

Capture and Modeling of Non-Linear Heterogeneous Soft Tissue

Bernd Bickel¹ Moritz Bächer² Miguel A. Otaduy³ Wojciech Matusik⁴ Hanspeter Pfister² Markus Gross¹
¹ETH Zurich ²Harvard University ³URJC Madrid ⁴Adobe Systems, Inc.




Figure 1: From left to right: Force-and-deformation capture of a non-linear heterogeneous pillow; synthesized deformation with fitted material parameters; and interactive deformation synthesized with our soft tissue modeling technique.

Abstract

This paper introduces a data-driven representation and modeling technique for simulating non-linear heterogeneous soft tissue. It simplifies the construction of convincing deformable models by avoiding complex selection and tuning of physical material parameters, yet retaining the richness of non-linear heterogeneous behavior. We acquire a set of example deformations of a real object, and represent each of them as a spatially varying stress-strain relationship in a finite-element model. We then model the material by non-linear interpolation of these stress-strain relationships in strain-space. Our method relies on a simple-to-build capture system and an efficient run-time simulation algorithm based on incremental loading, making it suitable for interactive computer graphics applications. We present the results of our approach for several non-linear materials and biological soft tissue, with accurate agreement of our model to the measured data.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: physically based animation and modeling, model acquisition, deformations, data-driven graphics

1 Introduction

Recent years have witnessed significant progress and popularity of physically-based deformation models. Numerous researchers have combined Newtonian mechanics, continuum mechanics, numerical computation and computer graphics, providing a powerful




Figure 2: Acquiring and modeling non-linear quasi-static soft tissue behavior. From left to right: An object is probed with a force sensor to acquire several example deformations, the applied force direction, and the force magnitude. For every measurement we estimate its stress-strain relationship and represent it as a sample in strain space. During runtime, we interpolate these samples in strain space using radial basis functions (RBFs) to synthesize deformations for novel force inputs.

toolkit for physically-based deformations and stunning simulations, with application in feature films, video games, and virtual surgery, among others.

However, achieving realistic soft-tissue deformations requires careful choices for material models and their parameters. Many real-world objects consist of *heterogeneous materials*, requiring spatially varying material parameters such as Young’s modulus and Poisson’s ratio. Setting them is a difficult and time-consuming process. Even more challenging is the problem of *material non-linearities*. Most materials, for example rubber or biological soft tissue, show non-linear constitutive behavior, i.e., a non-linear relationship between stress and strain. Despite the wide variety of non-linear constitutive models in the literature, such as the popular hyperelastic Neo-Hookean and Mooney-Rivlin models [Ogden 1997], this is still an active research area in material science. Nonetheless, non-linear physics equations are often simplified approximations to real material behavior, and choosing the appropriate model as well as tuning its parameters are extremely complex tasks.

We present a novel data-driven representation and modeling technique for simulating non-linear heterogeneous soft tissue that simplifies the construction of convincing deformable models (Fig. 2). Our technique employs finite element methods and exploits a set of measured example deformations of real-world objects, thereby avoiding complex selection of material parameters. We transfer every measured example deformation into a local element-wise strain space, and represent this example deformation as a locally linear

sample of the material's stress-strain relation. We then model the full non-linear behavior by interpolating the material samples in strain space using radial basis functions (RBFs). Finally, a simple elastostatic finite-element simulation of the non-linearly interpolated material samples based on incremental loading allows for efficient computation of rich non-linear soft-tissue simulations.

Other earlier work in computer graphics and robotics also proposed measurement-based model fitting as a means for obtaining deformable object representations [Pai et al. 2001; Lang et al. 2002; Schoner et al. 2004], but was limited to linear material models with global support. In contrast, our work is the first to represent *complex non-linear heterogeneous materials* through spatially varying non-linear interpolation of local material properties. Our complete soft tissue capture and modeling pipeline is also distinct for its simplicity. We present a simple-to-build capture system consisting of force probes and marker-based trinocular stereo, as well as an efficient and robust algorithm for fitting the local strain-space material samples. We demonstrate the effectiveness of our soft-tissue capture and modeling method for several non-linear materials and biological soft tissue. The combination of simplicity and efficiency, both in acquisition and computation, and the high-expressiveness of the results make our technique applicable for interactive applications in computer graphics and other fields.

2 Related Work

Researchers in many fields, ranging from mechanical engineering to biology, have long studied the problem of modeling complex elasticity properties. For a recent survey of deformation models in computer graphics, please refer to [Nealen et al. 2006].

Bio-Mechanical Models For soft tissue modeling, a common approach is to devise a constitutive model [Ogden 1997] that captures in a sufficiently accurate manner the various behavior regimes of the material, and then tune the model parameters until they best fit empirical data. This approach is, however, tedious and utterly complex, as it relies on accurate modeling of tissue geometry (e.g., the bones, fat, and muscles in facial tissue), rich excitation of material regimes, and accurate measurement of forces and deformations (even in typically inaccessible regions). Despite the complexity of the approach, it has seen large application in computer graphics since the pioneering work by Terzopoulos et al. [1987], as it can lead to stunning results with the appropriate amount of effort. Some examples of complex bio-mechanical models in computer graphics include the neck [Lee and Terzopoulos 2006], the torso [Zordan et al. 2004; Teran et al. 2005; DiLorenzo et al. 2008], the face [Koch et al. 1996; Magnenat-Thalmann et al. 2002; Terzopoulos and Waters 1993; Sifakis et al. 2005], and the hand [Sueda et al. 2008].

