

Data-Driven Estimation of Cloth Simulation Models

E. Miguel^{1,2} D. Bradley² B. Thomaszewski² B. Bickel² W. Matusik^{2,3} M. A. Otaduy¹ S. Marschner⁴

1. URJC Madrid 2. Disney Research Zurich 3. MIT CSAIL 4. Cornell University

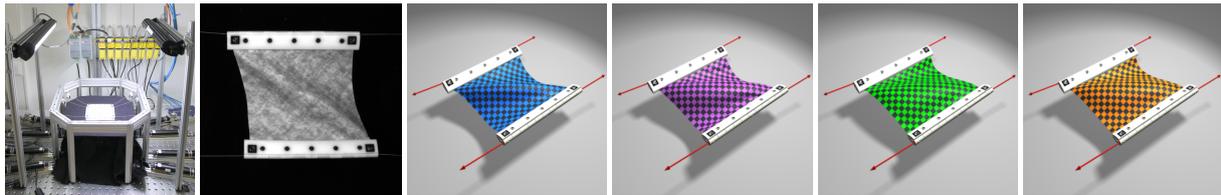


Figure 1: We capture deformation behaviors of cloth materials with a dedicated setup (column 1 from left). The measurement images (2) are reconstructed into 3D geometry (3) yielding dense deformation fields. We use this data to fit parameters and investigate approximation qualities of three common cloth models: springs (4), soft constraints (5), and the StVK model (6).

Abstract

Progress in cloth simulation for computer animation and apparel design has led to a multitude of deformation models, each with its own way of relating geometry, deformation, and forces. As simulators improve, differences between these models become more important, but it is difficult to choose a model and a set of parameters to match a given real material simply by looking at simulation results. This paper provides measurement and fitting methods that allow nonlinear models to be fit to the observed deformation of a particular cloth sample. Unlike standard textile testing, our system measures complex 3D deformations of a sheet of cloth, not just one-dimensional force–displacement curves, so it works under a wider range of deformation conditions. The fitted models are then evaluated by comparison to measured deformations with motions very different from those used for fitting.

1. Introduction

Today's cloth simulators for animation, visual effects, games, and apparel design can mimic real cloth to a high degree of fidelity. But to fully exploit their capabilities, the constitutive models for cloth deformation must be tuned with great care. During this tuning process it is difficult to tell which models and which parameters are giving results more like the real material. This paper aims to solve this problem by introducing new techniques to measure complete cloth behavior under controlled conditions and to estimate cloth deformation models from these measurements.

Most methods for testing cloth move the sample into a state of near-uniform strain, exercising one or at most two components of strain at once: pure stretching, pure shearing, or pure bending. One or two forces are measured to quantify the cloth's resistance to deformation, and the resulting force-displacement curves are valuable in studying the differences

between materials. However, this approach has certain limitations. The inevitable deviations from uniform strain create modeling error that cannot be quantified without knowing the actual strain variation; and force-displacement curves can be used directly to tune a cloth model, but do not provide any way to validate the resulting fit.

The contributions of this paper are, first, a new, general system for observing cloth properties that measures more complete data than previous work in cloth capture or textile testing, and second, a new method for fitting parametric models to this type of data. Finally we show results that illustrate the performance of several widely used cloth models.

Our measurement system applies forces to a sample of cloth using actuators and force sensors that let us know the complete applied force, in 3D. The resulting deformation is tracked by a stereo computer vision system that captures the complete deformation, also in 3D. Having deformation and

force information (rather than just force and overall displacement, or shape information without correspondence) makes our data well suited to model validation—the experiment measures the complete answer that should be predicted by a cloth simulator. Also, we do not need uniform strain, and in this paper we illustrate a range of tests, some that mimic traditional tests and some with more complex deformations.

Our approach to model estimation is to numerically optimize nonlinear stress-strain curves to minimize errors in force and position compared to the measurement. We have designed a general fitting method, suited for the vast majority of existing cloth models, that leverages equilibrium conditions to guide the iteration. By estimating model parameters under a sequence of deformations of increasing complexity, we alleviate problems with convergence in the presence of abundant local minima.

We have used our system to fit three membrane models and two bending models from the graphics literature, each based on a different strain measure, and to evaluate the resulting models against more complex motions.

2. Prior Work

Cloth simulation has a comparatively long history in computer graphics. Since the first physics-based approach by Terzopoulos et al. [TPBF87] a multitude of different cloth models have emerged, ranging from simple mass-spring systems [Pro95, CK02] over general particle systems [BHW94, BW98, EWS96] to elaborate models derived from continuum mechanics [EKS03, VMTF09] or even the discrete yarn structure [KJM08]. Considering the number of existing models, it is very hard to clearly identify or even quantify the advantages of individual approaches. Our goal is to define a platform for comparing cloth models to the observed behavior of real cloth.

As a central component of any cloth model, material models describe the relation between deformation and resulting forces. In the simplest case, this relationship is linear and thus completely described by a set of material constants, whose number depends on the cloth model: Mass-spring systems typically have four parameters (one per type of spring [Pro95]), whereas continuum-based methods can have two (isotropic) to six (completely anisotropic) parameters for planar deformation [EKS03] plus another one to three for bending [GGWZ07].

Continuum-based approaches can accurately describe the directional variation of material properties, but regardless of the cloth model, a single set of material coefficients for the entire deformation range is not sufficient to faithfully capture the nonlinear response of typical fabrics. Bi-phasic models, typically implemented as strain limiting methods [Pro95, BFA02, Mü08, TPS09, WOR10], improve on this by splitting the material behavior into an initial, weakly elastic range and a stiff, quasi-inextensible limit. At the extreme,

the elastic range can be replaced altogether by inextensibility constraints [GHF*07, EB08].

A better approximation to the true material response can be obtained by making the material parameters functions of the deformation, rather than constants, and by fitting these functions to measured data. To this end, previous work [BHW94, EWS96, VMTF09] has mainly relied on the Kawabata Evaluation System (KES) [Kaw80] and corresponding machinery. While the KES covers a comprehensive set of experiments, other devices have been used in more specific context such as the Picture Frame test [Cu179] for measuring shear properties and the Cantilever test [CPGE90] for measuring bending properties (see also Pabst et al. [PKST08]).

These measurement-based approaches establish a valuable link between simulation and real-world behavior, but they rely on experiments that isolate individual deformation modes. As an alternative, Bhat et al. [BTH*03] (and recently Kunitomo et al. [KNM10]) aim at avoiding the need for controlled conditions and try to extract parameters from casually captured videos of cloth. This approach appeals through a simple and inexpensive acquisition process, but it is not possible to accurately separate internal (i.e. material-specific) and external (e.g. friction, air drag) parameters. In a similar spirit, capture technology can be used to record time-varying geometry of complex cloth motions [WCF07, BPS*08, SGdA*10]. But while capturing can provide accurate deformation data, parameter fitting remains very difficult without explicit control over boundary conditions, in particular loading forces.

