

Building Accurate Physics-based Face Models from Data

PETR KADLEČEK, University of Utah, USA and Charles University, Czech Republic
LADISLAV KAVAN, University of Utah, USA

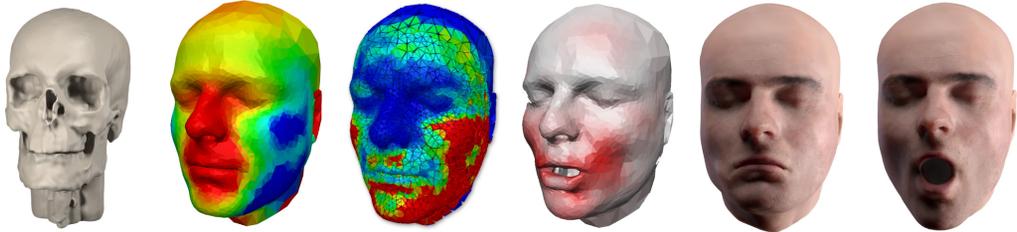


Fig. 1. We build a static anatomical face model from the MRI and use 3D surface scans as training data to learn mechanical parameters that explain deformations of the real face using physics-based simulation.

The human face is an anatomical system exhibiting heterogenous and anisotropic mechanical behavior. This leads to complex deformations even in a neutral facial expression due to external forces such as gravity. We start by building a volumetric model from magnetic resonance images of a neutral facial expression. To obtain data on facial deformations we capture and register 3D scans of the face with different gravity directions and with various facial expressions. Our main contribution consists in solving an inverse physics problem where we learn mechanical properties of the face from our training data (3D scans). Specifically, we learn heterogenous stiffness and prestrain (which introduces anisotropy). The generalization capability of our resulting physics-based model is tested on 3D scans. We demonstrate that our model generates predictions of facial deformations more accurately than recent related physics-based techniques.

CCS Concepts: • **Computing methodologies** → **Physical simulation**.

Additional Key Words and Phrases: physics-based simulation, anatomical modeling, facial animation

ACM Reference Format:

Petr Kadleček and Ladislav Kavan. 2019. Building Accurate Physics-based Face Models from Data. *Proc. ACM Comput. Graph. Interact. Tech.* 2, 2, Article 15 (July 2019), 15 pages. <https://doi.org/10.1145/3340256>

1 INTRODUCTION

Facial appearance plays an important role in many areas of life, including communication and interpersonal relationships. The human face is a fascinating mechanical system able to generate facial expressions by contractions of delicate facial muscles. Forces due to contractions of these

Authors' addresses: Petr Kadleček, University of Utah, Salt Lake City, UT, USA, Charles University, Prague, Czech Republic, petr.kadlecek@gmail.com; Ladislav Kavan, University of Utah, Salt Lake City, UT, USA, ladislav.kavan@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or 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 the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2577-6193/2019/7-ART15 \$15.00

<https://doi.org/10.1145/3340256>

muscles are transmitted to adjacent muscles, the bones and the skin via a network of connective tissues. Despite decades of progress, accurate biomechanical face modeling remains a challenge. Computer graphics typically employs direct deformation models such as blendshapes which capture only the surface (skin) shapes and do not attempt to model the internal biomechanics or anatomy of the face. Anatomical face models [Sifakis et al. 2005] have been recently gaining popularity [Bao et al. 2018; Cong et al. 2016; Ichim et al. 2017; Kozlov et al. 2017]. One of their advantages is the ability to add realistic physics-based effects such as collisions or inertia. Physics-based simulation in computer graphics typically involves tuning of geometric and material properties to achieve the desired visual effect. This process is non-trivial and time consuming even for experienced technical artists. Measurement-based approaches for determining mechanical parameters of biological soft tissues have been developed in biomechanics, typically involving mechanical loading experiments with excised samples of tissues [Lapeer et al. 2011] or using a specifically designed aspiration device [Luboz et al. 2014].

In this paper we propose a new approach to creating realistic physics-based face models, aiming to bridge the gap between graphics and biomechanics. We designed a simple synthetic experiment shown in Figure 2 to determine whether we can learn complex material properties of a cylinder affected by known external forces and observations of its surface deformations. The results showed (Figure 2 c) that even a single observation can be used for a non-trivial material parameter fitting. This sparked the idea of using various observations of facial soft tissue to discover its complex structure. We start by building a static face model in a neutral expression from an MRI scan. Note that MRI is not essential for our method and any template model adapted [Dicko et al. 2013] to the neutral scan can be used with some loss of accuracy depending on the distance of the adapted template model from the subject-specific skull geometry as shown in Section 5. We generate our segmentations and tet-meshes using existing tools [Fedorov et al. 2012] and claim no contributions in this part. Our main contribution consists of automated learning of mechanical properties of our subject’s face from data (3D scans of facial deformations).

An interesting fact attributable to the high elasticity of the face is that changing the direction of the gravity leads to surprisingly large deformations [Ozsoy et al. 2015]. By varying the direction of the gravity (by posing the subject’s head in various orientations) and by observing various facial expressions created by voluntary muscle activations, we can deduce the composition and mechanical properties of facial soft tissues. To be able to explain our captured facial deformations, we found that in addition to realistic muscle activation models, it is important to use heterogeneous elastic properties including prestrain. Prestrain refers to a biomechanical model of rest-pose stresses of biological tissues [Gee et al. 2010], analogous to the tension in the strings of a guitar. Adding prestrain introduces anisotropic trainable parameters which improves the explanatory power of our model while being well motivated biologically. Our method does not require any manual segmentation of individual muscles, nor any prior knowledge of muscle fiber directions of a template model.

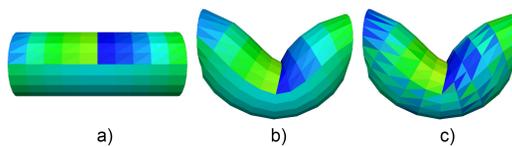


Fig. 2. Synthetic experiment for learning material parameters. a) Tetmesh in a rest pose with color coded varying stiffness not affected by gravity. b) The mesh in quasi-static equilibrium fixed on sides with applied gravity force c) Our prediction of stiffness parameters given the rest pose and surface of the tetmesh in b)

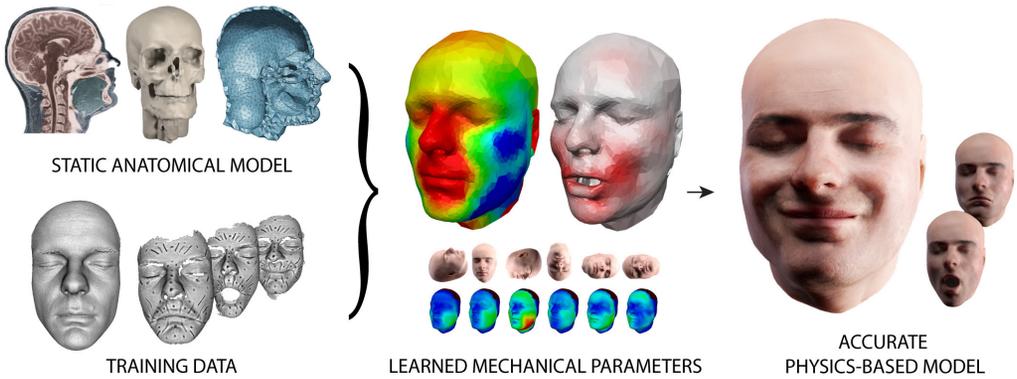


Fig. 3. Given a static face model, we use 3D surface scans as training data in our inverse modeling problem. The result is a physics-based mechanical model with material parameters that explain deformations of the real face.

Contributions. To our knowledge, computation of mechanical parameters of soft tissues from in-vivo surface measurements (3D scans) has not been attempted before. We present a formulation of a physics-based inverse problem that is estimating anatomical model parameters while addressing challenges stemming from fitting living tissue scans. Our problem can be robustly solved using state-of-the-art solvers. We believe that our work makes a step forward towards realistic physics-based modeling capable of explaining mechanical response of soft tissues and formation of facial expressions.