Measurement-Based Model Fitting To circumvent the complexity of parameter tuning, several authors have proposed measurement-based model fitting approaches. The seminal work of Pai et al. [2001] presents a capture and modeling system for a deformable object's shape, elasticity, and surface roughness. Their deformable model was based on a Green's functions matrix representation [James and Pai 1999], and was later extended to increase fitting robustness [Lang et al. 2002], and to handle viscoelasticity [Schoner et al. 2004]. Our approach shares their strategy for measuring surface displacements as the result of applied surface forces, but, unlike theirs, is not limited to linear material behavior and does not rely on global response functions. Sifakis et al. [2005] give a different spin to measurement-based modeling approaches, as they learn the relationship between facial muscle activation and skin positions. Others, particularly in biomechanics, have explored measurement-based fitting of the parameters of various constitutive models, such as Young modulus estimation based on a non-linear least squares problem [Schnur and Zabarar 1992], Young modulus and Poisson ratio estimation through linear least squares [Becker and Teschner 2007], estimation of non-linear viscoelastic materi-

als [Kauer et al. 2002], or even plasticity estimation [Kajberg and Lindkvist 2004]. Our work borrows from these approaches for the estimation of each individual sample of the stress-strain relationship. However, this alone is not sufficient for capturing the rich non-linear behavior of soft tissue. In contrast to previous work, the realism of our material model is greatly enhanced with spatially varying non-linear interpolation in strain space.

Data-driven Methods Purely data-driven techniques have gained large popularity in computer graphics, as they may produce highly realistic results for phenomena that are otherwise extremely complex to model. The interpolation of lightfield samples [Buehler et al. 2001] allows simulating the illumination of complex scenes, while recent data-driven reflection models [Matusik et al. 2003] represent each BRDF through a dense set of measurements. Data-driven methods have also been applied to several other aspects of deformation modeling in computer graphics, such as facial wrinkle formation from local skin deformations [Ma et al. 2008; Bickel et al. 2008], grasping of objects [Kry and Pai 2006], skeleton-driven cloth wrinkles [Kim and Vendrovsky 2008], body-skin deformation [Park and Hodgins 2006], or learning of skeleton-driven skin dynamics [Park and Hodgins 2008]. Our method is a mixture of model fitting techniques (i.e., estimating stress-strain parameters from local measurements) and data-driven methods (i.e., using tabulated stress-strain parameters and non-linear interpolation during runtime).

Shape Modeling Another common approach in computer graphics to model deformations is shape modeling [Botsch and Sorkine 2008]. Some of the existing approaches rely on predefined examples [Sloan et al. 2001; Allen et al. 2002; Sumner et al. 2005], or even exploit interpolation [Bergeron and Lachapelle 1985; Lewis et al. 2000; Blanz et al. 2003], but these techniques cannot model deformations as a reaction to contact in the way our technique does. Some recent approaches connect shape modeling with physically-based reactive models, by rigging using templates of forces [Capell et al. 2005] or by skeletal interpolation of elastic forces [Galoppo et al. 2009], but these approaches cannot model general non-linear soft tissue.

Deformation Capture Our work captures deformation examples by combining a stereo-vision acquisition system and force sensors, similar to earlier approaches [Pai et al. 2001]. Other techniques in material sciences also directly measure the parameters of constitutive models, such as the tensile test [Hart 1967], or apparatus for in-vivo measurement through tissue aspiration [Nava et al. 2003] or indentation [Ottensmeyer and Salisbury Jr. 2004].

3 Modeling of Non-Linear Materials

In this section, we describe our representation of non-linear heterogeneous elastic materials, and how this representation is used for modeling soft tissue deformations. We first give an overview of the representation, and then describe how we parameterize the materials and how this parameterization extends from the continuum setting to a finite element discretization. We also explain how we support material non-linearities through interpolation of local linear models, and finally we describe our algorithm for computing non-linear elastostatic deformations based on incremental loading.

3.1 Overview of our Approach

In materials science, (one-dimensional) elasticity properties have long been described through stress-strain curves. Inspired by this popular representation, we opt for modeling three-dimensional elastic properties by sampling the stress-strain function at various operating regimes and interpolating these samples in strain-space (See Fig. 2).

More specifically, we characterize each sample of the stress-strain

function using a (local) linear constitutive model. Then, in order to capture material non-linearity, we define the parameter values of the constitutive model at an arbitrary operating point through scattered-data interpolation in strain-space. Moreover, in order to capture material heterogeneity, we compute both the stress-strain samples and the scattered-data interpolation in a spatially varying manner. Fig. 3 shows example deformations with color-coded Young’s modulus, which varies both as a function of the location and the local strain.

It is worth noting that our model can capture elasticity properties, but not plasticity or viscosity, among others. Our model builds on FEM and linear elasticity theory, and we refer the interested reader to books on the topic [Bathe 1995; Hughes 2000].

3.2 Discretization and Parameterization

We use linear co-rotational FEM to locally represent a deformable object’s elastic properties. In other words, given an object’s deformed configuration, we model the stress-strain relationship with linear FEM. We capture non-linearity by varying the parameters of the stress-strain relationship as a function of the strain itself. Given a displacement field \mathbf{u} , the linear co-rotational FEM employs Cauchy’s linear strain tensor $\boldsymbol{\varepsilon}(\mathbf{u}) = \frac{1}{2} (\nabla \mathbf{u} + (\nabla \mathbf{u})^T)$. Invariance of the strain under rotations is obtained by extracting the rotational part of the deformation gradient through polar decomposition, and then warping the stiffness matrix [Müller and Gross 2004].