Closer to our work is the recent approach of Wang et al. [WRO11], who propose a data-driven piecewise linear elastic cloth model comprising 39 material parameters. These parameters are fitted to experimentally acquired data obtained from planar and bending deformations. Their capture setup is appealingly simple, but ours is more general and powerful: it produces a 3D surface, rather than a 2D deformation, and it measures all forces applied to the cloth as they change during a range of different deformations.

3. Measurement system

The design goals of our measurement system are to create deformations in a sample of cloth that explore a substantial range of the material's strain space, and to record complete information about the forces applied to the cloth and the deformation that it undergoes. Like other cloth testing systems, we focus primarily on tensile forces, because it is hard to repeatedly produce and measure compression forces in a sheet that is inclined to buckle.

Tests are performed on 100 mm square cloth samples using two kinds of plastic clips: small, rounded clips that grab a localized area, and long clips that grip one whole side of the sample. We measure the weights of all cloth samples as well as the clips (see Table 1) and use these values in the optimization process. Forces are applied to the clips by fine

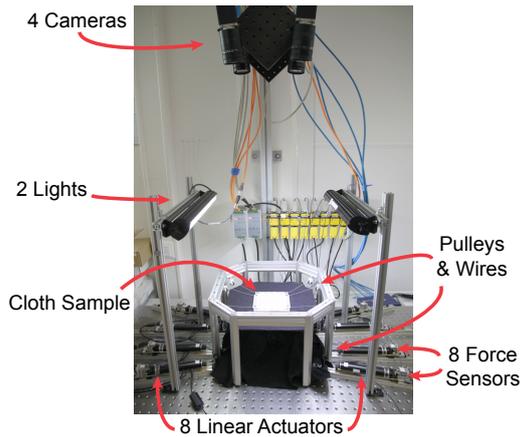


Figure 2: Acquisition setup for the measurement system.

wire cords that are pulled to defined displacements by eight linear actuators, and the tension in the cords is monitored by miniature load cells located at the actuator ends (see Figure 2). Our actuators and load cells are capable of applying and measuring tensions up to 45 N, but in our experiments the maximum force is typically on the order of 10 N.

The geometry of the cloth sample and the attached clips is monitored by a vision system composed of four high-resolution cameras. The location and orientation of the cords attached to the clips (which reveal the direction of the applied force) are also tracked. Each output frame of a measurement session contains:

- The configuration of the cloth sample, represented as a deformed mesh with 10K regularly sampled vertices.
- The positions and orientations of all clips attached to the cloth, including a list of clamped cloth vertices.
- The forces applied to all clips. The magnitudes are determined by the tension measurements, and the directions are determined by the observed directions of the cords.

Note that the actuator positions themselves are not part of the output, since they are superseded by the displacements measured at the clips. This prevents stretching of the cord, or other factors affecting the distance between the clip and the actuator, from affecting displacement accuracy.

3.1. Reconstruction

Our vision system recovers the space-time geometry of the deforming cloth and attached rigid clips, as well as the directions of the forces applied to the clips.

Initialization. The cloth sample starts flat on a table and we capture the rest pose without applied tensile forces. This initial frame serves to compute the geometry of the cloth without any occlusion from clips. We then attach the clips, and the measurement process continues automatically, following a defined script of actuations, and recording images

Cloth Sample	Id	Mass (g)
cotton satin	#4	1.2
rayon/spandex knit	#12	3.1
cotton denim	#14	4.6
wool/cotton blend	#18	2.4
plastic clips (3 sizes)		1.9, 10.1, 13.3

Table 1: Cloth and attachment clip masses.

and forces. We typically deform the cloth by moving the actuators at 0.5 mm/sec and capture a frame every 2 seconds.

Cloth Geometry Reconstruction. The raw data for a single deformation consists of 20 to 200 individual measurement frames, with a set of camera images and simultaneous force sensor readings for each frame.

We compute the per-frame geometry using a state-of-the-art stereo reconstruction technique [BBH08], which was specifically tailored for reconstructing cloth geometry [BPS*08]. If the inherent texture of the cloth is not sufficiently random, it is printed with a wavelet noise pattern [AIH*08] to provide texture that can be used for stereo reconstruction and tracking. The pattern is printed with a flatbed inkjet printer and does not have a noticeable effect on the material behavior.

To represent inter-frame correspondence, we use optical flow to obtain a single triangle mesh that deforms over time, akin to the human face tracking method of Bradley et al. [BHPS10]. To start, the cloth vertices in the rest pose frame (frame 0) are projected onto the input images, where optical flow predicts the projection of each vertex at the next time step. Back-projecting onto the reconstructed geometry for the next frame gives new position estimates for the cloth vertices. The process is then repeated using the result from frame n to obtain frame $n + 1$. As with all sequential tracking methods, very small errors can accumulate over time and cause temporal drift in the reconstruction. To avoid drift, we subsequently match each frame independently back to the rest pose frame using the approach described in Bradley et al. [BHPS10]. The final solution is smoothed using Laplacian regularization to remove noise.

Tracking Clips and Cords. In order to measure the complete answer that a simulator should predict, we need to determine the interaction between the rigid clips, the cloth, and the cords. The clips are produced, using rapid prototyping, with embedded codes [Fia05] that allow us to determine their identity, position, and orientation automatically. The area of cloth occluded by the clips is used to automatically determine which cloth vertices are clamped by each clip and will therefore be constrained to it in the simulator.

The vision system also finds the cords in the images and triangulates a 3D line for each cord. A few user scribbles on an input image indicate which cords are affecting each

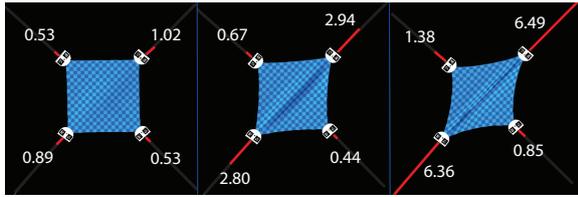


Figure 3: Force measurements for selected frames of a corner pulling sequence. Forces are rendered as red vectors with magnitudes proportional to their values (in Newtons).

clip. Figure 3 illustrates the force measurements and clip locations for three different frames from one experiment. The forces are rendered as red vectors with lengths proportional to the force magnitudes.

3.2. Measurements

The set of deformations to measure is motivated by the goals of the parameter fitting stage (Section 5): to fit model parameters for stretch, shear and bending that best describe the cloth, and to validate the parameter fits by comparing against other measurements.

