

Data-Driven Elastic Models for Cloth: Modeling and Measurement

Huamin Wang

James F. O'Brien

Ravi Ramamoorthi

University of California, Berkeley



Figure 1: When worn by the same mannequin model, shirts made of different cloth materials exhibit distinctive patterns of wrinkles and folds in our simulation. For example, the Gray Interlock shirt has many small wrinkles since it is compliant in stretching and bending, while the shirt made of the stiffer Pink Ribbon Brown material tends to form a few larger wrinkles.

Images copyright Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi.

Abstract

Cloth often has complicated nonlinear, anisotropic elastic behavior due to its woven pattern and fiber properties. However, most current cloth simulation techniques simply use linear and isotropic elastic models with manually selected stiffness parameters. Such simple simulations do not allow differentiating the behavior of distinct cloth materials such as silk or denim, and they cannot model most materials with fidelity to their real-world counterparts. In this paper, we present a data-driven technique to more realistically animate cloth. We propose a piecewise linear elastic model that is a good approximation to nonlinear, anisotropic stretching and bending behaviors of various materials. We develop new measurement techniques for studying the elastic deformations for both stretching and bending in real cloth samples. Our setup is easy and inexpensive to construct, and the parameters of our model can be fit to observed data with a well-posed optimization procedure. We have measured a database of ten different cloth materials, each of which exhibits distinctive elastic behaviors. These measurements can be used in most cloth simulation systems to create natural and realistic clothing wrinkles and shapes, for a range of different materials.

Keywords: Nonlinear elasticity, anisotropy, data-driven model, cloth simulation, parameter estimation.

Links: [DL](#) [PDF](#) [VIDEO](#) [WEBSITE AND DATA](#)

Contact email: {whmin, job, ravir}@eecs.berkeley.edu

From the conference proceedings of ACM SIGGRAPH 2011.
Appearing in ACM Transaction on Graphics Vol. 30, No. 4.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM SIGGRAPH, 2011, Vancouver, BC Canada
© Copyright ACM 2011

1 Introduction

Most real-world cloth materials exhibit nonlinear and anisotropic behavior due to their woven nature and fibrous composition. These properties distinguish different cloth materials by creating distinctive appearances when they drape, fold or wrinkle. For example, as shown in Figure 1, shirts composed of different materials will appear markedly different from each other even if they have the same cut. Unfortunately, most cloth simulation techniques in graphics ignore these properties for simplicity, and formulate cloth stiffness with linear isotropic models whose parameters are often manually selected. While using such a model simplifies the problem and generates physically plausible results, it is difficult to distinguish different cloth materials and many interesting wrinkling and folding effects cannot be accurately generated. A natural solution to this problem is to construct elastic models from real-world cloth data. Unfortunately, little research has been done in this direction even though data-driven approaches have been widely adopted in other areas of computer graphics.

There are two main approaches to capture real-world elastic behaviors. In materials science and textile engineering, it is common to design a device that isolates each material parameter and measures it directly. Since a large number of parameters are needed to describe the behavior of real cloth, designing such a device is complicated and typically results in large and expensive machines. A further complication arises because cross-terms may cause one parameter to depend on phenomena controlled by another. Some prior work in graphics has instead tried to estimate cloth material parameters from unconstrained motion in images or videos. While the uncontrolled nature of these experiments is appealing, there is a large parameter space, which is difficult to optimize for while avoiding local minima. Robustly tracking features from unconstrained cloth motion is another challenging problem. Feature tracking algorithms often suffer from noise and occlusion, which further affects the optimization result.

Our measurement methodology seeks to find a good balance between these two approaches. We build simple devices that deform samples in a controlled way so that their shapes can be easily measured. However, we do not require the cloth sample to be uniformly

deformed as with other measurement devices, nor do we require that measurements isolate individual material parameters. Instead, we reduce the parameter space by specifying an elastic model, and finding its stiffness parameters by formulating an optimization that minimizes the difference between the simulation and observed data. Because the experiment occurs in controlled configurations where features can be easily extracted, the whole parameter optimization process is well-defined. We find optimal parameters that closely match the actual displacement-force relationship of each material.