Thanks to symmetry of the strain and stress tensors, we can represent both as 6-vectors. Given the strain tensor, we construct the 6-vector as $\boldsymbol{\varepsilon} = (\varepsilon_{xx} \ \varepsilon_{yy} \ \varepsilon_{zz} \ \varepsilon_{xy} \ \varepsilon_{xz} \ \varepsilon_{yz})^T$, and similarly for the stress. The local linear material yields then a relationship

$$\boldsymbol{\sigma}(\mathbf{u}) = \mathbf{E}\boldsymbol{\varepsilon}(\mathbf{u}) \quad (1)$$

between strain and stress. For each element (in our case, a tetrahedron), assuming locally linear isotropic material, the 6×6 stress-strain relationship matrix \mathbf{E} can be represented by Young’s modulus E and Poisson ratio ν

$$\mathbf{E} = \frac{E}{(1 + \nu)(1 - 2\nu)} (\mathbf{G} + \nu\mathbf{H}), \quad (2)$$

with the two constant matrices

$$\mathbf{G} = \text{diag}(1, 1, 1, 0.5, 0.5, 0.5) \quad (3)$$

and

$$\mathbf{H} = \begin{bmatrix} -1 & 1 & 1 & 0 & 0 & 0 \\ 1 & -1 & 1 & 0 & 0 & 0 \\ 1 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 \end{bmatrix}. \quad (4)$$

This parametrization is intuitive, where the Poisson ratio ν is unitless and describes material compressibility, while Young’s modulus E defines material elasticity. However, we employ an alternative parameterization (λ, α) that allows us to describe the stress-strain relationship as a linear function of the parameters [Becker and Teschner 2007]:

$$\mathbf{E} = \lambda\mathbf{G} + \alpha\mathbf{H}, \quad (5)$$

with

$$\lambda = \frac{E}{(1 + \nu)(1 - 2\nu)} \quad \text{and} \quad \alpha = \lambda\nu. \quad (6)$$

The parameter α is also known as Lamé’s first parameter in elasticity theory, whereas λ is not directly related to any elasticity constant. With the (λ, α) parameterization, the stiffness matrix and




Figure 3: Two examples of a deformed pillow with color-coded Young’s modulus (‘blue’ is low, ‘red’ is high), which varies both as a function of location and the local strain. Probe pressure is higher on the right.

the elastic forces become linear in the parameters. We exploit this property in our parameter fitting algorithm in Section 4.1.

The per-element stiffness matrix can be written as

$$\mathbf{K}_e = \lambda_e V_e \mathbf{B}_e^T \mathbf{G} \mathbf{B}_e + \alpha_e V_e \mathbf{B}_e^T \mathbf{H} \mathbf{B}_e, \quad (7)$$

where V_e is the volume of the element (i.e., tetrahedron), and \mathbf{B}_e is a matrix dependent on the initial position of the element’s nodes. The complete stiffness matrix is obtained by assembling the warped per-element stiffness matrices $\mathbf{R}_e \mathbf{K}_e \mathbf{R}_e^T$, where \mathbf{R}_e is the element’s rotation. By grouping all material parameters $\{\lambda_e, \alpha_e\}$ in one vector \mathbf{p} , the stiffness matrix is parameterized as $\mathbf{K}(\mathbf{p})$.

3.3 Strain-Space Interpolation

As introduced earlier, we describe the non-linear material properties through scattered-data interpolation of known local linear parameters in an element-wise manner. We obtain these known local parameters from a set of example deformations, largely simplifying an artist’s job of tuning material parameters for complex non-linear constitutive models.

Let us assume a set of M known example measurements, each with a corresponding element-wise strain vector $\boldsymbol{\varepsilon}_i \in \mathbb{R}^6$ and a parameter vector $\mathbf{p}_i = (\lambda_i, \alpha_i)^T$. Recall that we use a rotationally-invariant strain by extracting the rotation of the deformation gradient through polar decomposition [Müller and Gross 2004]. Our non-linear strain-dependent material $\mathbf{p}(\boldsymbol{\varepsilon})$ is formed by interpolating linear material samples $\mathbf{p}_i(\boldsymbol{\varepsilon}_i)$. At a given deformed configuration, the non-linear material is represented by the corresponding linear material that achieves the same force-displacement relationship. Note that we do not exploit linearization in the more traditional way of capturing the local slope of a non-linear function.

For each element, we define the stress-strain relationship through scattered-data interpolation in the strain-space \mathbb{R}^6 using radial basis functions (RBFs). The element-wise function describing the material, $\mathbf{p}(\boldsymbol{\varepsilon}) : \mathbb{R}^6 \rightarrow \mathbb{R}^2$, has the form

$$\mathbf{p}(\boldsymbol{\varepsilon}) = \sum_{i=1}^M \mathbf{w}_i \cdot \varphi(\|\boldsymbol{\varepsilon} - \boldsymbol{\varepsilon}_i\|), \quad (8)$$

where φ is a scalar basis function, and $\mathbf{w}_i \in \mathbb{R}^2$ and $\boldsymbol{\varepsilon}_i$ are the weight and feature vector for the i ’th measurement, respectively.




Figure 4: Two left-most columns: Comparison of captured and synthesized deformations for a foam block. Two right-most columns: Examples of interactive deformations produced by sliding a cylinder on top of the model.