To reduce the risk of falling into local minima during parameter fits, we have designed deformation sequences that produce near-isolated strains, and allow estimating stretch, shear and bending properties in a separate and incremental manner. Unlike standard textile evaluation practices [Kaw80], we relax the requirement of uniform strains. To isolate stretching we perform a uni-axial tension experiment, with forces applied to two long bar clips attached to either side of the cloth (see Figure 4, 2nd column). The cloth is slowly stretched until a maximum force is reached and then slowly released back. The process is repeated three times, in both weft and warp directions separately.

Shearing is captured using an approximate *picture-frame* experiment [Cul79], where four long clips fix the cloth boundaries and shear stress is applied as the cords pull on opposite corners (Figure 4, 3rd column). To isolate bending deformation we slowly push the flat cloth sample off the edge of a table and measure its shape as it bends under its own weight (Figure 4, 4th column), for both weft and warp directions. Thus we have a total of five measurements per cloth sample that will be used for parameter fitting (two stretch, one shear, and two bending).

We also capture two sequences with more complex deformation (Figure 5) for validation after parameter fitting. In the first test, opposite edges of the cloth are pulled in opposite directions, causing shearing and buckling (Figure 5, top). The second is a four-corner pulling test, where opposite pairs of corners are pulled in alternation, resulting in diagonal wrinkles (Figure 5, bottom).

Figures 4 and 5 show that our acquisition system is able to recover the 3D cloth geometry including temporal tracking (illustrated with an overlaid checkerboard), tracked 3D

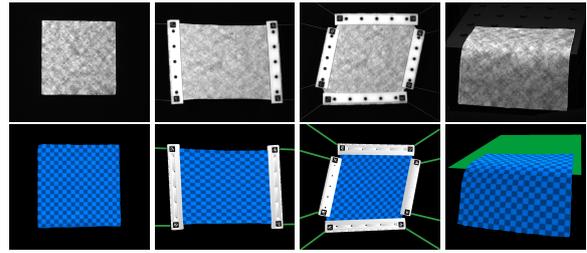


Figure 4: Selected frames from isolated measurements of stretching, shearing, and bending. The left column shows the cloth in its rest state. One input image is shown above each 3D reconstruction. The reconstruction includes parameterized cloth geometry, clip locations and the direction of the force vectors (shown as green lines).

clip locations, and individual 3D force directions (shown as green lines). To our knowledge, our method presents the first system able to record such extensive information about the behavior of a cloth sample.

3.3. Accuracy

In the vision system, the camera calibration accuracy is within 0.3 pixels, or about 0.075 millimeters at the distance of the cloth. The multi-view stereo algorithm of Bradley et al. [BBH08] is among the most accurate available according to the Middlebury evaluation benchmark. It is difficult to quantify the accuracy of the temporal flow computation, but it can be visualized by compositing the reconstructed deformation on top of the input images (see accompanying video).

The raw repeatability of our force sensors is about 3 millinewtons (RMS). The largest source of error in measuring the force indirectly through the cord is the internal friction in the cord as it bends around the pulleys, which introduces an artificial hysteresis of about 0.1 N.

4. Cloth Models

Our goal is to study the fidelity of constitutive models of cloth—models that predict the forces produced in the cloth

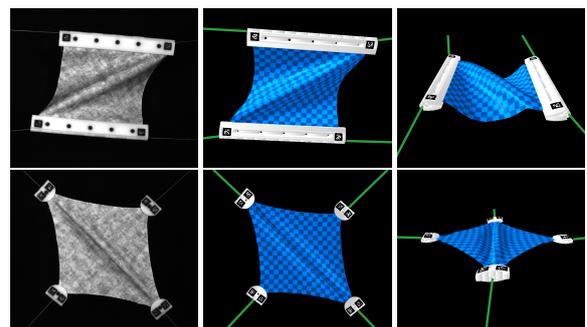


Figure 5: Selected frames from more elaborate manipulation of the cloth, demonstrating complex deformation behavior. The right column is rendered from a different viewpoint.

in response to deformations. The input of such a model is the positions of the vertices $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^3$ that define the deformation state of the sheet (analogous to strain in continuum mechanics) and the output is the forces that act between those vertices in response (analogous to stress). Although some of the models we look at are discrete in nature, we will use the convenient terms stress and strain to describe them.

4.1. Strain Metrics

Most elastic cloth models separate membrane (i.e., stretch and shear) and bending deformation energies. In both cases, deformation energy density can be described by the product of strain (ϵ) and stress (σ), i.e., $W = \frac{1}{2} \sigma \cdot \epsilon$. Furthermore, most of these models define separable scalar stress components as linear functions of individual scalar strain metrics. In that case, the energy density of each deformation component i can be written as $W_i = \frac{1}{2} k_i \epsilon_i^2$, where $k_i \epsilon_i = \sigma_i$ and k_i is the stiffness coefficient corresponding to the deformation component ϵ_i . The force density due to each ϵ_i follows as $\mathbf{F}_i = -\nabla W_i = -\sigma_i \nabla \epsilon_i = -k_i \epsilon_i \nabla \epsilon_i$. We have evaluated three models for membrane deformation that fit this description (spring systems, the soft constraint model by Baraff and Witkin [BW98] and the diagonalized St.Venant-Kirchhoff (StVK) model by Volino et al. [VMTF09]), and two bending models (spring systems and the edge-based bending model in Discrete Shells [GHDS03]).

Considering possible anisotropic behavior, we distinguish six different strain components on regularly triangulated cloth: weft-stretch ($\epsilon_{s,u}$), warp-stretch ($\epsilon_{s,v}$), shear ($\epsilon_{s,uv}$), weft-bend ($\epsilon_{b,u}$), warp-bend ($\epsilon_{b,v}$), and diagonal-bend ($\epsilon_{b,uv}$). Next, we describe in detail the strain metrics for the individual deformation components in the selected models. Note that not all force models define the quantities below explicitly as strains, as they often rely on the resolution of the discretization, or they differ simply by scale factors that can be embedded in the stiffness k_i . We use continuum strain definitions in all cases to fit them in a common formulation that allows us to easily compare the models. The details about the computation of strain gradients for the various cloth models, needed in the force computation, are given in their respective original papers.

Spring Membrane. All deformation components are modeled based on springs, with weft and warp ring-1 springs for stretch, and diagonal ring-1 springs for shear. The spring-based strain for each deformation component i can be defined as the relative change of edge length,

$$\epsilon_{s,i} = l/l_0 - 1, \quad (1)$$

with l the current length of the spring, and l_0 its rest length.