Based on this concept, we present a data-driven technique to construct elastic models of cloth. Our specific contributions are:

- **Modeling:** We propose a simple piecewise linear model that is a good approximation to nonlinear anisotropic deformations, and can be thought of as extending Hooke's law to match a range of real cloth behaviors. This leads to 24 stretching and 15 bending parameters per sample, that can be fit by optimization.
- **Measurement:** We develop novel measurement techniques for both stretching and bending behaviors of cloth samples. Our experimental setup is relatively simple and inexpensive to construct. (See Figures 7 and 11.) We validate the correctness of our measurements with additional experiments, and our cloth hanging examples demonstrate good qualitative agreement with real-world examples. (See Figure 14.)
- **Database:** We measure complete stretching and bending parameters for 10 examples ranging from compliant to stiff, and from anisotropic to isotropic. These measurements allow us to build a database that includes both raw measurements and fit parameters.

The acquired cloth models can be used with existing continuum-based cloth simulators to generate distinctive and natural elastic behaviors when pieces drape, wrinkle and fold. We show simulations of clothing for a virtual character in Figure 1, with clear differences apparent between different cloth materials.

2 Previous Work

Modeling: Accurate modeling of fabric elastic behavior remains an open problem. Significant efforts [Alsawaf, 1985; Zhang et al., 2005; King et al., 2005; Yeomana et al., 2010] have been devoted to finding numerical solutions to specific structures, but these approaches are insufficient for developing a general model for the wide range of real cloth materials. Our elastic model is closely related to generic and data-driven models in textile engineering. Previous work includes a two-step linearization method [Kawabata et al., 1982], a bi-linear model [Ng et al., 2005], a data-driven elastic model based on energy relaxation [McCartney et al., 2005], and a data-driven bending model [Zhou and Ghosh, 1998]. While these models closely approximate the displacement-force relationship, we found that a simple piecewise linear model demonstrates distinctive elastic behaviors, fits well in a parameter optimization framework, and can be easily included in existing cloth simulators.

Measurement: Devices that directly capture stiffness parameters from real cloth samples are used in materials science and textile engineering. These include biaxial tensile testers developed by KATO Tech [2011b], Deben [2011], Zwick [2011], and the KES-FB2 bending tester by KATO Tech [2011a]. These machines must be carefully designed to isolate the parameter to be measured. Avoiding interactions among different elastic behaviors is also required but difficult. For example, the sample should be homogeneously deformed to guarantee the accuracy of the measurement, which is challenging when sample edges are also deformed.

Volino and his collaborators [2009] proposed an elastic model based on one-dimensional strain-stress curves that are directly measured from unidirectional tensile tests. While their model approximates well elastic behaviors of stiff cloth material, they do not consider the relationship among stiffness parameters in different directions. Therefore, the model is not sufficient to handle compliant materials with obvious Poisson effects. Poisson effects can also affect the accuracy of unidirectional tensile tests.

Alternatively, other researchers [Bhat et al., 2003; Kunitomo et al., 2010] have presented automatic techniques to learn stiffness parameters from unconstrained cloth motions. While their experiments do not need specific devices and can be easily set up, the unconstrained nature of the motions makes it difficult to define meaningful features that can be tracked and used in the optimization process.

Bickel and collaborators [2009] proposed a system to capture and model nonlinear heterogeneous soft tissue in three dimensions. Although they only use two parameters for each data point under the locally isotropic assumption, they assign a data point to each tetrahedron in every example measurement, which greatly increases the number of parameters for fitting. In our case, we assume that the material is locally anisotropic but globally homogeneous, so we can sample data points over the strain space in a structured way. Further, our parameter fitting problem is well constrained and does not require an additional regularization term.