2 RELATED WORK

Facial animation dates back to the pioneering work [Blanz and Vetter 1999; Terzopoulos and Waters 1990]. More recently, high-fidelity capture setups of actor’s faces have become important in the film and game industries, such as Medusa system [Beeler et al. 2010; Beeler and Bradley 2014; Beeler et al. 2011] or other [Alexander et al. 2010; Smith et al. 2017; von der Pahlen et al. 2014]. Method for accurate registration and stabilization have been explored, including kinematics models of a jaw [Beeler and Bradley 2014; Yang et al. 2018; Zoss et al. 2018]. Industrial pipelines typically involve data-driven blendshape-type models which continue to be an important topic of research [Ichim et al. 2015; Lewis et al. 2014; Li et al. 2017; Yoon et al. 2017; Zell et al. 2017]. In general, the chief advantage of blendshape-type approaches is direct control of skin deformations which translates into high visual fidelity. Achieving high quality with biomechanical facial modeling is harder, because facial muscles influence the shape of the skin indirectly and are sensitive to the precise geometric and material modeling of anatomical shapes [Cong et al. 2015, 2016; Sifakis et al. 2005]. Despite these challenges, the quest for high-fidelity anatomical face models continues both in the academia and the industry [Lan et al. 2017].

Physics-based anatomical models have been explored in computer graphics especially in the context of full-body animation [Lee et al. 2009; Teran et al. 2003, [n. d.], 2005]. Recent methods explore new numerical simulation strategies such as Eulerian-on-Lagrangian simulation [Fan et al. 2013, 2014], Projective Dynamics [Kadlecěk et al. 2016; Saito et al. 2015], or a combination of data-driven and physics-based animation [Kim et al. 2017]. While sharing the same biological underpinnings, the shape and function of facial muscles are quite different from large skeletal muscles such as the biceps [Blemker et al. 2005]. In particular, facial muscles are very thin and are

attached via connective tissues to each other and the skin; the primary function of facial muscles is generation of facial expressions [Ekman and Friesen 1977].

Building upon seminal work [Terzopoulos and Waters 1990], in recent years there has been a resurgence of interest in physics-based facial animation in computer graphics. Blendshape animation has been enriched with dynamics effects using mass-spring systems [Ma et al. 2012], “Projective Dynamics”-based surface simulation [Barrielle et al. 2016], or finite element simulation of an outer volumetric layer attached to an inner blendshape model [Kozlov et al. 2017] (a hybrid approach effective also for full-body animation [Kim et al. 2017]). Anatomical modeling is a non-trivial task which can be facilitated using automatic methods [Stavness et al. [n. d.]], which is especially challenging for facial modeling of stylized characters such as Yoda [Cong et al. 2015].

Accurate control of facial expressions via simulated facial muscles is a hard problem. Cong and colleagues introduced “art-directed muscles” that allowed artists to sculpt facial muscle activations [Cong et al. 2016]. This was further improved with careful modeling of facial anatomy [Lan et al. 2017]. A different approach was explored by [Ichim et al. 2016], who proposed “volumetric blendshapes”, allowing all tetrahedrons in a finite element simulation to activate and thus achieve desired shapes while enjoying the benefits of physics-based simulation. A blendshape-driven muscle control approach was also successfully applied for fitting monocular RGB images with a fully differentiable pipeline [Bao et al. 2018]. Recently, [Ichim et al. 2017] fit a generalized muscle model to training data (3D scans of an actor in different facial expressions). While related to our approach, the muscle model of [Ichim et al. 2017] can produce biologically unrealistic activations which compensate for the simplifying assumptions in material modeling. Specifically, the elasticity model used in [Ichim et al. 2017] relies on a homogenous isotropic material with zero prestrain, which is only a crude approximation of the real mechanical behavior of facial soft tissues. In this paper, we propose methods to learn heterogeneous, anisotropic materials with prestrain from data, allowing our model to predict realistic mechanical effects, such as changing directions of gravity.

Methods for learning material parameters from data have a long history in graphics [Pai et al. 2001]. [Bickel et al. 2009] captured non-linear heterogeneous soft tissues by probing a deformable object and measuring its elastic response using computer vision techniques. Even though their results include a face example, the face model is passive (not actuated) and does not model muscles or prestrains. Robotic actuation has been applied to create animatronic facial replicas with the aid of 3D printing [Bickel et al. 2012]. Despite some similarities, the mechanics of man-made robotic systems is quite different from living tissues. 3D printing of deformable objects that assume the desired shapes under gravity motivated development of novel numerical methods [Chen et al. 2014], later extended with data-driven modeling of dynamics [Wang et al. 2015].

Mechanical modeling of biological soft tissues has been extensively studied in biomechanics [Weiss et al. 1996] and remains an active area of research aided with simulation platforms such as the open source ArtiSynth [Lloyd et al. [n. d.]], SOFA [Faure et al. 2012], and FEBio [Maas et al. 2012] or the commercial Abaqus [Hibbitt et al. 2001]. In-vivo material parameter identification of the breast was studied from gravity loading data and plate compression [Han et al. 2011]. Synthetic experiments showed that three gravity loading orientations were required to identify heterogeneous parameters of a silicon gel beam [Gamage et al. 2011]. [Barbarino et al. 2009] modeled elasticity of the face using an advanced Rubin-Bodner material [Rubin and Bodner 2002], but did not consider muscle activations or formation of facial expressions. In a similar vein as [Sifakis et al. 2005], an MRI-based biomechanical face model was constructed by Wu [2013]. In contrast to this paper, [Wu 2013] finds muscle activations for target expressions manually and relies on accurate muscle segmentations from the MRI and material (Mooney-Rivlin) parameters from the literature [Nazari et al. 2010; Tran et al. 2007].

3 DATA CAPTURE AND PROCESSING

Even though modern MRI scanners are very powerful and provide good resolution volumetric scans of the human body without any radiation, they have limitations. In addition to the high cost, long scanning times and limited availability of MRI machines, an important limitation is that the subject must remain motionless inside the MRI scanner for several minutes (depending on the sequence). We scanned a neutral face in a supine position and segmented the bones and skin as shown in Figure 4.

Long scanning time means that MRI scanning of most facial expressions is practically impossible because muscle fatigue would prevent the subject from staying still for minutes. Instead, we captured geometry of deformed facial shapes using a structured light scanner (Artec Spider), producing detailed 3D scans of the skin (Figure 4). Specifically, in our deformed facial shapes we vary gravity directions (by changing the subject's head orientation) and facial expressions (by asking the subject to smile, frown, etc.). The face is quite supple and varying gravity directions results in surprisingly large skin displacements, see Figure 9. Please see the supplemental material for more details about MRI data processing, segmentation, 3D surface scan processing and registration.

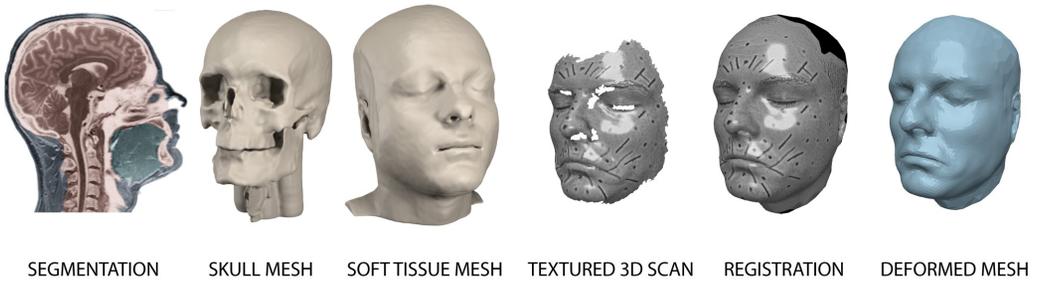


Fig. 4. Segmentation of soft tissues and the bones from our MRI scans and registration of textured 3D scans.

4 METHOD

The main technical challenge we address is an inverse problem where we solve for mechanical parameters of a model that will match acquired 3D scans of real facial shapes under various external loads and muscle activations. Solving it leads to a large optimization problem which can not be solved separately per expression.

