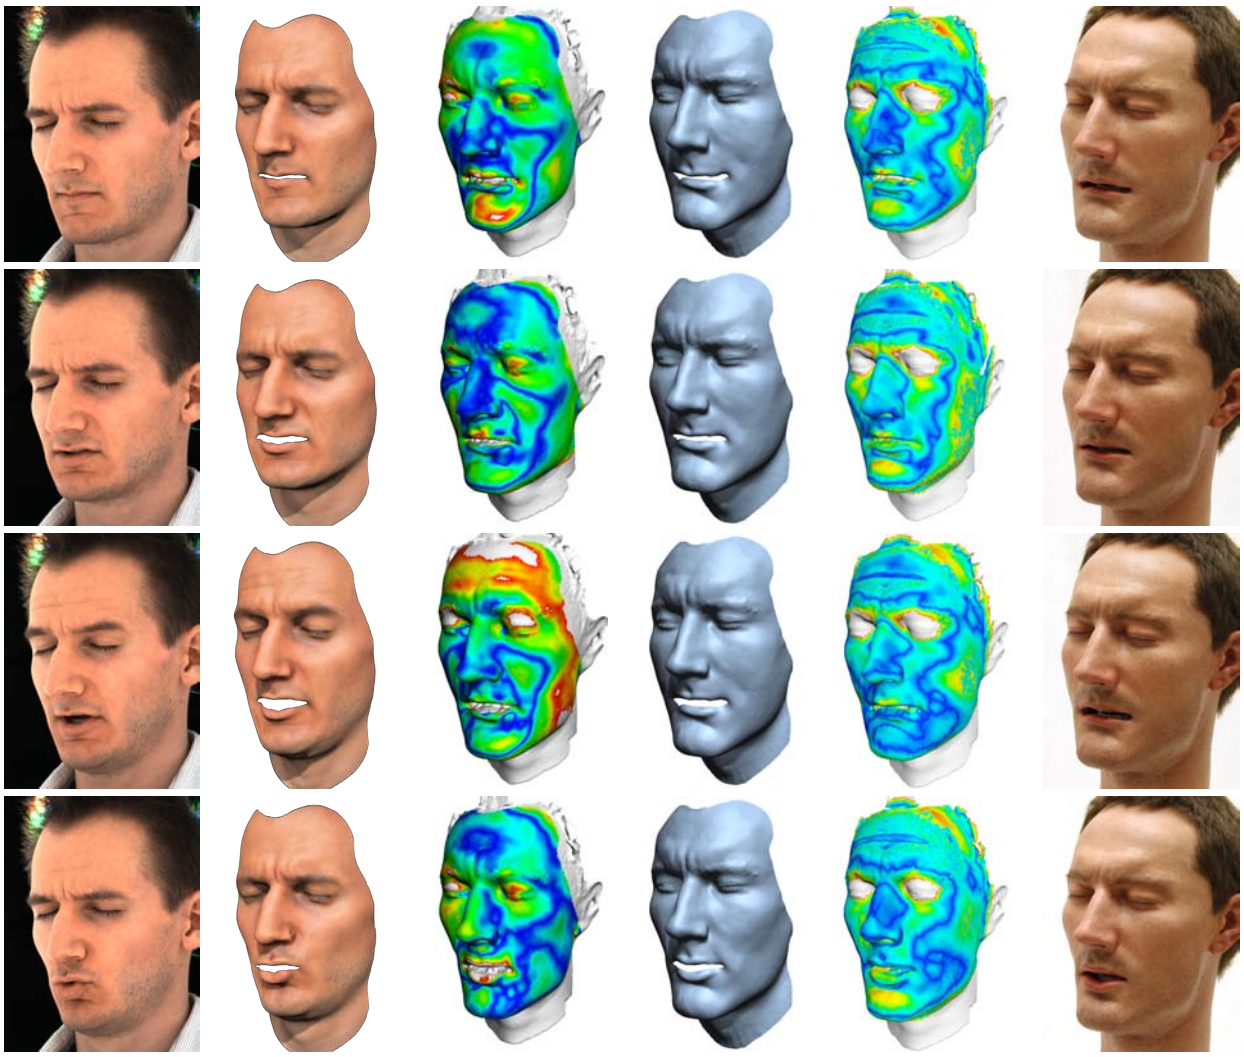

Figure 9: *Side-by-side comparison. From left to right: Photographs of captured facial expressions; Renderings of 3D reconstructions; error comparison of 3D reconstructions to the fabricated skin (blue 0mm, red 4mm); 3D simulation of the animatronic head for the corresponding pose; error comparison of our simulation to the fabricated skin (blue 0mm, red 4mm); expressions generated with the optimized fabricated skin actuated by an animatronic device.*

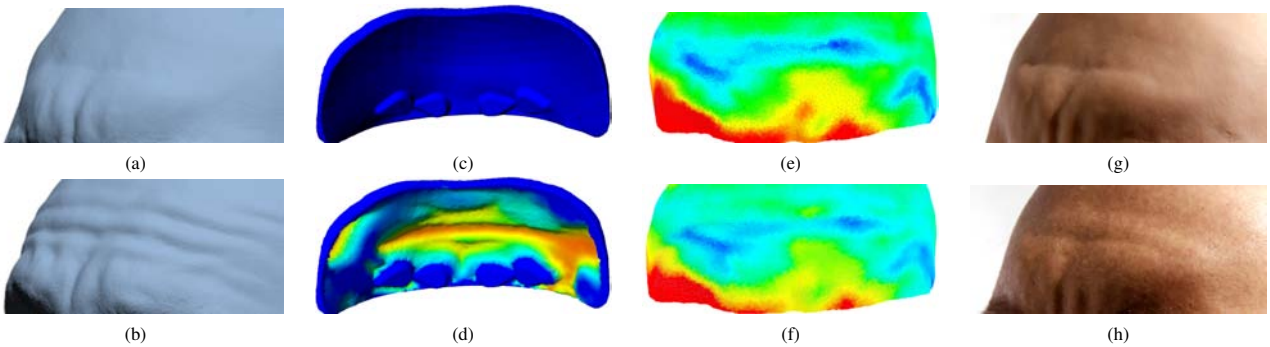

Figure 10: *Forehead Optimization: (a) Rest shape. (b) Target geometry. (b,c) Rendering of the inner side of the unoptimized and optimized forehead with color-coded thickness variation. Comparison of the unoptimized and optimized skin: Error plot of the (e) unoptimized and (f) optimized skin (blue 0mm, red 4mm). Although the variations are subtle, a comparison of the fabricated unoptimized (g) and optimized (h) skin shows visible differences.*