We employ the biharmonic RBF kernel $\varphi(r) = r$. This globally supported kernel allows for smoother interpolation of sparsely scattered example poses than locally supported kernels, and avoids difficult tuning of the support radius [Carr et al. 2001].

As a preprocess, we compute the RBF weights w_i . This reduces to solving $2T$ linear $M \times M$ systems for a deformable object with T elements due to the fact that the stress-strain relationship is an element-wise description of the material. This also leads to scattered-data interpolation of the material parameters in a rather low-dimensional \mathbb{R}^6 domain. In contrast, interpolation of material properties is much more complicated in earlier approaches based on linear models with global support [Pai et al. 2001] due to the extremely high dimensionality of the parameterization.

3.4 Elastostatic FEM Simulation

We compute novel deformations using an elastostatic FEM formulation $\mathbf{K}\mathbf{u} = \mathbf{F}$, where the force \mathbf{F} includes, among others, the load produced by a contact probe. To correctly capture the material’s non-linearity during the deformation, we apply the load of the probe gradually, and solve the elastostatic FE problem for each load increment. In other words, at each loading step we measure the current strain ε , we compute the material parameters $\mathbf{p}(\varepsilon)$ by means of the interpolation described above, we formulate the elastostatic problem, and we solve it for the new deformations. The incremental loading procedure ensures that the non-linearity of the material is correctly captured during the complete deformation process, with the material parameters depending on the strain at all times.

For contact handling, we compute a distance field for the rigid probe object that produces the deformations. We test for collisions between points on the deformable object and the distance field and, upon collision, we compute the penetration depth and direction. We then define a linear force field at each colliding point and solve the FEM simulation through iterative quasi-static simulation. At each iteration of the quasi-static FEM simulation, we first compute the material parameters for the current configuration based on the interpolation algorithm described above. Then, given the stiffness matrix and the linear collision force field, we define a quasi-static problem and solve for the new displacements. We compute several iterations until an equilibrium is reached.

4 Fitting the Material Parameters

We now describe how we compute the actual material parameters for a given object. This consists of two parts: First, estimating parameter values for each deformation example, and second, selecting a suitable basis from all the deformation examples.

4.1 Parameter Estimation Algorithm

In order to estimate a sample of the stress-strain relationship, we apply a known input force to the object under study. For each captured deformation we can distinguish three different regions on the object’s surface: (i) the probing region, with measured non-zero forces and measured displacements, (ii) the attached region, with unknown forces and zero displacements, and (iii) the free region, with zero forces and measured displacements. We use $\bar{\mathbf{x}}$ and $\bar{\mathbf{F}}$ to denote the vectors of known displacements and forces, respectively, at the points corresponding to mesh nodes in the model.

Given measured displacements and forces, we compute spatially varying material parameters \mathbf{p} as:

$$\hat{\mathbf{p}} = \arg \min_{\mathbf{p}} \left\{ \sum_{i=1}^n \|\mathbf{x}_i(\mathbf{p}, \bar{\mathbf{F}}) - \bar{\mathbf{x}}_i\|^2 + \gamma \|\mathbf{L}\mathbf{p}\|^2 \right\}, \quad (9)$$

where $\mathbf{x}_i(\mathbf{p}, \bar{\mathbf{F}})$ denotes the position of a mesh node as a function of material parameters and the measured forces. The sparse Laplacian matrix \mathbf{L} enforces spatial smoothness of parameters. We employ the umbrella operator [Zhang 2004] $(\mathbf{L}\mathbf{p})_i = \sum_j w_{i,j}(\mathbf{p}_i - \mathbf{p}_j)$, where i and j refer to tetrahedron labels, and $w_{i,j} = 1$ iff two tetrahedra share a vertex. This regularization is required to prevent overfitting due to noise in the acquired data. This is also mathematically required to obtain a well-posed problem because the number of parameters is always twice the number of tetrahedra, $|\mathbf{p}| = 2T$, whereas the number of measured positions $|\bar{\mathbf{x}}| = n$ may be smaller, which would result in an underconstrained problem. We also considered scattered data interpolation of material parameters in object space as an alternative for addressing the underconstrained problem, but it would be difficult to decide where to place the samples for highly heterogeneous objects.

We minimize the non-linear residual Eq. (9) iteratively using the Levenberg-Marquardt algorithm [Levenberg 1944]. We derive the Jacobian matrix in the Appendix. Instead of defining the residual in

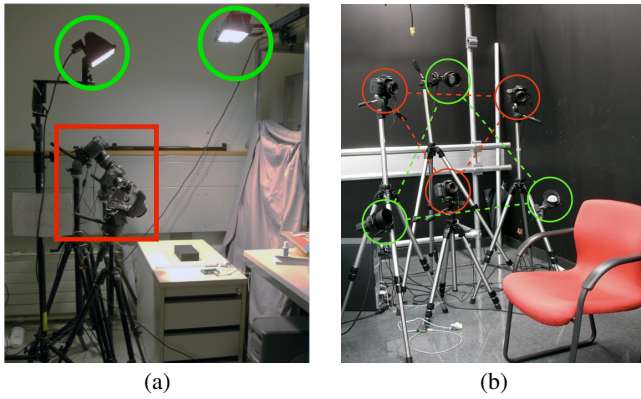



Figure 5: Our trinocular stereo vision system consists of three high-resolution cameras (indicated in red) and two to three light sources (indicated in green). The cameras are arranged in a triangular setup, which helps maximize visibility during capture of a contact interaction. The light sources ensure uniform illumination during the acquisition.

terms of measured positions, the error functional could also be described in terms of measured forces [Becker and Teschner 2007], yielding a linear optimization problem. However, our observations have shown that this approach is unstable when the force-displacement relationship is not close to linear material behavior.