4.1 Mechanical model

We start by discussing the key mechanical modeling concepts used in our method. Building upon the finite element method [Sifakis and Barbic 2012], we are optimizing for a rest shape \mathbf{x}_{rest} in a static equilibrium subject to no external forces (including gravity). The rest shape can, however, be subject to internal forces representing inherent residual biological strain which in addition to varying material parameters μ, λ affect the behavior of soft tissue elastic deformation.

Modeling passive tissue. When excised from the body, biological soft tissues retract, i.e., release stored elastic potential energy. This mechanical effect can be modeled with prestrain and it plays an important role in the body, e.g., in stabilizing joints [Ellis et al. 2006]. We use linear finite elements and model prestrain by introducing a material space-dependent function $\mathbf{G} \in \mathbb{R}^{3 \times 3}$ which modulates the deformation gradient. The total internal elastic potential energy of a prestrained deformable

solid can be approximated as:

$$E_{\text{elast}}(\mathbf{x}) = \sum_{i=1}^m v_i \Psi(\mathbf{F}_i(\mathbf{x}, \mathbf{x}_{\text{rest}}) \mathbf{G}_i)$$

where $v_i > 0$ is rest-pose volume of i -th tetrahedron, Ψ is a hyperelastic energy density function, $\mathbf{F}_i \in \mathbb{R}^{3 \times 3}$ is deformation gradient (function of the current deformed state $\mathbf{x} \in \mathbb{R}^{3n}$ and the rest pose $\mathbf{x}_{\text{rest}} \in \mathbb{R}^{3n}$) and $\mathbf{G}_i \in \mathbb{R}^{3 \times 3}$ is the prestrain for i -th tetrahedron (m is the number of tets and n the number of vertices of our tet-mesh). With $\mathbf{G}_i = \mathbf{I}$, E_{elast} becomes equivalent to a standard FEM setup [Sifakis and Barbic 2012]. With general \mathbf{G}_i , there can be non-zero potential energy stored in the configuration \mathbf{x}_{rest} , i.e., $E_{\text{elast}}(\mathbf{x}_{\text{rest}}) > 0$, making the naming convention less intuitive. Observing the object configuration \mathbf{x}_{rest} without prestrain or measuring the prestrain directly is difficult. One would have to make infinitely small dissections corresponding to each material point of our model [Maas et al. 2016]. An observable rest pose $\mathbf{x}_{\text{rest}O}$ affected by prestrain and gravity is a solution to $\nabla E_{\text{elast}}(\mathbf{x}_{\text{rest}O}) + \nabla E_{\text{grav}}(\mathbf{x}_{\text{rest}O}) = 0$, subject to boundary conditions (e.g., fixed skull vertices).

Modeling muscles. Unfortunately, capturing the fiber direction of thin facial muscles is a difficult task even with diffusion tensor imaging [Levin et al. 2011]. The 3D surface of facial expressions are very sensitive to the directions of contractile muscle forces, which has been identified to be a major obstacle in realistic facial modeling [Cong et al. 2016]. In our method, we use general deformation model that can exert arbitrary internal force (i.e., not constrained by an average fiber direction per one tetrahedron), formulated as a transformation of the rest shape matrix: $\mathcal{A}(\mathbf{A}_i) \in \mathbb{R}^{3 \times 3}$ where $\mathbf{A}_i \in \mathbb{R}^6$. This flexibility allows fitting arbitrary surface scans making the method robust, but it can easily lead to overfitting of input data by treating passive deformation as muscle contraction which leads to overly stiff model that does not generalize well to unseen poses [Ichim et al. 2017]. To build as accurate of a model as possible, we need to clearly differentiate between deformation due to internal passive, active and external forces. We achieve this by solving simultaneously for both passive and active deformation and by regularization of muscle activation inspired by the principle of least action in biomechanics [Nubar and Contini 1961]. This is important even for neutral expressions. Even though a subject can be relaxed, some muscles are unconsciously activated (e.g. pushing lower teeth against upper teeth). A large difference in neutral expression can be observed in persons affected by Bell's palsy.

Elastic deformation. Our internal elastic potential uses corotated linear elasticity [Sifakis and Barbic 2012] augmented with prestrain and a muscle activation model:

$$E_{\text{elast}}(\mathbf{x}, \mathbf{A}, \mathcal{H}) = \sum_{i=1}^m v_i \mu_i \|\mathbf{P}_i(\mathbf{x}, \mathbf{x}_{\text{rest}}, \mathbf{G}_i) - \mathcal{A}(\mathbf{A}_i)\|_F^2 + v_i \frac{\lambda_i}{2} \text{tr}^2(\mathbf{P}_i(\mathbf{x}, \mathbf{x}_{\text{rest}}, \mathbf{G}_i) - \mathbf{I}) \quad (1)$$

The vertex positions $\mathbf{x} \in \mathbb{R}^{3n}$ and muscle activations $\mathcal{A}(\mathbf{A}_i)$ are time-varying parameters (where n is the number of vertices, m the number of tets in our mesh and $\mathcal{A}(\mathbf{A}_i) \in \mathbb{R}^{3 \times 3}$ is a symmetric matrix). The notation $\mathcal{H} = (\mathbf{x}_{\text{rest}}, \mathbf{E}, \mathbf{G})$ collects all time-invariant, subject-specific parameters: the rest pose \mathbf{x}_{rest} , Young's modulus \mathbf{E} and prestrain \mathbf{G} . While constant in time, these parameters vary spatially, as indicated by the tetrahedron index i . Lamé coefficients μ, λ are linear functions of Young's modulus \mathbf{E} and Poisson's ratio set to constant $\nu = 0.45$. The prestrained corotated term $\mathbf{P}_i(\mathbf{x}, \mathbf{x}_{\text{rest}}, \mathbf{G}_i) = \mathbf{R}_i \mathbf{F}_i(\mathbf{x}, \mathbf{x}_{\text{rest}}) \mathbf{G}_i$ is obtained by factoring out rotational component of $\mathbf{F}_i \mathbf{G}_i$. The first term is enforcing the shape of the tetrahedron to match the muscle activation, the second term is minimizing relative volumetric change of the deformation. Passive tissues can be modeled

by setting \mathbf{A}_i to the 3×3 identity matrix. In practice, our tets are relatively coarse and may thus contain a mix of passive and active tissues, so \mathbf{A}_i can be seen as their blending.

Gravity. Accuracy of the gravity direction is one of the important factors in our material parameter estimation. Some 3D scanning devices (e.g., handheld scanners) do not provide information about a gravity vector and it is therefore necessary to optimize for it. We initialize a gravity potential $E_{\text{grav}}(\mathbf{x}, \mathbf{g})$ with a best guess of a gravity force vector $\mathbf{g} \in \mathbb{R}^3$ corresponding to each of our surface scan.

Constraints. We use Dirichlet boundary conditions corresponding to the bones (the skull and the jaw), expressed as $\mathbf{c}_{\text{bone}}(\mathbf{x}, \mathbf{J}) = 0$, where $\mathbf{J} \in SE(3)$ is a time-varying rigid transformation of the jaw relative to the skull. The $\mathbf{c}_{\text{bone}}(\mathbf{x}, \mathbf{J})$ function fixes the vertices of the skull in place, while vertices of the jaw are rigidly transformed by \mathbf{J} .

Forward problem. With this model, quasi-static facial deformations can be computed by minimizing $E_{\text{elast}}(\mathbf{x}, \mathbf{A}, \mathcal{H}) + E_{\text{grav}}(\mathbf{x}, \mathbf{g})$ subject to $\mathbf{c}_{\text{bone}}(\mathbf{x}, \mathbf{J}) = 0$. The inverse problem, i.e., computing both \mathcal{H} (expression-independent parameters) as well as muscle activations \mathbf{A} and jaw transformations \mathbf{J} (expression-dependent parameters) which best explain our training 3D scans of facial expressions is explained in the following subsection.

4.2 Inverse problem formulation

The input data to our inverse modeling module are registered 3D scans of facial deformations and tetrahedral template mesh as shown in Figure 3. The data capture and processing is discussed in more detail in the supplemental material.