Diagonalized StVK. The membrane deformation is defined using the Green-Lagrange strain tensor, a formulation introduced to computer graphics by Terzopoulos et al. [TPBF87]. Given a per-triangle mapping function \mathbf{w} from the undeformed 2D configuration $(\mathbf{x}_{a,0}, \mathbf{x}_{b,0}, \mathbf{x}_{c,0})$ to the deformed

3D configuration $(\mathbf{x}_a, \mathbf{x}_b, \mathbf{x}_c)$, the deformation gradient can be computed as

$$(\mathbf{w}_u \ \mathbf{w}_v) = (\mathbf{x}_b - \mathbf{x}_a \ \mathbf{x}_c - \mathbf{x}_a) (\mathbf{x}_{b,0} - \mathbf{x}_{a,0} \ \mathbf{x}_{c,0} - \mathbf{x}_{a,0})^{-1}. \quad (2)$$

Then, the components of the membrane Green-Lagrange strain are defined as:

$$\epsilon_{s,u} = \frac{\|\mathbf{w}_u\|^2 - 1}{2}, \quad \epsilon_{s,v} = \frac{\|\mathbf{w}_v\|^2 - 1}{2}, \quad \epsilon_{s,uv} = \mathbf{w}_u^T \mathbf{w}_v. \quad (3)$$

Volino et al. [VMTF09] approximate the standard StVK model zeroing out off-diagonal terms in the matrix that relates strain and stress, $\sigma = \mathbf{E}\epsilon$. Then, in the diagonalized StVK, each membrane stress component depends only on its corresponding strain component, $\sigma_{s,i}(\epsilon_{s,i})$.

Soft Constraints. Weft- and warp-stretch are measured through a subtle modification of the Green-Lagrange strain tensor, defining terms that are quadratic in positions instead of quartic:

$$\epsilon_{s,u} = \|\mathbf{w}_u\| - 1, \quad \epsilon_{s,v} = \|\mathbf{w}_v\| - 1, \quad \epsilon_{s,uv} = \mathbf{w}_u^T \mathbf{w}_v. \quad (4)$$

Spring Bending. The deformation is measured based on weft and warp ring-2 springs for weft- and warp-bend, and diagonal ring-2 springs for diagonal-bend. Same as for membrane deformation, strain is measured as the relative change of edge length (1).

Discrete Shells Bending. It is measured as the change of dihedral angle at edges:

$$\epsilon_{b,i} = 1/h_0 (\theta - \theta_0), \quad (5)$$

with θ the current angle, and θ_0 the undeformed angle. Grinspun et al. [GHDS03] and Bridson et al. [BMF03] discovered concurrently the appropriate weighting of the angle change in order to model homogeneous bending on irregular triangle meshes with a homogeneous stiffness. Grinspun et al. define h_0 as a third of the average of the heights of the two triangles incident to the edge. This definition implies that bending energy density is integrated over edge-centered rectangles of size $l_0 \times h_0$. With our separation of weft-, warp- and diagonal-bending to capture anisotropy, the bending models in Discrete Shells and by Baraff and Witkin [BW98] are equivalent up to a stiffness scale factor.

4.2. Nonlinear Stress-Strain Curves

The generic force density model $\mathbf{F} = -\sigma \nabla \epsilon$ defined above assumes a linear stress-strain curve $\sigma = k\epsilon$. However, stress-strain curves are potentially nonlinear functions. Then, for each deformation component, we model stress as a function $\sigma_i = k_i(\epsilon_i)\epsilon_i$, with a strain-dependent stiffness k_i encoded using Hermite splines. We enforce non-negative constraints on the stiffness values at control points. The resulting nonlinear force density function, $\mathbf{F}_i = -k_i(\epsilon_i)\epsilon_i \nabla \epsilon_i$ yields a conservative force field, but note that the elastic energy density can

no longer be defined simply as $\frac{1}{2}k\epsilon^2$, and would now require the integration of the stiffness function.

Although only Volino et al. [VMTF09] propose a general nonlinear stress-strain relationship (though many systems use some form of strain limiting instead), the same construction can easily be built on any of our selected models. Because linear models fit the data poorly, we used the nonlinear model in all cases, resulting in a consistent set of models, parameterized by the number of spline control points, which reduces to the widely used linear models when each spline has a single control point.

5. Fitting

The key question of how well a given model describes a particular piece of cloth is answered by fitting the model to the measurement data: adjusting its parameters to minimize the difference between the model's predictions and the measured behavior, both in position and force. We do this by solving an optimization problem, leveraging that the cloth is at static equilibrium at the measured configurations.

In principle all parameters of a cloth model can be fit to a sufficiently rich single deformation sequence, but this can result in a problem fraught with local minima. In order to achieve stable fits, we have designed an incremental optimization procedure that fits model parameters a few at a time using the isolated deformations described in Section 3.2.

5.1. Optimization Problem

For each different cloth sample, we have created a simulated replica with the same mass, uniformly distributed, and the same 100mm square geometry, discretized with a regular 25×25 -node mesh, connected either with springs or with quadrilaterals split into triangles, depending on the model. In each measurement sequence, a different set of nodes is fixed to rigid bodies representing the clips. For the bending measurement sequences (see Figure 4), we fix all cloth nodes above the edge of the table. The measured pulling forces of the cords are applied as point forces on the rigid bodies at known locations, with known magnitudes and orientations.

Given a set of captured static deformation frames, we wish to know the (nonlinear) stress-strain curves for the deformation components of a cloth model, such that a simulated cloth matches known positions and forces as well as possible. Specifically, we minimize the weighted error of cloth positions and clip forces over a sequence of measurement frames, subject to the constraint of static equilibrium on all frames. For the formulation of the objective function, we concatenate in vectors the positions, \mathbf{x}_n , and net forces, \mathbf{F}_n , of free cloth nodes at all frames, as well as the forces, \mathbf{F}_c , applied by the cords on the clips. Due to equilibrium, the net force on the clips, produced by cord forces, gravity, and forces from fixed cloth nodes, must be zero. We indicate with $\tilde{\mathbf{x}}_n$ and $\tilde{\mathbf{F}}_c$, respectively, the known cloth node positions and clip forces, measured as described in Section 3.