4.2 Strain-Space Basis Selection

A material capture session consists of capturing N example deformations, from which we obtain the training dataset of N parameter vectors for each element in the mesh. However, this dataset may be rather large, and we are interested in selecting a compact set of M basis parameter vectors for each element. Note that M need not be the same for all elements.

We select the basis in the same greedy manner as proposed by [Carr et al. 2001]. We start by setting a parameter vector at zero strain with the average parameters computed for very small-strain deformations. We then add the parameter vector with largest error, until a given error tolerance is achieved. After each parameter vector is added to the basis, we need to compute the RBF weights that best fit the parameter vectors for all N example deformations in a least-squares manner, as described in Eq. (8). This figure shows the evolution of the fitting error for the foam block in Fig. 4. This error plot accumulates the error for all captured deformations, not only those added to the basis. The error drops quickly after adding the second parameter vector to the basis because the first vector may not represent the average material behavior well. See Section 6 for more details on the validation of our method.



5 Data Acquisition

We developed a simple data acquisition system consisting of force probes and a marker-based trinocular stereo system. Deformations are induced by physical interaction with the object. We decided to use a marker-based system due to its simplicity, robustness, and independence of the object’s surface properties.

Figure 5 shows our trinocular stereo vision system, consisting of




Figure 6: (a) Contact probe with integrated force sensor. (b) From left to right: USB Interface Kit, Force Sensing Resistor (red circle), Phidget Voltage Divider, and connection cable.

three Canon 40D cameras that capture images at a resolution of 3888×2592 . These cameras are placed in a triangular configuration to minimize occlusions caused by the contact probes during data acquisition. We built an external trigger device to synchronize the three cameras, and use additional light sources to ensure uniform illumination during the acquisition process. The surface displacement during static deformations is measured using a set of markers that we paint on the object’s visible surface. Our system is capable of measuring viewpoint-registered marker positions to an accuracy of < 1 mm.

We built contact probes with arbitrary shapes and circular disks of different diameters attached to the tip of a long screwdriver (see Fig. 6). We estimate the position and orientation of the contact probe using two markers on the white shaft of the screwdriver. To measure the magnitude of the contact forces we use a 0.2 inch Force Sensing Resistor (FSR) (Item S-20-1000-FS2) connected to a Phidget Voltage Divider (Item S-50-P1121) and USB Interface Kit 8/8/8 (Item C-200-P1018) by Trossen Robotics. The force sensor’s read operation is synchronized with the external camera trigger signal.

6 Results

Model Evaluation. We have evaluated the quality of our soft-tissue capture and modeling technique on several real-world objects, including two foam blocks, a heterogeneous soft pillow, and a human face. The data is publicly available at the authors’ web sites.

Fig. 4 shows a foam block with homogeneous material. We acquired 48 deformation examples, well distributed over the foam to induce deformations in all 1,805 tetrahedra of our model. We then constructed the non-linear material representation, with bases of 8 samples per tetrahedron on average, using the procedure in Section 4.1. Even though the object is homogeneous, it should be noted that the material parameters that were estimated for each input example are non-homogeneous due to non-linearities in the stress-strain relationship. The average fitting error for the captured deformations is less than 1 mm (see figure in Section 4.2). Fig. 4 shows synthesized deformations produced with our technique using a probe with a larger, different contact area than the probe used for data acquisition.

To compare our model to a uniform linear co-rotational model we use the homogeneous foam shown in Fig. 7 and Fig. 8. We captured 12 deformation examples with the probe near the center of the block and modeled the object with 3,240 tetrahedra. We computed an average-fit linear co-rotational model that best approximates all the input deformations. As shown in Fig. 8, our model (blue) accurately captures the hyperelastic behavior of the foam, while the average-fit linear co-rotational model (green) underestimates the deformation at small force values and overestimates it at large ones. In addition, the linear co-rotational model suffers from element inversion for large forces.

Our model is of course not confined to the contact shapes that were




Figure 7: Comparing real (top) and modeled (bottom) deformations with a different contact probe than the one used in the data acquisition phase.




Figure 8: Comparison of deformations using our method vs. an average-fit linear co-rotational model.

used during data acquisition. Fig. 7 shows a side-by-side comparison of our model (bottom) to real deformations (top) using a different contact probe than the circular one we used for data acquisition. We captured the applied force with the new contact probe, and then distribute it uniformly in the simulated setting. The figure shows high correspondence between the real and simulated scenarios. We refer the reader to the accompanying video for an animated side-by-side comparison.

To evaluate the sensitivity of our capture and modeling approach to measurement noise we created example deformations of a virtual block with three layers of user-defined non-linear materials. We then evaluated the accuracy in matching these deformations with our model under different levels of noise in the input data. Specifically, we applied Gaussian noise with a variance of 10%, 20% and 30% to the input displacements and then measured the L^2 error for all deformations and error levels. On average, we obtain an error of 0.3% of the maximum displacement for the case without error, and 2.1%, 3.1% and 4.4% for the cases with 10%, 20% and 30% input noise, respectively.




Fig. 9 shows a pillow object with heterogeneous behavior even in its rest state. The screenshots compare the captured deformations with the deformations of the 1,691 tetrahedra model synthesized with our algorithm. The figure also shows screenshots of deformations at interactive frame rates of about 10 Hz on a standard PC.