We denote the vertices of the “target” 3D scans as $\mathbf{t}_k \in \mathbb{R}^{3p}$, where k is the target scan index and p the number of vertices in the triangle mesh representing the skin (the same for all targets due to registration). Because our tet-mesh is volumetric and $\mathbf{x} \in \mathbb{R}^{3n}$ thus includes internal vertices, we define a sparse selector matrix $\mathbf{S} \in \mathbb{R}^{3p \times 3n}$ such that $\mathbf{S}\mathbf{x}$ “selects” the skin vertices corresponding to the targets \mathbf{t}_k . In the inverse optimization problem we will also solve for the rest pose $\mathbf{x}_{\text{rest}} \in \mathbb{R}^{3n}$ discussed in Section 4.1. Note that \mathbf{x}_{rest} is *not* subject to gravity. We formulate our inverse modeling as the following optimization problem:

$$\begin{aligned} \min_{\mathbf{x}_{\text{rest}}, \mathbf{x}_k, \mathbf{A}_k, \mathbf{J}_k, \mathbf{g}_k, \mathcal{H}} \quad & \|\mathbf{S}\mathbf{x}_k - \mathbf{T}_k \mathbf{t}_k\|^2 + E_{\text{reg}}(\mathbf{A}_k, \mathbf{x}_{\text{rest}}, \mathbf{E}, \mathbf{G}, \mathbf{D}) \\ \text{subj. to} \quad & \nabla E_{\text{elast}}(\mathbf{x}_{\text{rest}}, \mathbf{x}_k, \mathbf{A}_k, \mathcal{H}) + \nabla E_{\text{grav}}(\mathbf{x}_k, \mathbf{g}_k) = 0 \\ & \mathbf{c}_{\text{bone}}(\mathbf{x}_{\text{rest}}, \mathbf{I}) = 0, \mathbf{c}_{\text{bone}}(\mathbf{x}_k, \mathbf{J}_k) = 0, \mathbf{c}_{\text{col}}(\mathbf{x}_k) \geq 0 \end{aligned} \quad (2)$$

This problem simultaneously solves for multiple tet-mesh deformations $\mathbf{x}_k \in \mathbb{R}^{3n}$, where k indexes target 3D scans. In the objective function $\|\mathbf{S}\mathbf{x}_k - \mathbf{T}_k \mathbf{t}_k\|^2$, the $\mathbf{T}_k \in SE(3)$ is a rigid transformation applied to the target triangle mesh \mathbf{t}_k in order to compensate for global rigid motions of the skull (a process known as “rigid stabilization” [Beeler and Bradley 2014]). E_{reg} is a regularization term discussed in the Section 4.3. The constraint $\mathbf{c}_{\text{bone}}(\mathbf{x}_{\text{rest}}, \mathbf{I}) = 0$ means that the jaw in the rest pose is closed (its rigid transformation is identity). In contrast, in non-neutral expressions the jaw vertices are rigidly transformed by $\mathbf{J}_k \in SE(3)$. The constraints $\mathbf{c}_{\text{col}}(\mathbf{x}_k) \geq 0$ are used to resolve collisions: if we find a vertex in \mathbf{x}_k penetrating a tetrahedron, we find projection onto the closest surface triangle and add a half-space constraint to $\mathbf{c}_{\text{col}}(\mathbf{x}_k) \geq 0$ which pushes the vertex out of the collision. Finally, the constraint $\nabla E_{\text{elast}}(\mathbf{x}_k, \mathbf{A}_k, \mathcal{H}) + \nabla E_{\text{grav}}(\mathbf{x}_k, \mathbf{g}_k) = 0$ requests the deformed pose \mathbf{x}_k to be in quasi-static equilibrium subject to gravity and boundary conditions. Our target 3D scans intentionally contain different head orientations, i.e., the gravity directions for each target can be different and are represented by trainable parameter $\mathbf{g}_k \in \mathbb{R}^3$.

4.3 Solving the inverse problem

Solving Eq. 2 numerically is challenging. Our optimization problem has many variables because we are solving for deformed vertex positions $\mathbf{x}_1 \in \mathbb{R}^{3n}, \mathbf{x}_2 \in \mathbb{R}^{3n}, \dots$ for all of our targets. Furthermore, the deformed vertex positions are coupled through the rest pose \mathbf{x}_{rest} and shared mechanical parameters \mathbf{E}, \mathbf{G} , which makes the problem *not* separable per each target. To further complicate matters, the problem is ill-posed, because our data \mathbf{t}_k correspond only to skin (surface) measurements, but we are solving for unknown volumetric properties.

Our optimization strategy is to use the idea of homotopy. In our case, we use a progression from an initial state \mathbf{x}_{init} which is in a quasi-static equilibrium, but does not match target scans, to a state \mathbf{x} fitting all target scans as close as possible in a sense of Eq. 2. Even though asymptotic numerical method [Chen et al. 2014] can provide better convergence than standard Newton-type optimization, the implementation for a general deformation problem is challenging. We found that a linear increase of the target term weight (from 0 to 1), iterative update of regularization terms E_{reg} and block coordinate descent optimization proved to be a good strategy for stable convergence and allowed us to prepare an automatic procedure for the complete inverse optimization.

In the first phase, we fit a set of six scans with neutral expression but with varying directions of gravity (see the top row in Figure 9). We compute an initial guess of the gravity-free rest pose $\mathbf{x}_{\text{rest}}^{\text{init}}$ by averaging all of our 3D scans of the neutral expression under different directions of gravity. This is a good initial guess because for each gravity direction \mathbf{d} , our dataset also contains the opposite direction $-\mathbf{d}$, e.g., for a supine head orientation (facing upwards), we also have a prone orientation (facing downwards). We also use the $\mathbf{x}_{\text{rest}}^{\text{init}}$ as a regularization term for \mathbf{x}_{rest} .

Next, we set $\mathbf{G} = \mathbf{I}$ and optimize for the heterogeneous material properties \mathbf{E} . We assume constant density of soft tissues. In the case of quasi-statics, variation of densities and stiffness results in ambiguous fitting. If an MRI scanner is available, some sequences may provide a better initialization of densities. Optimizing for spatially varying material parameters has been recognized as a challenge in previous work [Kim et al. 2017; Wang et al. 2015], which proposed to simplify the problem by sampling material properties only in a sparse set of points (such as 130 points). This reduces the number of variables, but also reduces expressivity and relies on well-chosen locations of these sample points, which is non-trivial in the face. Instead, we solve for a separate \mathbf{E} for each tet by progressive refinement. The process starts by optimizing a single \mathbf{E} for all tets (homogeneous material), which theoretically corresponds to an infinitely strong Laplacian regularization term. Specifically, we use the standard umbrella operator penalizing $\|\mathbf{E}_a - \mathbf{E}_b\|^2$ summed over all pairs (a, b) of adjacent tets, as in [Bickel et al. 2009]. Next, we use the homogeneous \mathbf{E} as an initial guess and progressively reduce the weight of the Laplacian regularization term. This way, we find heterogeneous stiffness parameters E but so far only with isotropic material. Next, we allow the optimizer to optimize for prestrains \mathbf{G} (along with \mathbf{E} and \mathbf{x}_{rest}), introducing anisotropy into our model. We applied additional regularization on \mathbf{E} and \mathbf{G} with a closeness term to an initial guess which we update in each iteration k : $\|\mathbf{G}_i^k - \mathbf{G}_i^{k-1}\|_F^2$ (regularization plasticity). Finally, we enable optimization of the muscle parameters \mathbf{A} . We use an order of magnitude stronger weight for the muscle activation regularization $\|\mathbf{A}_i - \mathbf{I}\|_F^2$ to enforce principle of least action (Section 4.1).

Additional free variables of our optimization process include kinematic parameters of the jaw (\mathbf{J}_k). We allow arbitrary rigid transformation, parameterized with three DoFs for translation and three Euler angles for rotation. We did not encounter any singularities with the Euler angles because the set of physiological rotations of the jaw is limited. We used the same parametrization for global head transformation (rigid stabilization) \mathbf{T}_k , taking advantage of our “pre-stabilization” to ensure that changes of global orientation at solver time will be small. Similarly, we parameterize gravity directions with two angles (polar and azimuthal angle, the magnitude is fixed to $\|\mathbf{g}_k\| = 9.81\text{m/s}^2$).