We also concatenate in a vector \mathbf{k} the (unknown) stiffness values at the control points of the nonlinear stress-strain curves for the deformation components of the cloth. Since the pieces of cloth are homogeneous, we use a single curve for each deformation component for all frames and all cloth elements. Then, the computation of model parameters based on the minimization of position and force errors subject to the static equilibrium condition can be formulated as the following nonlinear constrained least-squares problem:

$$\begin{aligned} \mathbf{k} &= \arg \min_{\mu} \mu \|\mathbf{x}_n(\mathbf{k}) - \tilde{\mathbf{x}}_n\|^2 + \lambda \|\mathbf{F}_c(\mathbf{x}_n, \mathbf{k}) - \tilde{\mathbf{F}}_c\|^2, \\ \text{s.t. } \mathbf{F}_n(\mathbf{x}_n, \mathbf{k}) &= 0. \end{aligned} \quad (6)$$

In this optimization problem, we use the measured clip positions, $\tilde{\mathbf{x}}_c$, as known boundary conditions. For stretch tests, the objective function is based only on clip forces, i.e., $\mu = 0, \lambda = 1$, while for bend tests it is based only on cloth positions (since there are no measured forces), i.e., $\mu = 1, \lambda = 0$. For shear tests, the objective function is based only on clip forces parallel to the direction of the clips themselves. We observed that, in situations of near-homogeneous shear, the clip-parallel forces are dominated by shear, while clip-orthogonal forces are dominated by stretch. Then, by fitting only clip-parallel forces we reduce the sensitivity to potential errors in stretch stiffness.

The optimization problem contains two unknowns: the parameter vector \mathbf{k} and cloth node positions \mathbf{x}_n . We solve the optimization in an iterative manner, refining \mathbf{k} and \mathbf{x}_n separately on two nested loops. In an outer loop, we refine \mathbf{k} by local minimization of the error function and, in an inner loop, we recompute \mathbf{x}_n to satisfy the equilibrium constraint.

Outer Loop. Assuming cloth positions that satisfy the (nonlinear) equilibrium constraints on all captured frames, we locally linearize those constraints w.r.t. both \mathbf{k} and \mathbf{x}_n . As a result, we obtain a linear expression that relates node positions to parameter values:

$$\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n} \Delta \mathbf{x}_n + \frac{\partial \mathbf{F}_n}{\partial \mathbf{k}} \Delta \mathbf{k} = 0 \Rightarrow \Delta \mathbf{x}_n = -\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n}^{-1} \frac{\partial \mathbf{F}_n}{\partial \mathbf{k}} \Delta \mathbf{k}. \quad (7)$$

We also locally linearize clip forces,

$$\Delta \mathbf{F}_c = \frac{\partial \mathbf{F}_c}{\partial \mathbf{x}_n} \Delta \mathbf{x}_n + \frac{\partial \mathbf{F}_c}{\partial \mathbf{k}} \Delta \mathbf{k}, \quad (8)$$

and we turn (6) into a linear least squares problem, which we solve to refine the stiffness parameters:

$$\begin{aligned} \mathbf{k}(i+1) &= \arg \min_{\mu} \mu \|\mathbf{x}_n(i) - \tilde{\mathbf{x}}_n - \frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n}^{-1} \frac{\partial \mathbf{F}_n}{\partial \mathbf{k}} \Delta \mathbf{k}\|^2 + \\ &\lambda \|\mathbf{F}_c(i) - \tilde{\mathbf{F}}_c + \left(\frac{\partial \mathbf{F}_c}{\partial \mathbf{k}} - \frac{\partial \mathbf{F}_c}{\partial \mathbf{x}_n} \frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n}^{-1} \frac{\partial \mathbf{F}_n}{\partial \mathbf{k}} \right) \Delta \mathbf{k}\|^2. \end{aligned} \quad (9)$$

We terminate the outer loop (and hence the overall optimization) when the residual is reduced by less than 1% between two consecutive iterations. To ensure convergence of the Newton-like iterations and to enforce non-negativity

constraints on the components of \mathbf{k} , we execute a line search from $\mathbf{k}(i)$ to the solution of (9) if the residual grows or if the solution violates some constraint. The solution to the linear least squares problem requires solving a system $\mathbf{A}\mathbf{k} = \mathbf{b}$, where the size of \mathbf{A} is given by the number of unknown stiffness values, $|\mathbf{k}|$. In our test examples, this number was always below 10, and we solved the linear systems using LDL factorization. The formulation of \mathbf{A} , on the other hand, requires solving $|\mathbf{k}|$ linear systems of type $\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n} \mathbf{y} = \mathbf{b}$, which we did using the conjugate gradient method.

Inner Loop. Once the parameter values $\mathbf{k}(i+1)$ are refined, we bring the cloth to a static equilibrium position, $\mathbf{x}_n(i+1)$. We do this by solving quasi-static simulations until convergence on all captured frames, starting always from the measured configuration \mathbf{x}_n and using the measured clip positions $\tilde{\mathbf{x}}_c$ as boundary conditions. We consider that a piece of cloth has converged to equilibrium when $\|\mathbf{F}_n\| < 10\mu\text{N}$. The quasi-static simulations involve linear-system solves with the cloth stiffness matrix $\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n}$. We found that, during intermediate iterations, the stiffness matrix may not always be well conditioned, therefore we have solved the quasi-static equilibrium problems using additive Levenberg-Marquardt, which effectively produces a modified stiffness matrix of the form $\frac{\partial \mathbf{F}_n}{\partial \mathbf{x}_n} + \mu\mathbf{I}$. For improved conditioning, we also use this modified stiffness matrix in the outer loop.

5.2. Incremental Parameter Fitting

The nonlinearity of cloth deformation, together with the complex interplay of various deformation components in the resulting forces and positions, make the optimization problem above extremely complex in the general case, prone to falling in local minima and sensitive to initialization values. However, we largely alleviate these issues with the design of the five isolated deformation measurements described in Section 3.2, which allow us to separately fit stiffness curves for the six deformation components described in Section 4.1, following an incremental parameter fitting procedure.

First, we fit in parallel the weft-stretch stiffness curve, $k_{s,u}(\epsilon_{s,u})$, for the weft-stretch sequence, and the warp-stretch stiffness, $k_{s,v}(\epsilon_{s,v})$, for the warp-stretch sequence. We ignore shear and bend parameters for stretch fits, as we have observed that they have little effect. Second, using known stretch stiffness curves, we fit the shear stiffness $k_{s,uv}(\epsilon_{s,uv})$, for the shear sequence. Third, we fit in parallel the weft-bending stiffness $k_{b,u}(\epsilon_{b,u})$, for the weft-bending measurement sequence, and the warp-bending stiffness $k_{b,v}(\epsilon_{b,v})$, for the warp-bending sequence. Finally, we fit the diagonal-bending stiffness curve $k_{b,uv}(\epsilon_{b,uv})$, using both weft- and warp-bending measurements. To better account for cross-influence of shear and bending, we use their estimated values as initial guesses and run another fitting iteration.