Facial Deformation. We have also applied our soft-tissue capture and modeling technique to the challenging task of facial deformations, as shown in Fig. 10. We have modeled the facial tissue with a single layer of 8,261 tetrahedra that are attached to a low-resolution skull model. To model the sliding contacts between the tissue and the skull we use the same contact handling as for the probe object (see Section 3.4). Given the deformation of the tetrahedral mesh, we compute the deformation of a high-resolution triangle mesh using a smooth embedding based on moving least squares interpolation like Kaufmann et al. [2008].

Note that our face model does not correctly capture all types of deformations because we use a model with closed lips, and all the deformation examples in the training dataset were captured with relaxed muscles and closed jaw. Nevertheless, the model is able to produce compelling deformations even without anatomically correct modeling of the musculoskeletal structure of the face.

7 Discussion

We have presented a novel data-driven method for modeling non-linear heterogeneous soft tissue. The major practical contribution of our work is the ability to model rich non-linear deformations in a very simple manner, without the complex task of carefully choosing material models and parameters. Instead, our data-driven method relies on a simple-to-build acquisition system, a novel representation of the material through spatially varying interpolation of fitted linear models, and a simple deformation synthesis method.

Our work suggests a highly innovative approach to non-linear material modeling, but it also suffers from limitations. Due to its formulation, our technique is currently limited to capturing elastic properties. A fully dynamic simulation of soft tissue would require capturing other properties such as viscosity and plasticity. One interesting conclusion of our work is that it is often possible to obtain compelling surface deformations with a volumetric meshing unaware of an object’s actual volumetric structure. This is of course not valid for all situations. For example, our face model could be greatly enhanced with accurate lip contact and jaw motion models.

There are several aspects of our model that deserve further exploration. One of them is its ability for capturing anisotropic behavior. The underlying linear co-rotational material model that we use for representing deformation samples can only capture isotropic behavior, but deformation samples with the same total strain but in different directions will lead to anisotropic behavior. In other words, we locally model the material isotropic in strain space, yet strain-space interpolation of material parameters provides global anisotropic behavior. It is worth exploring to what extent our approach captures anisotropy.




Figure 9: Two left-most columns: Comparisons of captured and synthesized deformations for a heterogeneous non-linear pillow. Right column: Interactive deformations of the model produced by pushing (top) and pulling (bottom).




Figure 10: Left: Capture of facial deformations; Middle: Synthesized deformations for the captured examples; Right: Frames of an animation with a cylindrical probe pressing on the cheek.

Another aspect that deserves further analysis is the formulation of the quasi-static deformation problem. Given a certain strain, we employ a local linear co-rotational model to formulate a quasi-static deformation problem. However, our model is not strictly a local linearization, which means that the stiffness matrix of the quasi-static deformation problem does not employ correct force derivatives. At the same time, our linear model is more robust than a model obtained by local differentiation and avoids non-passive regimes.

Similar to other approaches, our parameter fitting algorithm is formulated as a minimization problem and may end up in a local minimum. In fact, we have identified fitting error as the major source of potential inaccuracies in the deformation synthesis. Sometimes, fitting error also appears because we limit Poisson's ratio to physically valid values during the minimization. Robust parameter identification is still an open research problem in materials science, and some recent approaches explore alternative solutions including particle filters [Burion et al. 2008]. Multi-resolution fitting may be another way of increasing robustness.

Currently, we only measure the force in direction of the probe's shaft without measuring tangential forces and friction behavior. More accurate force sensors could capture such effects. Finally, one could build a fully automated capture system using robotics. Using a more efficient parameter estimation algorithm for material fitting, one could evaluate the need for further samples of the stress-strain relationship online, and determine the optimal probing patterns on the fly.

Acknowledgements

We would like to thank the anonymous reviewers for their helpful comments, and thank Manuel Lang, Peter Kaufmann, Sebastian Martin, Giuseppe Barbarino, Guillermo Diez-Caas, Kalyan Sunkavalli, and the members of the CG Lab in Zurich. This research was supported by the NCCR Co-Me grant of the Swiss National Science Foundation. Moritz Bächer was partially supported by the Initiative in Innovative Computing (IIC) and the School of Engineering and Applied Sciences (SEAS) at Harvard.