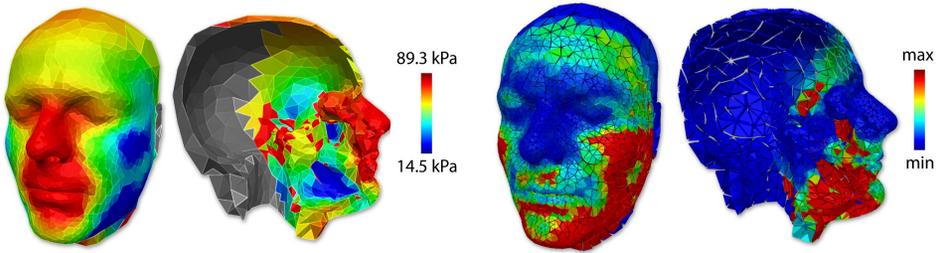


Fig. 5. Visualization of heterogeneous stiffness of our model and prestrain at each tetrahedron.

We solve our optimization problems using an open source interior point solver IPOPT [Wächter and Biegler 2006] with PARDISO [Schenk and Gärtner 2004] as the underlying linear solver. Our mechanical model is twice differentiable. In most cases, we provide the solver with exact gradient and Hessian. In the inverse problem, the quasi-static constraint Hessian is approximated. An exact Hessian contains third derivative of the mechanical model resulting from the chain rule which we ignore [Bickel et al. 2012]. Our problem converges in total of 25 iterations (5 iterations per each phase) and takes approximately 20 minutes for the highest resolution tetmesh (21k tetrahedra) and 13 targets scans on a desktop CPU.

5 RESULTS

Stiffness. We visualize our resulting stiffness parameter in Figure 5. The computed Young’s modulus values are in range of $E \in [14.5, 89.3]$ with an average of $E = 68.2kPa$. These values are aligned with in-vivo stiffness measurements [Luboz et al. 2014; Shinohara et al. [n. d.]]. However, material properties reported in the literature must be interpreted judiciously, e.g., Lapeer et al. [2011] surveyed literature on skin stiffness measurements and found large differences between the reported values (attributed to differences between in-vivo/in-vitro measurements, nonlinearity of the skin and differences in elastic models). In Figure 5 we can see that the stiffness is lowest in areas such as the cheeks which contain the thickest layers of adipose tissues. Note that stiffness in the lips region seem to be too large. We observed that in case of large noise in the data (e.g., lip motion fitted by muscle deformation apparent in Figure 9), our isotropic stiffness fitting failed to capture this deformation as passive. This was not the case for experiments using only selected scans with minimal lip motion. Fortunately, the anisotropic prestrain model is more flexible and is able to generalize better as shown in Fig. 9.

Prestrain. Prestrain models the fact that soft tissues retract when excised from the body. This means that each tet in our mesh has a relaxed configuration with zero prestrain. Visualization of such configuration becomes more complicated because tet vertices corresponding to neighboring tets might not be aligned. Nevertheless, we can visualize prestrains by disconnecting our tet-mesh into isolated tets, scaling them down by a factor of 0.8, and transforming them by G , see Figure 5. The color corresponds to prestrain magnitude computed as $\|G - I\|_F^2 \in [0, 3.7]$.

Muscle activations. Recall that our A matrices specify the deformation of each tet due to muscle contractions. Similarly to prestrain, we visualize our resulting muscle contractions as a magnitude $\alpha = \|A - I\|_F^2$ in Figure 6.

Evaluation. Validation of our method is challenging due to residual muscle activation which is present in all facial expressions of healthy human beings. While measuring EMG signal is technically possible, each sensor measures electrical activity integrated over multiple muscles at the surface and the electrode placement affects the deformation. Ideal validation data would contain surface scans of a completely paralyzed face under different external loads.

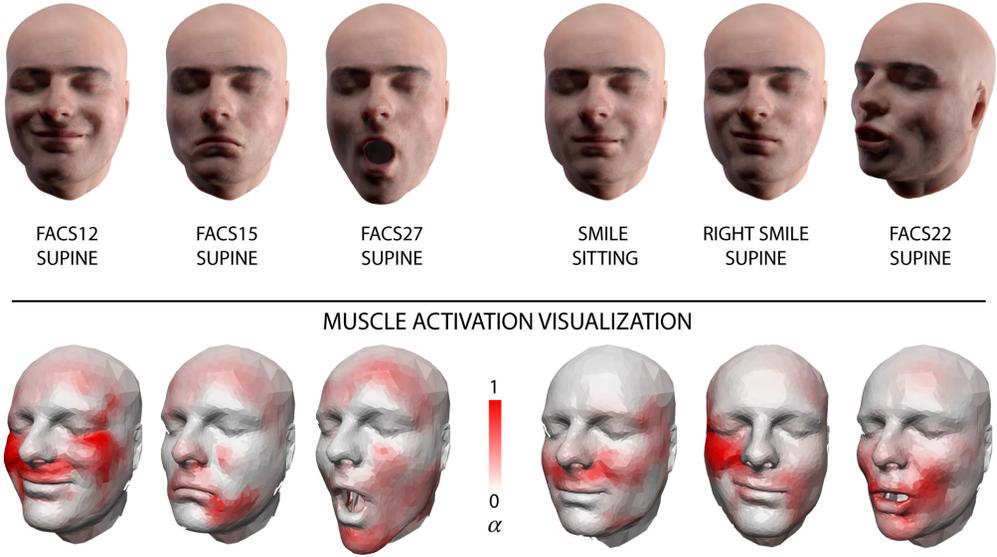


Fig. 6. Explaining various facial expressions with our model using muscle deformations A. Second row: visualization of normalized muscle activation magnitude α used to generate the corresponding facial expressions with our method.

In Figure 9 we try to analyze how well our model predicts 3D scans under varying gravity directions by ignoring the muscle activation. Although some scans clearly show deformation due to motion and muscle activation (e.g., lips), we still find the comparison interesting. Each validation scan was removed from training data. In the first row, we calculate the distance between the original 3D scans against optimized homogeneous material (single E for all tets). The largest difference of 13mm is observed with prone head orientation (facing downwards), where the soft tissues sag the most under gravity. Methods using elastic model with homogeneous material (e.g., [Ichim et al. 2017; Mollemans et al. 2007]) either can not fit similar scans properly or need to compensate by artificial muscle deformation. In the second row, we show the result of our method *without* fitting prestrain parameters G. In the third row is the final result of our method with prestrain enabled, achieving best generalization to all poses and prediction closest to all the 3D scans.

Input data ambiguity. We test our method with different subsets of input scans to show the effect of ambiguity. Figure 7 shows optimized stiffness parameters for subsets **a-e**. The subset **a** uses only one scan. Since we're also optimizing for the rest-pose, the result mostly depends on regularization weights. Subsets **b** and **c** use different subset of size two - showing different results depending on how close input scans are with respect to gravity direction. Subset **d** with three scans shows a result similar to a set using all six gravity directions.

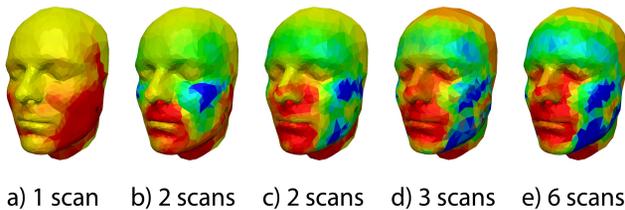


Fig. 7. Stiffness parameters computed from limited scan data.

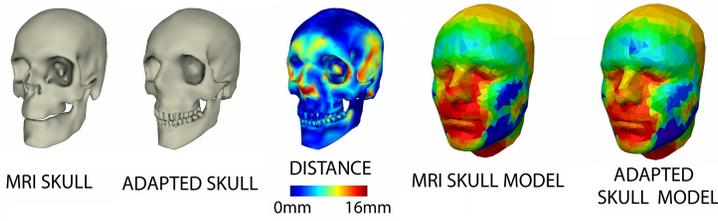


Fig. 8. The first row shows our original skull segmented from MRI scan, the adapted skull from a template model and color-coded distance. The second row shows stiffness computed using the original model and the model with the adapted skull.