Directly capturing and reconstructing cloth animation from videos is also a problem that has been addressed in graphics, including marker-based techniques by White and his collaborators [2007] and markerless approaches by Bradley [2008]. Our work is also related to the cloth optimization framework proposed by Wojtan and his colleagues [2006]. While their work is focused on optimizing cloth motion for a mass-spring system, we are interested in finding stiffness parameters for elastic models in continuum-based cloth simulation.

Simulation: Physically based cloth simulation has been an active topic in graphics for more than twenty years. Much work has focused on improving simulation and collision techniques. (For example, [Baraff and Witkin, 1998; Bridson et al., 2002; Bridson et al., 2003].) In contrast, our focus is on accurately modeling nonlinear and anisotropic cloth behaviors. Strain limiting techniques approximate the nonlinearity of cloth in a bi-phasic way [Provot, 1995; Bridson et al., 2003; Goldenthal et al., 2007; Müller, 2008; English and Bridson, 2008; Thomaszewski et al., 2009; Wang et al., 2010]. Volino and his colleagues [2009] pushed this direction even further by using a hyperelastic St. Venant-Kirchhoff model. Breen and his colleagues [1994a; 1994b] modeled different cloth materials by approximating one-dimensional Kawabata testing results using a particle based system. Choi and Ko [2002] proposed a nonlinear bending model for simulating wrinkles and folds. Grinspun and collaborators [2003] presented a thin shell model to simulate bending behaviors of stiff shells. Kaldor and colleagues [2008; 2010] demonstrated that nonlinear anisotropic behaviors of knitted cloth can be simulated at the yarn level. All of these methods can benefit significantly from data-driven elastic measurements of real cloth, as described in this paper.

3 Piecewise Linear Elastic Model

The constitutive law in an elastic model defines the relationship between the strain/displacement and the stress/force. Ideally the elastic model should approximate well the relationship observed in real materials, including the nonlinear and anisotropic behaviors characteristic of real cloth materials. In addition, the model should be easy to incorporate into existing cloth simulators, and able to fit




Figure 2: Polar space spanned by λ_{\max} and φ sampled by six data points.

accurately to real data with a simple optimization procedure. We found a piecewise linear model well suited to these requirements.

3.1 Planar Stretching Model

Assuming that the scale of threads and their interweaving patterns are significantly smaller than elastic behaviors we are interested in, we treat cloth samples as a two-dimensional continuum. The planar displacement can be described by a strain tensor and the planar force can be described by a stress tensor. A planar elastic model defines the relationship between these two tensors.

We start with the simple linear anisotropic model obtained by generalizing Hooke's law [Slaughter, 2002]:

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_{uu} \\ \sigma_{vv} \\ \sigma_{uv} \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} & c_{13} \\ c_{12} & c_{22} & c_{23} \\ c_{13} & c_{23} & c_{33} \end{bmatrix} \begin{bmatrix} \varepsilon_{uu} \\ \varepsilon_{vv} \\ \varepsilon_{uv} \end{bmatrix} = \mathbf{C} \boldsymbol{\varepsilon} \quad (1)$$

in which $\boldsymbol{\sigma}$ and $\boldsymbol{\varepsilon}$ are the stress and strain tensor, and \mathbf{C} is a 3×3 symmetric stiffness tensor matrix with six parameters. From experimental observations, Boisse and his colleagues [2001] pointed out that woven composite fabrics are orthotropic and when the local coordinate system is aligned with the warp-weft directions, \mathbf{C} can be simplified into a matrix of four parameters:

$$\mathbf{C} = \begin{bmatrix} c_{11} & c_{12} & 0 \\ c_{12} & c_{22} & 0 \\ 0 & 0 & c_{33} \end{bmatrix}. \quad (2)$$

Intuitively, it means the shearing resistance is separable from the stretching resistance. We found that the orthotropic assumption appears to be valid for most cloth materials we tested.