To fit each stiffness curve $k_i(\epsilon_i)$, we iteratively subdivide the Hermite spline adding more control points until the residual error function (6) is reduced by less than 1% or a speci-

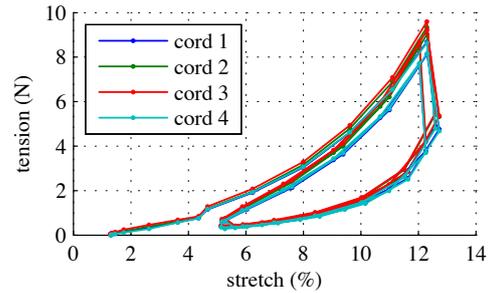


Figure 6: Raw force measurements for X stretching in Sample #18. Cloth is held by two bar clips, with two cords attached to each clip.

fied maximum number of points, usually 4 or 5, is reached. First, we evaluate the strain histogram for the corresponding measurement sequence, and we determine maximum and minimum strains after removing outliers. We initialize the stiffness curve with one control point (i.e., constant stiffness), and subsequently we subdivide the strain range with equidistant control points.

6. Results

We tested our system on four fabric samples, including a knit and the three common weave patterns (plain weave, twill, and satin), and three fiber types (cotton, wool, and synthetic): cotton satin (#4), rayon/spandex knit (#12), cotton denim (#14), and wool/cotton blend (#18). Each fabric was tested with seven deformations (see Section 3.2): for fitting, stretch in X and Y, simple shear, and bending in X and Y; and for evaluation, complex shearing and corner pulling. Figure 6 shows the four individual tension measurements for a typical stretching test, plotted as a function of the overall extension of the cloth. The measurement shows the typical behavior of a woven fabric: a nonlinear curve with increasing stiffness for higher strain, and large hysteresis. The test repeats three times, retracing the same loop each time after the initial extension from rest.

We worked with three cloth models built from the components described in Section 4. The Springs model uses the spring membrane model with the spring bending model; the Soft Constraints model uses Baraff and Witkin's membrane model with the Discrete Shells bending model; and the St. VK model uses the diagonalized St. Venant-Kirchoff membrane model with the Discrete Shells bending model. We fit all the models in four variants: linear (constant stiffness for each deformation mode), isotropic (identical stiffness in warp and weft), linear and isotropic (the simplest variant), and nonlinear orthotropic (the most general variant). The results are too numerous to include in the paper; we refer the reader to the supplementary material, which illustrates the behavior of the nonlinear orthotropic variant of all three models for all four fabrics, and the behavior of the variants

of the Soft Constraints model for denim, a largely nonlinear and anisotropic material.

Observations. In Figure 7 and Table 2 we present results for fitting the Soft Constraints model to Sample #12. For each test we show a selected frame (near maximum distortion) with renderings illustrating the captured and fitted cloth geometry and forces. To illustrate the fitting residuals more quantitatively, we show a force-displacement plot comparing a summary of the measured forces to the predictions of the fitted model and a vector-field plot illustrating the position error over the geometry of the fitted mesh (see caption for details). Similar tables are shown for all cloth samples and several model variants in the supplementary document.

The four selected fabrics span a large range of possible cloth behaviors. In a nutshell, #12 is isotropic and very compliant in stretch and bending; #4 is also isotropic, very stiff in stretch but compliant in bending; #14 is stiff and quite isotropic in stretch, but extremely anisotropic in bending (with 33/1 stiffness ratio in weft and warp); and #18 is anisotropic both in stretch (with 10/1 stiffness ratio) and in bending (with 13/1 stiffness ratio). The maximum stretch stiffness for #4 is 250 times higher than for #12, while #14 is 10 times stiffer in shear than any other fabric. All four fabrics show similar hysteresis behavior, with loading-to-unloading stretch stiffness ratios ranging from 1.4/1 to 1.8/1. Sample #12 is nearly linear in the test deformation range, while all other three fabrics exhibit nonlinearity. Interestingly, nonlinearity may arise in some deformation modes but not in others, with no clear pattern.

For stretching, all three cloth models fit nicely to the average of the hysteresis bands, even in highly nonlinear cases. The fitting residual is larger for stiffer fabrics, and the non-linear orthotropic model variants fit anisotropic fabrics best, as expected, while linear and/or isotropic variants reach a reasonable compromise but are not always able to remain inside the hysteresis band. For shearing, the fitting force residual is larger for #14, the stiffest fabric. Across models, the Soft Constraints and St. VK models fit to the average of the shearing hysteresis band, while the Springs model deviates at times. For bending, no forces are available, and we evaluate the position residual as well as profiles of sample curves orthogonal to the support plane. The fitting residual is similar for all fabrics, but distinctly higher for the Springs model. Often, the residual is dominated by a difference in curl near the edge of the sample, while the overall shape is well fit.

The last two rows of Figure 7 show the evaluation tests; these data were not used in fitting the models. The behavior of sample #12, the most linear fabric, is predicted well in all cases, as seen in the force-displacement plots, the buckling behavior in corner pulling, and the (lower) effective shear stiffness of the sheet when allowed to buckle in the complex shear test. In the three other samples, however, the force-displacement plot of the corner pulling test reveals that the stiffness is underestimated for the Soft Constraints and

Scenario	Stretch-X	Stretch-Y	Shear	Bend-X	Bend-Y
Residual	0.53N	0.62N	0.22N	21.3mm	22.0mm

Table 2: RMS residuals of non-linear orthotropic fitting to Sample #12 with the Soft Constraints model.

St. VK models, and overestimated for the Springs model. Visually, the mismatch is more apparent in the complex shear test, where models with underestimated stiffness exhibit wider folds than the real fabrics.

We have also evaluated the fitted models on new test samples of each fabric, to validate their generality. Specifically, we have tested stretching on new samples of rayon/spandex knit (#12.2) and cotton denim (#14.2), and shearing on new samples of cotton satin (#4.2) and wool/cotton blend (#18.2). The force-displacement plots of the real cloth samples, shown in the supplementary document, indicate very similar behavior between fitting and test samples for #12 and #14, and a larger disparity for #4 and #18. The evaluation plots for the simulation models behave similar for the test and fitting cases, but the matching quality depends of course on the actual disparity across cloth samples.

Discussion. While overall force-displacement behavior is nicely matched, the actual folding shapes of simulated cloth may deviate largely from the captured cloth, because even a small change in material properties may lead to distant stable configurations in the L^2 sense. For this reason, the traditional L^2 metric is not appropriate for evaluating error in this case. The discontinuity of stable configurations is also the cause of flickering and twitches in some of our examples.

The Springs model exhibits the worst fitting quality in shearing force-displacement curves, and the highest fitting residual for bending. This is probably due to the inherent coupling of stretch and bending deformation components in this model. Nevertheless, the overall deformations in complex shearing fit reasonably well. In contrast to continuum models, complex parameter tuning has often been regarded as a caveat of mass-spring models; but our results indicate that satisfactory parameter estimation is possible by incorporating anisotropy and nonlinearity into the model. The Soft Constraints and St. VK models produce results with very similar quality, which is expected as the models present only subtle differences as described in Section 4.1.