References

- ALLEN, B., CURLESS, B., AND POPOVIĆ, Z. 2002. Articulated body deformation from range scan data. *ACM Trans. Graph.* 21, 3, 612–619.
- BATHE, K.-J. 1995. *Finite Element Procedures*. Prentice-Hall.
- BECKER, M., AND TESCHNER, M. 2007. Robust and efficient estimation of elasticity parameters using the linear finite element method. In *SimVis*, 15–28.
- BERGERON, P., AND LACHAPPELLE, P., 1985. Controlling facial expression and body movements in the computer generated short “Tony de Peltrie”. *Siggraph Course Notes*.
- BICKEL, B., LANG, M., BOTSCH, M., OTADUY, M. A., AND GROSS, M. 2008. Pose-space animation and transfer of facial details. In *Proc. of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, 57–66.
- BLANZ, V., BASSO, C., POGGIO, T., AND VETTER, T. 2003. Re-animating faces in images and video. *Computer Graphics Forum* 22, 3 (Sept.), 641–650.
- BOTSCH, M., AND SORKINE, O. 2008. On linear variational surface deformation methods. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 14, 1, 213–230.
- BUEHLER, C., BOSSE, M., McMILLAN, L., GORTLER, S., AND COHEN, M. 2001. Unstructured lumigraph rendering. In *Proc. of ACM SIGGRAPH*, ACM, 425–432.
- BURION, S., CONTI, F., PETROVSKAYA, A., BAUR, C., AND KHATIB, O. 2008. Identifying physical properties of deformable objects by using particle filters. In *Proc. of the International Conference on Robotics and Automation*, 1112–1117.
- CAPELL, S., BURKHART, M., CURLESS, B., DUCHAMP, T., AND POPOVIĆ, Z. 2005. Physically based rigging for deformable characters. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, 301–310.
- CARR, J. C., BEATSON, R. K., CHERRIE, J. B., MITCHELL, T. J., FRIGHT, W. R., MCCALLUM, B. C., AND EVANS, T. R. 2001. Reconstruction and representation of 3D objects with radial basis functions. In *Proc. of ACM SIGGRAPH*, 67–76.
- DILORENZO, P., ZORDAN, V., AND SANDERS, B. 2008. Laughing Out Loud: Control for modeling anatomically inspired laughter using audio. *ACM Trans. Graph. (Proc. of ACM SIGGRAPH Asia)* 27, 5.
- GALOPPO, N., OTADUY, M. A., MOSS, W., SEWALL, J., CURTIS, S., AND LIN, M. C. 2009. Controlling deformable material with dynamic morph targets. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*.
- HART, E. W. 1967. Theory of the tensile test. *Acta Metallurgica* 15, 351–355.
- HUGHES, T. J. R. 2000. *The Finite Element Method. Linear Static and Dynamic Finite Element Analysis*. Dover Publications.
- JAMES, D. L., AND PAI, D. K. 1999. ArtDefo: Accurate real time deformable objects. In *Proc. of ACM SIGGRAPH*, ACM Press/Addison-Wesley Publishing Co., 65–72.
- KAJBERG, J., AND LINDKVIST, G. 2004. Characterisation of materials subjected to large strains by inverse modelling based on in-plane displacement fields. *International Journal of Solids and Structures* 41, 13, 3439–3459.
- 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.
- KAUFMANN, P., MARTIN, S., BOTSCH, M., AND GROSS, M. 2008. Flexible simulation of deformable models using discontinuous galerkin fem. *Proceedings of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, 105–115.
- KIM, T.-Y., AND VENDROVSKY, E. 2008. Drivenshape - a data-driven approach to shape deformation. In *Proc. of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation*.
- KOCH, R. M., GROSS, M. H., CARLS, F. R., VON BÜREN, D. F., FANKHAUSER, G., AND PARISH, Y. 1996. Simulating facial surgery using finite element methods. In *Proc. of ACM SIGGRAPH*, 421–428.
- KRY, P. G., AND PAI, D. K. 2006. Interaction capture and synthesis. *ACM Trans. Graph. (Proc. of ACM SIGGRAPH)* 25, 3, 872–880.
- LANG, J., PAI, D. K., AND WOODHAM, R. J. 2002. Acquisition of elastic models for interactive simulation. *International Journal of Robotics Research* 21, 8, 713–733.
- LEE, S.-H., AND TERZOPOULOS, D. 2006. Heads up!: Biomechanical modeling and neuromuscular control of the neck. *ACM Trans. Graph. (Proc. of ACM SIGGRAPH)* 25, 3.
- LEVENBERG, K. 1944. A method for the solution of certain nonlinear problems in least squares. *The Quarterly of Applied Mathematics*, 2, 164–168.
- LEWIS, J. P., CORDNER, M., AND FONG, N. 2000. Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. In *Proc. of ACM SIGGRAPH*, 165–172.

MA, W.-C., JONES, A., CHIANG, J.-Y., HAWKINS, T., FREDERIKSEN, S., PEERS, P., VUKOVIC, M., OUHYOUNG, M., AND DEBEVEC, P. 2008. Facial performance synthesis using deformation-driven polynomial displacement maps. *ACM Trans. Graph. (Proc. of ACM SIGGRAPH Asia)* 27, 5.

MAGNENAT-THALMANN, N., KALRA, P., LÉVÊQUE, J. L., BAZIN, R., BATISSE, D., AND QUELEUX, B. 2002. A computational skin model: fold and wrinkle formation. *IEEE Trans. on Information Technology in Biomedicine* 6, 4, 317–323.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)* 22, 3, 759–770.

MÜLLER, M., AND GROSS, M. 2004. Interactive virtual materials. In *GI '04: Proceedings of Graphics Interface 2004*, Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 239–246.

NAVA, A., MAZZA, E., KLEINERMANN, F., AVIS, N. J., AND MCCLURE, J. 2003. Determination of the mechanical properties of soft human tissues through aspiration experiments. In *Proc. of MICCAI*, 222–229.

NEALEN, A., MLLER, M., KEISER, R., BOXERMAN, E., AND CARLSON, M. 2006. Physically based deformable models in computer graphics. *Computer Graphics Forum* 25, 4 (Dec.), 809–836.

OGDEN, R. W. 1997. *Non-Linear Elastic Deformations*. Courier Dover Publications.

OTTENSMEYER, M. P., AND SALISBURY JR., J. K. 2004. In-vivo data acquisition instrument for solid organ mechanical property measurement. In *Proc. of MICCAI*, 975–982.

PAI, D. K., VAN DEN DOEL, K., JAMES, D. L., LANG, J., LLOYD, J. E., RICHMOND, J. L., AND YAU, S. H. 2001. Scanning physical interaction behavior of 3d objects. In *Proceedings of ACM SIGGRAPH*, 87–96.