Peng and Cao [2005] provided a formula to transform \mathbf{C} into the corresponding stiffness matrix if a non-orthogonal local coordinate system is not aligned with the warp-weft coordinate system. In our case the local coordinate system is always orthogonal and the difference between those two coordinate systems is only caused by a bias angle, so we can simply define the resting cloth mesh in the warp-weft coordinate system while the deformed mesh is in the local coordinate system. This allows us to avoid extra computational cost when applying the elastic model to a biased cloth piece.

Instead of treating \mathbf{C} as a constant matrix, the piecewise linear elastic model formulates \mathbf{C} as a piecewise linear function $\mathbf{C}(\boldsymbol{\varepsilon})$ of the strain tensor $\boldsymbol{\varepsilon}$. Since the original Green-Lagrangian strain tensor in the Voigt form $\boldsymbol{\varepsilon} = [\varepsilon_{uu}, \varepsilon_{vv}, \varepsilon_{uv}]^T$ is coordinate dependent and its values are not intuitive enough to demonstrate the actual deformation, we re-parameterize $\boldsymbol{\varepsilon}$ into principal strains λ_{\max} , λ_{\min} and a strain angle φ using Eigenvalue Decomposition:

$$2 \begin{bmatrix} \varepsilon_{uu} & \varepsilon_{uv} \\ \varepsilon_{uv} & \varepsilon_{vv} \end{bmatrix} = \mathbf{R}_\varphi^T \begin{bmatrix} (\lambda_{\max} + 1)^2 - 1 & 0 \\ 0 & (\lambda_{\min} + 1)^2 - 1 \end{bmatrix} \mathbf{R}_\varphi \quad (3)$$




Figure 3: The four stiffness parameters c_{11} , c_{12} , c_{22} , c_{33} for the Gray Interlock material shown in the polar space.




Figure 4: The four stiffness parameters c_{11} , c_{12} , c_{22} , c_{33} for the Ivory Rib Knit material shown in the polar space.

in which \mathbf{R}_φ is a rotational matrix defined by strain angle φ . We noticed from our experiments that λ_{\min} has significantly less influence on \mathbf{C} , so we simplify the parametrization $\mathbf{C}(\boldsymbol{\varepsilon})$ by only using two variables: λ_{\max} and φ . Figure 2 shows a standard scheme that uses six data points to sample the polar space spanned by λ_{\max} and φ . Each data point contains four parameters, c_{11} , c_{12} , c_{22} and c_{33} as used in Equation 2, and the full stretching model has 24 parameters. $\mathbf{C}(\lambda_{\max}, \varphi)$ is then defined by linearly interpolating data points over λ_{\max} and φ , respectively. The interpolated result is a stiffness tensor matrix that maps a strain tensor to a stress tensor. Figures 3 and 4 show all four variables of the stiffness tensor matrix in a polar space for the Gray Interlock material and the Ivory Rib Knit material respectively. Each variable demonstrates the nonlinear elastic properties using the piecewise linear model and it is clear that those two materials have different elastic properties. The difference between c_{11} and c_{22} also shows that both materials are anisotropic.

Exceptions: Some materials are not structured symmetrically to their warp and weft directions. Further, for some materials these directions are not easily recognizable from the cloth sample, and some materials may not even have these directions defined in their structure (such as tri-axially woven fabrics or nonwoven fabrics). A commonly known example is denim that is asymmetrically woven




Figure 5: The bending angle and its neighboring triangles.

in a twill structure. Fortunately, our experiments show that an elastic model under the orthotropic assumption still can approximate the strain-stress relationship of denim well. Although we did not find such cases in our experiments, for cloth materials that greatly conflict with this assumption, we can use all 6 parameters of the stiffness tensor \mathbf{C} in equation 1 (instead of the four parameter approximation in equation 2), and use more data points to sample the strain angle φ from 0° to 180° in the polar space.