At least three effects are missed by the tested models: hysteresis, Poisson effect (due to the diagonalization of the standard StVK model), and cross-modal stiffening (e.g., shear stiffening due to stretching). We indeed identified stretch stiffening in the shearing deformations, therefore we chose clip-parallel forces as objective function to minimize the effect of stretch errors on shear optimization. We conjecture that missing cross-modal stiffening may also be, to a large extent, the reason for stiffness underestimation in the corner pulling test for the Soft Constraints and St. VK models. An

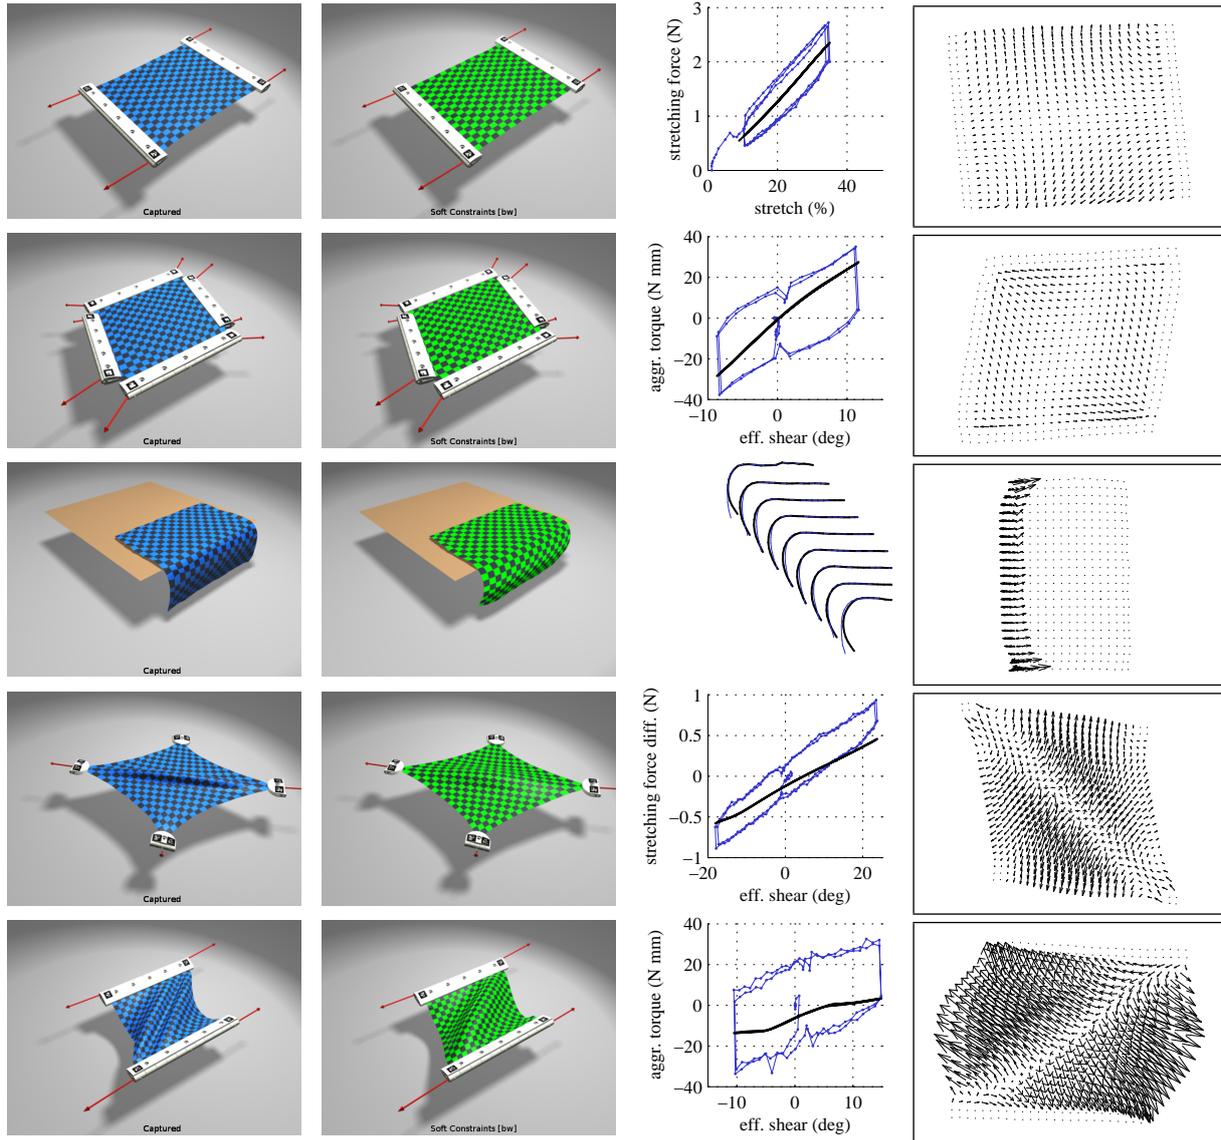


Figure 7: Fitting results for Sample #12 with the Soft Constraints model. Top to bottom; Stretch-X, Simple shear, Bend-X, Corner pull, Complex shear. Left to right: captured geometry, equilibrium of fitted model, force comparison (thin line: measurement; thick line: model), position residual (vertex position minus corresponding measured position, magnified 5x). “Effective shear” is the shear angle of the best-fit transformation to the motion of the clamped cloth vertices. “Aggregated torque” is the torque about the center of the cloth applied by the clamps, with every other clamp counted negatively. “Stretching force difference” is the difference in the two corner-to-corner stretching forces.

extension to the nonlinear model of Wang et al. [WRO11] could help alleviate these problems.

7. Conclusion

This paper has demonstrated a novel system for observing cloth behavior, including complete information about deformation and forces, and a new method for fitting and evaluating cloth models using the measurements. Our system is different from standard textile testing systems because it

captures detailed geometry information; it is different from previous cloth capture systems in that it captures complete force information and measures deformations of a 3D surface. The combination of very complete position *and* force information provides an unprecedented view into the complex behavior of cloth.

Our measurement setup offers very accurate control over membrane deformations, but the bending tests require manual intervention and are thus less precise. Furthermore, the

bending tests are most accurate for samples with straight edges, but some cloth materials (in particular knit) tend to curl up at free boundaries. In order to eliminate these problems, we would like to investigate alternative ways of controlling bending deformations in the future.

The data from our experiments shows some of the limitations of current models. The most obvious of these is hysteresis—all widely used cloth models are elastic, but cloth is clearly far from elastic, resulting in quite large errors for any given point in the experiment. There are many paths for future work in measurement, including more complete exploration of strain space (including compression) and capture of dynamic properties, and in fitting, where new ways of evaluating fitting error are needed that can work when the cloth's equilibrium state is unstable or non-deterministic.