Template skull adaptation. We tested our method on a less accurate skull and mandible geometry adapted from a template model to assess the importance of precise segmentation of bones from MRI data. We started by smoothing high-frequency details on both template and target model. Next, we registered both models with manually picked keypoints and prepared a tetrahedral mesh that corresponds to the original surface mesh and the adapted skull. The result of our inverse problem in Figure 8 shows that stiffness parameters near the surface are similar and only small changes near areas of thin soft structures are present. Please note that the template skull geometry after the adaptation is not very different from the segmented skull (maximum distance is 16mm). Adapting skull to different subjects with varying fat distribution is a challenging task [Achenbach et al. 2018].

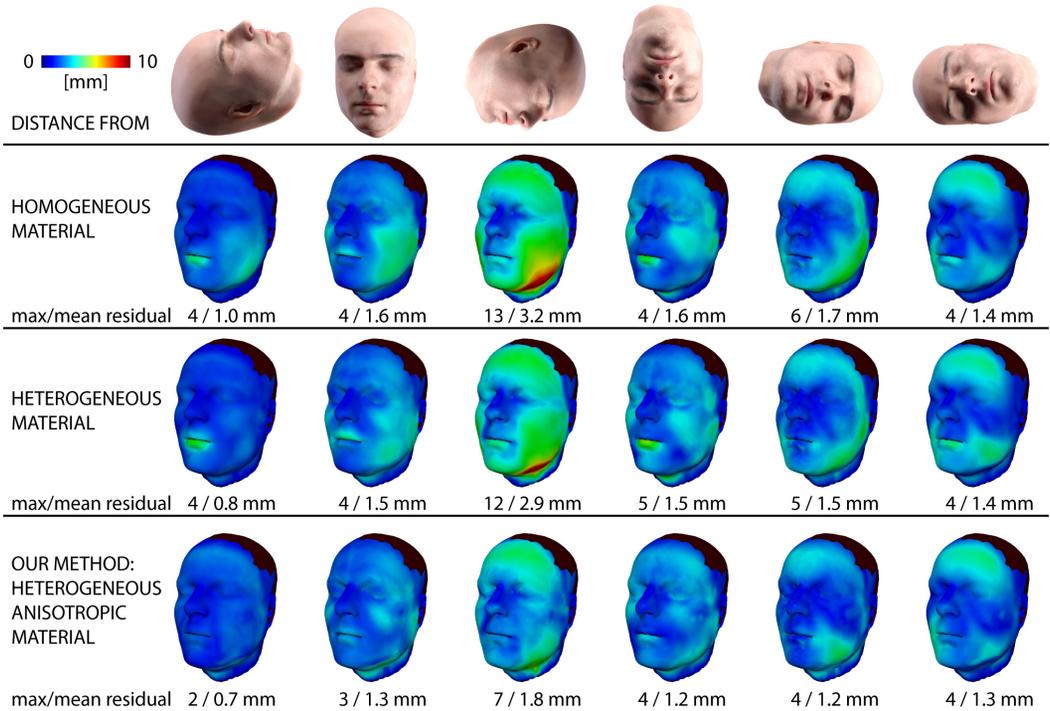


Fig. 9. Comparison of elastic models – distances from 3D scans in various gravity directions (individually excluded from training data) shown in the first row (reference measurements). Second row: simulation with homogeneous material; Third row: heterogenous material without prestrain; last row: heterogeneous material with anisotropic prestrain (our method).

6 LIMITATIONS AND FUTURE WORK

Even though we believe our project is a significant step towards realistic facial modeling, there are several limitations leading to opportunities for future work.

Input data. We tested our method with only one human subject. Our 3D scanning process also needs to be improved, because our structured-light scanner (Artec Spider) is not ideal for capturing facial expressions. One scan can take up to several minutes which makes scanning challenging due to muscle fatigue.

Registration. Incorrect registration (e.g., around lips) and rigid stabilization is causing some errors in the fitting. Ideally, our algorithm should run the inverse problem together with corrections for inaccurate correspondences.

Mechanical model. Our current mechanical model also leaves room for improvement. We note that many different constitutive equations could be used instead of corotated elasticity. Our choice was motivated mainly by tractability of the resulting inverse problem. The additional prestrain can be understood as a simple way to improve an arbitrary elastic model. Whether the anisotropic behavior would generalize well under larger loads is subject to further testing.

Muscle control. We briefly experimented with a model without kinematic jaw constraints where the motion of the jaw was driven purely by muscle activation. The model was able to fit all of our scans with jaw open and we are planning to experiment with purely muscle-driven control without kinematic constraints in the future.

Even though an intuitive blendshape-driven control of muscle activations [Ichim et al. 2017] can be directly applied to our method, our muscle model is overparameterized and might not be suitable for accurate prediction of rest pose changes, e.g., during craniofacial surgical evaluation. Anatomically plausible subject-specific muscle model is subject to future work.

Surface scan ambiguity. The current parameter learning is limited by surface scan observations. Simple synthetic experiments can show that multiple material configurations can lead to almost identical surface deformation under various conditions. Different ways of applying and measuring external loads should be therefore studied. Note that aspiration device testing might not reveal deep non-linearity without damaging the skin [Luboz et al. 2014].

Dynamics. Our current physics-based modeling methodology assumes quasi-static deformations corresponding to slowly moving faces. Even though we can of course add ad-hoc dynamics effects, the proper approach would involve learning viscoelastic properties from data – dynamic deformations of real faces, which would require a dynamic (4D) capture studio, as was recently explored in the case of human bodies [Kim et al. 2017]. In this paper we focused on volumetric modeling of facial soft tissues; production of realistic facial animations would require adding accurate models of hair, teeth, tongue, eyes and eyelids.

7 CONCLUSION

Passive soft tissue deformation may initially seem less important than e.g. muscle activations. However, one may be convinced otherwise when taking a selfie when lying on the back or even upside-down (we encourage the reader to try!). Apart from differences in blood flow, the change in shape of the face can be quite surprising.

The core idea of our method is to overcome the challenging problem of capturing mechanical properties of biological tissues in-vivo from limited data. We believe our approach is an interesting complement to general mainstream machine learning models such as neural networks which can achieve impressive results, but need very large training data sets to do so. Our method has been designed to minimize the amount of training data by taking advantage of domain-specific priors, such as the principles of soft tissue elasticity. These priors contain trainable parameters such as soft

tissue material parameters, which we estimated from training data. We hope that this approach will contribute to the long-term goal of creating realistic models of the human face.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant Numbers IIS-1617172, IIS-1622360 and IIS-1764071. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The work was also supported by grants SVV-2017-260452, GAUK 1524217 and we gratefully acknowledge the support of Activision, Adobe, and hardware donation from NVIDIA Corporation. We thank the Willard Marriott Library at the University of Utah and Tereza Kadleckova for a help with 3D scanning. Finally, we thank the Utah Center for Advanced Imaging Research for their support.