IRVING, G., TERAN, J., AND FEDKIW, R. 2004. Invertible finite elements for robust simulation of large deformation. In *2004 ACM SIGGRAPH / Eurographics SCA*, 131–140.

KAUER, M., VUSKOVIC, V., DUAL, J., SZEKELY, G., AND BAJKA, M. 2002. Inverse finite element characterization of soft tissues. *Medical Image Analysis 6*, 3, 257–287.

KHAREVYCH, L., MULLEN, P., OWHADI, H., AND DESBRUN, M. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Trans. Graph. 28* (July), 51:1–51:8.

KOCH, R. M., GROSS, M. H., CARLS, F. R., VON BÜREN, D. F., FANKHAUSER, G., AND PARISH, Y. I. H. 1996. Simulating facial surgery using finite element models. In *Proc. of Comp. graph. and int. tech.*, ACM, SIGGRAPH '96, 421–428.

LEE, S.-H., AND TERZOPOULOS, D. 2006. Heads up!: biomechanical modeling and neuromuscular control of the neck. In *ACM SIGGRAPH 2006 Papers*, ACM, 1188–1198.

LEE, S.-H., SIFAKIS, E., AND TERZOPOULOS, D. 2009. Comprehensive biomechanical modeling and simulation of the upper body. *ACM Trans. Graph. 28* (September), 99:1–99:17.