Acknowledgments. This work was funded in part by the Spanish Ministry of Science and Innovation, project TIN2009-07942, and by the European Research Council, project ERC-2011-StG-280135 Animetrics.

References

- [AIH*08] ATCHESON B., IHRKE I., HEIDRICH W., TEVS A., BRADLEY D., MAGNOR M., SEIDEL H.-P.: Time-resolved 3d capture of non-stationary gas flows. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 27, 5 (2008), 132. 3
- [BBH08] BRADLEY D., BOUBEKEUR T., HEIDRICH W.: Accurate multi-view reconstruction using robust binocular stereo and surface meshing. In *Proc. CVPR* (2008). 3, 4
- [BFA02] BRIDSON R., FEDKIW R., ANDERSON J.: Robust treatment of collisions, contact and friction for cloth animation. In *Proceedings of ACM SIGGRAPH '02* (2002), pp. 594–603. 2
- [BHPS10] BRADLEY D., HEIDRICH W., POPA T., SHEFFER A.: High resolution passive facial performance capture. *ACM Trans. Graphics (Proc. SIGGRAPH)* (2010), 41. 3
- [BHW94] BREEN D., HOUSE D., WOZNY M.: Predicting the drape of woven cloth using interacting particles. In *Proceedings of ACM SIGGRAPH '94* (1994), pp. 365–372. 2
- [BMF03] BRIDSON R., MARINO S., FEDKIW R.: Simulation of clothing with folds and wrinkles. In *Proc. ACM SIGGRAPH/Eurographics SCA* (2003), pp. 28–36. 5
- [BPS*08] BRADLEY D., POPA T., SHEFFER A., HEIDRICH W., BOUBEKEUR T.: Markerless garment capture. *ACM Trans. Graphics (Proceedings of SIGGRAPH)* 27, 3 (2008), 99. 2, 3
- [BTH*03] BHAT K. S., TWIGG C. D., HODGINS J. K., KHOSLA P. K., POPOVIĆ Z., SEITZ S. M.: Estimating cloth simulation parameters from video. In *Proc. ACM SIGGRAPH/Eurographics SCA* (2003), pp. 37–51. 2
- [BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. In *Proc. SIGGRAPH* (1998), pp. 43–54. 2, 5
- [CK02] CHOI K.-J., KO H.-S.: Stable but responsive cloth. In *Proceedings of ACM SIGGRAPH '02* (2002), pp. 604–611. 2
- [CPGE90] CLAPP T. G., PENG H., GHOSH T. K., EISCHEN J. W.: Indirect measurement of the moment-curvature relationship for fabrics. *Textile Research J.* 60, 9 (1990). 2
- [Cul79] CULPIN M. F.: The shearing of fabrics: A novel approach. *J. Textile Institute* 70, 3 (1979), 81–88. 2, 4
- [EB08] ENGLISH E., BRIDSON R.: Animating developable surfaces using nonconforming elements. In *Proceedings of ACM SIGGRAPH '08* (2008). 2
- [EKS03] ETZMUSS O., KECKEISEN M., STRASSER W.: A Fast Finite Element Solution for Cloth Modelling. In *Proc. Pacific Graphics* (2003), pp. 244–251. 2
- [EWS96] EBERHARDT B., WEBER A., STRASSER W.: A fast, flexible, particle-system model for cloth draping. *IEEE Computer Graphics and Applications* 16, 5 (1996), 52–59. 2
- [Fia05] FIALA M.: Artag, a fiducial marker system using digital techniques. In *IEEE CVPR* (2005), pp. 590–596. 3
- [GGWZ07] GARG A., GRINSPUN E., WARDETZKY M., ZORIN D.: Cubic shells. In *Proceedings of ACM SIGGRAPH/Eurographics SCA* (2007), pp. 91–98. 2
- [GHDS03] GRINSPUN E., HIRANI A., DESBRUN M., SCHRÖDER P.: Discrete shells. In *Proceedings of ACM SIGGRAPH/Eurographics SCA* (2003), pp. 62–67. 5
- [GHF*07] GOLDENTHAL R., HARMON D., FATTAL R., BERCOVIER M., GRINSPUN E.: Efficient simulation of inextensible cloth. In *Proc. SIGGRAPH* (2007), pp. 281–290. 2
- [Kaw80] KAWABATA S.: The standardization and analysis of hand evaluation. *Textile Machinery Soc. Japan* (1980). 2, 4
- [KJM08] KALDOR J. M., JAMES D. L., MARSCHNER S.: Simulating knitted cloth at the yarn level. In *Proceedings of ACM SIGGRAPH '08* (2008). 2
- [KNM10] KUNITOMO S., NAKAMURA S., MORISHIMA S.: Optimization of cloth simulation parameters by considering static and dynamic features. In *ACM SIGGRAPH Posters* (2010), p. 15:1. 2
- [Mül08] MÜLLER M.: Hierarchical position based dynamics. In *Proc. of VRIPHYS* (2008), pp. 1–10. 2
- [PKST08] PABST S., KRZYWINSKI S., SCHENK A., THOMASZEWSKI B.: Seams and bending in cloth simulation. In *Proc. of VRIPHYS* (2008). 2
- [Pro95] PROVOT X.: Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Graphics Interface '95* (1995), pp. 147–154. 2
- [SGdA*10] STOLL C., GALL J., DE AGUIAR E., THRUN S., THEOBALT C.: Video-based reconstruction of animatable human characters. *ACM Trans. Graph.* 29, 5 (2010). 2
- [TPBF87] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. In *Proceedings of ACM SIGGRAPH '87* (1987), pp. 205–214. 2, 5
- [TPS09] THOMASZEWSKI B., PABST S., STRASSER W.: Continuum-based strain limiting. *Computer Graphics Forum (Proceedings of Eurographics '09)* 28 (2009), 569–576. 2
- [VMTF09] VOLINO P., MAGNENAT-THALMANN N., FAURE F.: A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Trans. Graph.* 28, 4 (2009). 2, 5, 6
- [WCF07] WHITE R., CRANE K., FORSYTH D. A.: Capturing and animating occluded cloth. *ACM Trans. Graph.* 26, 3 (July 2007). 2
- [WOR10] WANG H., O'BRIEN J., RAMAMOORTHY R.: Multi-resolution isotropic strain limiting. In *Proc. SIGGRAPH Asia* (2010), ACM. 2
- [WRO11] WANG H., RAMAMOORTHY R., O'BRIEN J.: Data-driven elastic models for cloth: Modeling and measurement. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 30, 4 (2011), 71. 2, 9