REFERENCES

- Jascha Achenbach, Robert Brylka, Thomas Gietzen, Katja zum Hebel, Elmar Schömer, Ralf Schulze, Mario Botsch, and Ulrich Schwanecke. 2018. A multilinear model for bidirectional craniofacial reconstruction. In *Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine*. Eurographics Association, 67–76.
- Oleg Alexander, Mike Rogers, William Lambeth, Jen-Yuan Chiang, Wan-Chun Ma, Chuan-Chang Wang, and Paul Debevec. 2010. The digital Emily project: Achieving a photorealistic digital actor. *Computer Graphics and Applications, IEEE* 30, 4 (2010), 20–31.
- Michael Bao, Matthew Cong, Stéphane Grabli, and Ronald Fedkiw. 2018. High-Quality Face Capture Using Anatomical Muscles. *CoRR* abs/1812.02836 (2018). arXiv:1812.02836 <http://arxiv.org/abs/1812.02836>
- GG Barbarino, M Jabareen, J Trzewik, A Nkengne, G Stamatas, and E Mazza. 2009. Development and validation of a three-dimensional finite element model of the face. *Journal of biomechanical engineering* 131, 4 (2009), 041006.
- Vincent Barrielle, Nicolas Stoiber, and Cedric Cagniard. 2016. Blendforces, a Dynamic Framework for Facial Animation. *Comp. Graph. Forum* (2016).
- Thabo Beeler, Bernd Bickel, Paul Beardsley, Bob Sumner, and Markus Gross. 2010. High-quality single-shot capture of facial geometry. In *ACM Transactions on Graphics (ToG)*, Vol. 29. ACM, 40.
- Thabo Beeler and Derek Bradley. 2014. Rigid stabilization of facial expressions. *ACM Trans. Graph.* 33, 4 (2014), 44.
- Thabo Beeler, Fabian Hahn, Derek Bradley, Bernd Bickel, Paul Beardsley, Craig Gotsman, Robert W Sumner, and Markus Gross. 2011. High-quality passive facial performance capture using anchor frames. In *ACM Transactions on Graphics (TOG)*, Vol. 30. ACM, 75.
- Bernd Bickel, Moritz Bächer, Miguel A Otaduy, Wojciech Matusik, Hanspeter Pfister, and Markus Gross. 2009. Capture and modeling of non-linear heterogeneous soft tissue. *ACM Trans. Graph.* 28, 3 (2009).
- Bernd Bickel, Peter Kaufmann, Mélina Skouras, Bernhard Thomaszewski, Derek Bradley, Thabo the, Phil Jackson, Steve Marschner, Wojciech Matusik, and Markus Gross. 2012. Physical face cloning. *ACM Trans. Graph.* 31, 4 (2012), 118.
- Volker Blanz and Thomas Vetter. 1999. A morphable model for the synthesis of 3D faces. In *Proc. of the 26th annual conf. on Comp. graph. and interactive techniques*. 187–194.
- Silvia S Blemker, Peter M Pinsky, and Scott L Delp. 2005. A 3D model of muscle reveals the causes of nonuniform strains in the biceps brachii. *Journal of biomechanics* 38, 4 (2005), 657–665.
- Xiang Chen, Changxi Zheng, Weiwei Xu, and Kun Zhou. 2014. An asymptotic numerical method for inverse elastic shape design. *ACM Trans. Graph.* 33, 4 (2014), 95.
- Matthew Cong, Michael Bao, Kiran S Bhat, Ronald Fedkiw, et al. 2015. Fully automatic generation of anatomical face simulation models. In *Proc. of the EG/SIGGRAPH Symposium on Comp. Anim.* ACM, 175–183.
- Matthew Cong, Kiran S Bhat, and Ronald Fedkiw. 2016. Art-directed muscle simulation for high-end facial animation. In *Proc. of the EG/SIGGRAPH Symposium on Comp. Anim.* 119–127.
- Ali-Hamadi Dicko, Tiantian Liu, Benjamin Gilles, Ladislav Kavan, François Faure, Olivier Palombi, and Marie-Paule Cani. 2013. Anatomy transfer. *ACM Trans. Graph.* 32, 6 (2013), 188.
- Paul Ekman and Wallace V Friesen. 1977. Facial action coding system. (1977).
- Benjamin J Ellis, Trevor J Lujan, Michelle S Dalton, and Jeffrey A Weiss. 2006. Medial collateral ligament insertion site and contact forces in the ACL-deficient knee. *Journal of Orthopaedic Research* 24, 4 (2006), 800–810.
- Ye Fan, Joshua Litven, David IW Levin, and Dinesh K Pai. 2013. Eulerian-on-lagrangian simulation. *ACM Transactions on Graphics (TOG)* 32, 3 (2013), 22.

- Ye Fan, Joshua Litven, and Dinesh K Pai. 2014. Active volumetric musculoskeletal systems. *ACM Trans. Graph.* 33, 4 (2014), 152.
- François Faure, Christian Duriez, Hervé Delingette, Jérémie Allard, Benjamin Gilles, Stéphanie Marchesseau, Hugo Talbot, Hadrien Courtecuisse, Guillaume Bousquet, Igor Peterlik, et al. 2012. Sofa: A multi-model framework for interactive physical simulation. In *Soft Tissue Biomechanical Modeling for Computer Assisted Surgery*. Springer, 283–321.
- Andriy Fedorov, Reinhard Beichel, Jayashree Kalpathy-Cramer, Julien Finet, Jean-Christophe Fillion-Robin, Sonia Pujol, Christian Bauer, Dominique Jennings, Fiona Fennessy, Milan Sonka, et al. 2012. 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic resonance imaging* 30, 9 (2012), 1323–1341.
- Babarenda Gamage, P Thiranjai, Vijayaraghavan Rajagopal, Matthias Ehrigott, Martyn P Nash, and Poul MF Nielsen. 2011. Identification of mechanical properties of heterogeneous soft bodies using gravity loading. *International journal for numerical methods in biomedical engineering* 27, 3 (2011), 391–407.
- Michael W Gee, Ch Förster, and WA Wall. 2010. A computational strategy for prestressing patient-specific biomechanical problems under finite deformation. *International Journal for Numerical Methods in Biomedical Engineering* 26, 1 (2010), 52–72.
- Lianghao Han, John H Hipwell, Christine Tanner, Zeike Taylor, Thomy Mertzaniidou, Jorge Cardoso, Sebastien Ourselin, and David J Hawkes. 2011. Development of patient-specific biomechanical models for predicting large breast deformation. *Physics in Medicine and Biology* 57, 2 (2011), 455.
- Hibbitt, Karlsson, and Sorensen. 2001. *ABAQUS/standard User's Manual*. Vol. 1. Hibbitt, Karlsson & Sorensen.
- Alexandru Ichim, Ladislav Kavan, Merlin Nimier-David, and Mark Pauly. 2016. Building and Animating User-Specific Volumetric Face Rigs. In *Proc. of the EG/SIGGRAPH Symposium on Comp. Anim.*
- Alexandru Eugen Ichim, the Bouaziz, and Mark Pauly. 2015. Dynamic 3D Avatar Creation from Hand-held Video Input. *ACM Trans. Graph.* (2015).
- Alexandru-Eugen Ichim, Petr Kadlec, Ladislav Kavan, and Mark Pauly. 2017. Phace: Physics-based Face Modeling and Animation. *ACM Trans. Graph.* 36, 4 (2017).
- Petr Kadlec, Alexandru-Eugen Ichim, Tiantian Liu, Jaroslav Krivanek, and Ladislav Kavan. 2016. Reconstructing Personalized Anatomical Models for Physics-based Body Animation. *ACM Trans. Graph.* 35, 6 (2016).
- Meekyoung Kim, Gerard Pons-Moll, Sergi Pujades, Seungbae Bang, Jinwook Kim, Michael J Black, and Sung-Hee Lee. 2017. Data-driven physics for human soft tissue animation. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 54.
- Yera Kozlov, Derek Bradley, Moritz Bäcker, Bernhard Thomaszewski, Thabo Beeler, and Markus Gross. 2017. Enriching Facial Blendshape Rigs with Physical Simulation. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 75–84.
- Lana Lan, Matthew Cong, and Ronald Fedkiw. 2017. Lessons from the evolution of an anatomical facial muscle model. In *Proceedings of the ACM SIGGRAPH Digital Production Symposium*. ACM, 11.
- Rudy J Lapeer, Paul D Gasson, and Vasudev Karri. 2011. A hyperelastic finite-element model of human skin for interactive real-time surgical simulation. *IEEE Transactions on Biomedical Engineering* 58, 4 (2011), 1013–1022.
- Sung-Hee Lee, Eftychios Sifakis, and Demetri Terzopoulos. 2009. Comprehensive biomechanical modeling and simulation of the upper body. *ACM Trans. Graph.* 28, 4 (2009), 99.
- David IW Levin, Benjamin Gilles, Burkhard Mädler, and Dinesh K Pai. 2011. Extracting skeletal muscle fiber fields from noisy diffusion tensor data. *Medical Image Analysis* 15, 3 (2011), 340–353.
- John P Lewis, Ken Anjo, Taehyun Rhee, Mengjie Zhang, Frederic H Pighin, and Zhigang Deng. 2014. Practice and Theory of Blendshape Facial Models. In *Eurographics (State of the Art Reports)*. 199–218.
- Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics* 36, 6 (Nov. 2017), 194:1–194:17. Two first authors contributed equally.
- John E Lloyd, Ian Stavness, and Sidney Fels. [n. d.]. ArtiSynth: A fast interactive biomechanical modeling toolkit combining multibody and finite element simulation. In *Soft tissue biomechanical modeling for computer assisted surgery*.
- Vincent Luboz, Emmanuel Promayon, and Yohan Payan. 2014. Linear elastic properties of the facial soft tissues using an aspiration device: towards patient specific characterization. *Annals of biomedical engineering* 42, 11 (2014), 2369–2378.
- Wan-Chun Ma, Yi-Hua Wang, Graham Fyffe, Bing-Yu Chen, and Paul Debevec. 2012. A blendshape model that incorporates physical interaction. *Computer Animation and Virtual Worlds* 23, 3-4 (2012).
- Steve A Maas, Benjamin J Ellis, Gerard A Ateshian, and Jeffrey A Weiss. 2012. FEBio: finite elements for biomechanics. *Journal of biomechanical engineering* 134, 1 (2012), 011005.
- Steve A Maas, Ahmet Erdemir, Jason P Halloran, and Jeffrey A Weiss. 2016. A general framework for application of prestrain to computational models of biological materials. *Journal of the mechanical behavior of biomedical materials* (2016).
- Wouter Mollemans, Filip Schutyser, Nasser Nadjmi, Frederik Maes, and Paul Suetens. 2007. Predicting soft tissue deformations for a maxillofacial surgery planning system: from computational strategies to a complete clinical validation. *Medical image analysis* 11, 3 (2007), 282–301.
- Mohammad Ali Nazari, Pascal Perrier, Matthieu Chabanas, and Yohan Payan. 2010. Simulation of dynamic orofacial movements using a constitutive law varying with muscle activation. *Computer Methods in Biomechanics and Biomedical*