LEVIN, D. 1998. The approximation power of moving least-squares. *Math. Comput. 67* (October), 1517–1531.

LI, X.-Y., SHEN, C.-H., HUANG, S.-S., JU, T., AND HU, S.-M. 2010. Popup: automatic paper architectures from 3d models. *ACM Trans. Graph. 29* (July), 111:1–111:9.

MORI, Y., AND IGARASHI, T. 2007. Plushie: an interactive design system for plush toys. In *ACM SIGGRAPH 2007 papers*, ACM, SIGGRAPH '07.

NESME, M., KRY, P. G., JEŘÁBKOVÁ, L., AND FAURE, F. 2009. Preserving topology and elasticity for embedded deformable models. In *ACM SIGGRAPH 2009 papers*, ACM, 52:1–52:9.

NISHIO, S., ISHIGURO, H., AND HAGITA, N. 2007. *Humanoid Robots: New Developments*. I-Tech, ch. Geminoid: Teleoperated Android of an Existing Person.

NOCEDAL, J., AND WRIGHT, S. J. 2000. *Numerical Optimization*. Springer, August.

OH, J.-H., HANSON, D., KIM, W.-S., HAN, I. Y., KIM, J.-Y., AND PARK, I.-W. 2006. Design of android type humanoid robot albert hubo. In *Proc. of IEEE/RSJ Int. Conference on Intelligent Robots and Systems*, 1428–1433.

PAI, D. K., DOEL, K. V. D., JAMES, D. L., LANG, J., LLOYD, J. E., RICHMOND, J. L., AND YAU, S. H. 2001. Scanning physical interaction behavior of 3d objects. In *Proc. of Comp. graph. and int. tech.*, ACM, SIGGRAPH '01, 87–96.

SCHENK, O., AND GÄRTNER, K. 2006. On fast factorization pivoting methods for symmetric indefinite systems. In *Elec. Trans. Numer.*, no. 23, 158–179.

SIFAKIS, E., NEVEROV, I., AND FEDKIW, R. 2005. Automatic determination of facial muscle activations from sparse motion capture marker data. In *ACM SIGGRAPH 2005*, 417–425.

SUEDA, S., KAUFMAN, A., AND PAI, D. K. 2008. Musculotendon simulation for hand animation. *ACM Trans. Graph. (Proc. SIGGRAPH) 27*, 3.

TERAN, J., SIFAKIS, E., BLEMKER, S. S., NG-THOW-HING, V., LAU, C., AND FEDKIW, R. 2005. Creating and simulating skeletal muscle from the visible human data set. *IEEE Trans. on Vis. and Comp. Graph. 11*, 3, 317–328.

TERZOPOULOS, D., AND WATERS, K. 1993. Analysis and synthesis of facial image sequences using physical and anatomical models. *IEEE Trans. Pattern Anal. Mach. Intell. 15* (June), 569–579.

TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K. 1987. Elastically deformable models. In *Proc. Comp. graph. and int. tech.*, ACM, SIGGRAPH '87, 205–214.

THOUTIREDDY, P., AND ORTIZ, M. 2004. A variational r-adaption and shape-optimization method for finite-deformation elasticity. *Int. J. Numer. Meth. Engng. 61*, 1–21.

VAN GELDER, A. 1998. Approximate simulation of elastic membranes by triangulated spring meshes. *J. Graph. Tools 3*, 2, 21–42.

WANG, Y., HUANG, X., LEE, C.-S., ZHANG, S., LI, Z., SAMARAS, D., METAXAS, D., ELGAMMAL, A., AND HUANG, P. 2004. High resolution acquisition, learning and transfer of dyanmic 3-d facial expressions. In *Comp. Graph. Forum*, 677–686.

WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Trans. Graph. 28* (July), 32:1–32:6.

ZHANG, L., SNAVELY, N., CURLESS, B., AND SEITZ, S. M. 2004. Spacetime faces: high resolution capture for modeling and animation. *ACM Transactions on Graphics 23*, 3 (Aug.), 548–558.