PARK, S. I., AND HODGINS, J. K. 2006. Capturing and animating skin deformation in human motion. *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)* 25, 3.

PARK, S. I., AND HODGINS, J. K. 2008. Data-driven modeling of skin and muscle deformation. *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)* 27, 3.

SCHNUR, D. S., AND ZABARAS, N. 1992. An inverse method for determining elastic material properties and a material interface. *International Journal for Numerical Methods in Engineering* 33, 10, 2039–2057.

SCHONER, J. L., LANG, J., AND SEIDEL, H.-P. 2004. Measurement-based interactive simulation of viscoelastic solids. *Computer Graphics Forum (Proc. Eurographics)* 23, 3, 547–556.

SIFAKIS, E., NEVEROV, I., AND FEDKIW, R. 2005. Automatic determination of facial muscle activations from sparse motion capture marker data. *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)* 24, 3, 417–425.

SLOAN, P.-P. J., ROSE, III, C. F., AND COHEN, M. F. 2001. Shape by example. In *I3D '01: Proceedings of the 2001 symposium on Interactive 3D graphics*, ACM, New York, NY, USA, 135–143.

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

SUMNER, R. W., ZWICKER, M., GOTSMAN, C., AND POPOVIĆ, J. 2005. Mesh-based inverse kinematics. In *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH)*, vol. 24, 488–495.

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

TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K. 1987. Elastically deformable models. In *Proc. of ACM SIGGRAPH* 87, 205–214.

TERZOPOULUS, D., AND WATERS, K. 1993. Analysis and synthesis of facial image sequences using physical and anatomical models. *IEEE Trans. PAMI* 14 (June), 569–579.

TOLEDO, S., CHEN, D., AND ROTKIN, V. 2003. Taucs: A library for sparse linear solvers.

ZHANG, H. 2004. Discrete combinatorial laplacian operators for digital geometry processing. In *Proc. of SIAM Conference on Geometric Design and Computing*, Nashboro Press, 575–592.

ZORDAN, V., CELLY, B., CHIU, B., AND DILORENZO, P. C. 2004. Breathe easy: Model and control of human respiration for computer animation. In *Proc. of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, 29–38.

Appendix: Jacobian for Parameter Fitting

During fitting of material parameters $\mathbf{p} = \{\lambda_e, \alpha_e\}$ through minimization of Eq. (9), we need to compute the Jacobian of the deformed vertex positions w.r.t. the parameters, i.e., $\mathbf{J} = \frac{\partial \mathbf{x}}{\partial \mathbf{p}}$, in each iteration of the Levenberg-Marquardt algorithm.

Given external forces \mathbf{F} and initial positions \mathbf{x}_0 , the deformed positions under the linear co-rotational elastostatic problem [Müller and Gross 2004] are

$$\mathbf{x} = \mathbf{K}^{-1} (\mathbf{F} + \mathbf{K}' \mathbf{x}_0), \quad (10)$$

$$\text{with } \mathbf{K} = \sum_e [\mathbf{R}_e \mathbf{K}_e \mathbf{R}_e^T]_e \text{ and } \mathbf{K}' = \sum_e [\mathbf{R}_e \mathbf{K}'_e]_e.$$

Here $[\dots]_e$ denotes the assembly of the submatrix of the e -th element into the complete stiffness matrix.

The Jacobian w.r.t. each parameter $p_i \in \{\lambda_e, \alpha_e\}$ can then be computed as

$$\mathbf{J}_i = \frac{\partial \mathbf{K}^{-1}}{\partial p_i} (\mathbf{F} + \mathbf{K}' \mathbf{x}_0) + \mathbf{K}^{-1} \frac{\partial \mathbf{K}'}{\partial p_i} \mathbf{x}_0, \quad (11)$$

$$\text{with } \frac{\partial \mathbf{K}^{-1}}{\partial p_i} = -\mathbf{K}^{-1} \frac{\partial \mathbf{K}}{\partial p_i} \mathbf{K}^{-1}.$$

Note that we do not compute the inverse of \mathbf{K} . Instead, we compute a sparse Cholesky factorization [Toledo et al. 2003], and then use this factorization many times for solving the linear systems above.

Recall the expression for the (unwarped) per-element stiffness matrix in Eq. (7). The remaining terms are defined as:

$$\frac{\partial \mathbf{K}}{\partial \lambda_e} = [\mathbf{V}_e \mathbf{R}_e \mathbf{B}_e^T \mathbf{G} \mathbf{B}_e \mathbf{R}_e^T]_e, \quad \frac{\partial \mathbf{K}'}{\partial \lambda_e} = [\mathbf{V}_e \mathbf{R}_e \mathbf{B}_e^T \mathbf{G} \mathbf{B}_e]_e,$$

$$\frac{\partial \mathbf{K}}{\partial \alpha_e} = [\mathbf{V}_e \mathbf{R}_e \mathbf{B}_e^T \mathbf{H} \mathbf{B}_e \mathbf{R}_e^T]_e, \quad \frac{\partial \mathbf{K}'}{\partial \alpha_e} = [\mathbf{V}_e \mathbf{R}_e \mathbf{B}_e^T \mathbf{H} \mathbf{B}_e]_e. \quad (12)$$

In the case when some nodes are constrained not to deform (e.g., when the bottom of the captured objects is fixed), their known positions move to the right-hand side in Eq. (10), and the Jacobians must be slightly modified.