- Engineering* 13, 4 (2010), 469–482.
- Yves Nubar and Renato Contini. 1961. A minimal principle in biomechanics. *The bulletin of mathematical biophysics* 23, 4 (1961), 377–391.
- U Ozsoy, R Sekerci, and E Ogut. 2015. Effect of sitting, standing, and supine body positions on facial soft tissue: Detailed 3D analysis. *International journal of oral and maxillofacial surgery* 44, 10 (2015), 1309–1316.
- Dinesh K Pai, Kees van den Doel, Doug L James, Jochen Lang, John E Lloyd, Joshua L Richmond, and Som H Yau. 2001. Scanning physical interaction behavior of 3D objects. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 87–96.
- MB Rubin and SR Bodner. 2002. A three-dimensional nonlinear model for dissipative response of soft tissue. *International Journal of Solids and Structures* 39, 19 (2002), 5081–5099.
- Shunsuke Saito, Zi-Ye Zhou, and Ladislav Kavan. 2015. Computational Bodybuilding: Anatomically-based Modeling of Human Bodies. *ACM Trans. Graph.* 34, 4 (2015).
- Olaf Schenk and Klaus Gärtner. 2004. Solving unsymmetric sparse systems of linear equations with PARDISO. *Future Generation Computer Systems* 20, 3 (2004), 475–487.
- Minoru Shinohara, Karim Sabra, Jean-Luc Gennisson, Mathias Fink, and Mickaél Tanter. [n. d.]. Real-time visualization of muscle stiffness distribution with ultrasound shear wave imaging during muscle contraction. *Muscle & nerve* ([n. d.]).
- Eftychios Sifakis and Jernej Barbic. 2012. FEM simulation of 3D deformable solids: a practitioner’s guide to theory, discretization and model reduction. In *ACM SIGGRAPH 2012 Courses*. 20.
- Eftychios Sifakis, Igor Neverov, and Ronald Fedkiw. 2005. Automatic determination of facial muscle activations from sparse motion capture marker data. In *ACM Trans. Graph.*, Vol. 24. 417–425.
- Alex Smith, Sven Pohle, Wan-Chun Ma, Chongyang Ma, Xian-Chun Wu, Yanbing Chen, Etienne Danvoye, Jorge Jimenez, Sanjit Patel, Mike Sanders, and Cyrus A. Wilson. 2017. Emotion Challenge: Building a New Photoreal Facial Performance Pipeline for Games. In *Proceedings of the ACM SIGGRAPH Digital Production Symposium (DigiPro ’17)*. ACM, New York, NY, USA, Article 8, 2 pages. <https://doi.org/10.1145/3105692.3105695>
- Ian Stavness, Mohammad Ali Nazari, Cormac Flynn, Pascal Perrier, Yohan Payan, John E Lloyd, and Sidney Fels. [n. d.]. Coupled biomech. modeling of the face, jaw, skull, tongue, and hyoid bone. In *3D Multiscale Physiological Human*.
- Joseph Teran, Sylvia Blemker, V Hing, and Ronald Fedkiw. 2003. Finite volume methods for the simulation of skeletal muscle. In *Proc. of the EG/SIGGRAPH Symposium on Comp. Anim.* Eurographics Association, 68–74.
- Joseph Teran, Eftychios Sifakis, Silvia S Blemker, Victor Ng-Thow-Hing, Cynthia Lau, and Ronald Fedkiw. [n. d.]. Creating and simulating skeletal muscle from the visible human data set. *Vis. and Computer Graphics, IEEE Trans. on* ([n. d.]).
- Joseph Teran, Eftychios Sifakis, Geoffrey Irving, and Ronald Fedkiw. 2005. Robust quasistatic finite elements and flesh simulation. In *Proc. of the EG/SIGGRAPH Symposium on Comp. Anim.* ACM.
- Demetri Terzopoulos and Keith Waters. 1990. Physically-based facial modelling, analysis, and animation. *Computer Animation and Virtual Worlds* 1, 2 (1990), 73–80.
- HV Tran, F Charleux, M Rachik, Alain Ehrlicher, and MC Ho Ba Tho. 2007. In vivo characterization of the mech. properties of human skin derived from MRI and indentation techniques. *Comp. methods in biomech. and biomed. eng.* (2007).
- Javier von der Pahlen, Jorge Jimenez, Etienne Danvoye, Paul Debevec, Graham Fyffe, and Oleg Alexander. 2014. Digital Ira and Beyond: Creating Real-time Photoreal Digital Actors. In *ACM SIGGRAPH 2014 Courses (SIGGRAPH ’14)*. Article 1, 384 pages. <https://doi.org/10.1145/2614028.2615407>
- Andreas Wächter and Lorenz T Biegler. 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming* 106, 1 (2006).
- Bin Wang, Longhua Wu, KangKang Yin, Uri Ascher, Libin Liu, and Hui Huang. 2015. Deformation capture and modeling of soft objects. *ACM Trans. Graph.* 34, 4 (2015), 94.
- Jeffrey A Weiss, Bradley N Maker, and Sanjay Govindjee. 1996. Finite element implementation of incompressible, transversely isotropic hyperelasticity. *Computer methods in applied mechanics and engineering* 135, 1 (1996), 107–128.
- Tim Wu. 2013. *A computational framework for modelling the biomechanics of human facial expressions*. Ph.D. Dissertation. ResearchSpace@ Auckland.
- Wenwu Yang, Nathan Marshak, Daniel Sýkora, Srikumar Ramalingam, and Ladislav Kavan. 2018. Building Anatomically Realistic Jaw Kinematics Model from Data. *CoRR abs/1805.05903* (2018). arXiv:1805.05903 <http://arxiv.org/abs/1805.05903>
- Seung-Hyun Yoon, John Lewis, and Taehyun Rhee. 2017. Blending Face Details: Synthesizing a Face Using Multiscale Face Models. *IEEE computer graphics and applications* 37, 6 (2017), 65–75.
- Eduard Zell, JP Lewis, Junyong Noh, Mario Botsch, et al. 2017. Facial retargeting with automatic range of motion alignment. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 154.
- Gaspard Zoss, Derek Bradley, Pascal Bérard, and Thabo Beeler. 2018. An empirical rig for jaw animation. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 59.