3.2 Bending Model

Our piecewise linear bending model is extended from the bending force equation by Bridson and his colleagues [2003]:

$$F_i = k \sin(\theta/2) (h_1 + h_2)^{-1} |E| u_i \quad (4)$$

for $i=1, \dots, 4$, in which E is the edge vector, F_i is the bending force applied on the i -th vertex, h_1 and h_2 are the heights of two triangles, u_i is a vector defined for the i -th vertex, and k is the mesh-independent bending stiffness coefficient. Figure 5 shows the configuration of the bending angle and its neighborhood. As with stretching, we consider a piecewise-linear curvature-force relationship, and simply define k as a piecewise linear function $k(\alpha)$ using a mesh-independent variable α :

$$\alpha = \sin(\theta/2)(h_1 + h_2)^{-1} \quad (5)$$

The variable α approximates half of the curvature in a simple fashion, using $\sin(\theta/2)$ and $(h_1 + h_2)^{-1}$, since they are already calculated for Equation 5 and do not introduce extra computational cost. So far the model in Equations 4 and 5 is isotropic since it does not consider any bias angle. To make the model anisotropic, we sample the angle space, construct the piecewise linear elastic model separately for each orientation, and define the complete model by linear interpolation in the polar space spanned by the angle and α .

Our experience shows that using five data points to sample α can approximate the bending behaviors well without overfitting. In total we use 15 data points for three samples of each cloth material. Figure 6 plots the estimated bending stiffness curves for a Pink Ribbon Brown material. The parameters are fit using the optimization framework that will be discussed later in Section 4. This graph shows that the bending stiffness relationship is different when the bias angle changes.

4 Measurement

The planar model proposed in Section 3.1 requires 24 parameters and the bending model proposed in Section 3.2 contains 15 parameters. With this number of material parameters it is impractical to manually adjust them in order to match the model to a specific cloth material. Instead, we propose a measurement and optimization framework that allows us to model stiffness parameters automatically from captured cloth behaviors. In this section,




Figure 6: Bending stiffness curves of a Pink Ribbon Brown cloth material.

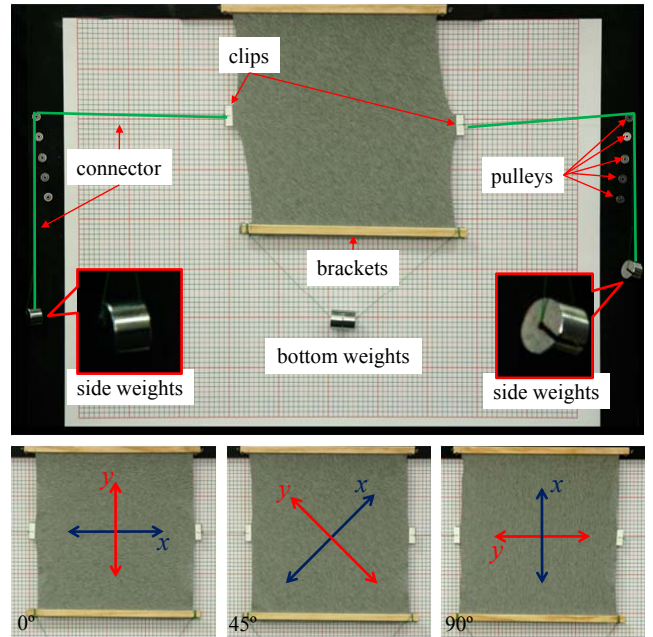


Figure 7: The experiment setup for stretching tests. For each material, a set of three cloth samples with different bias angles are tested in the experiment.

Images copyright Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi.

we describe our measurement setup to determine those parameters. We conduct a separate set of measurements for stretching (4.1) and bending (4.2). Compared with existing data-driven techniques, our system is easy and inexpensive to build and configure. It does not completely isolate each parameter as would be required for direct measurement, but it generates test configurations that are controlled enough to enable the stiffness parameters to be fit within a simple optimization framework using well-defined features in the observed cloth samples.

4.1 Stretching Measurement

Our stretching measurements are made using a novel experimental setup, with parameters fit using optimization. The objective of this setup is not to isolate individual material properties. Instead it creates controlled and measurable configurations that can be used as targets for optimization.

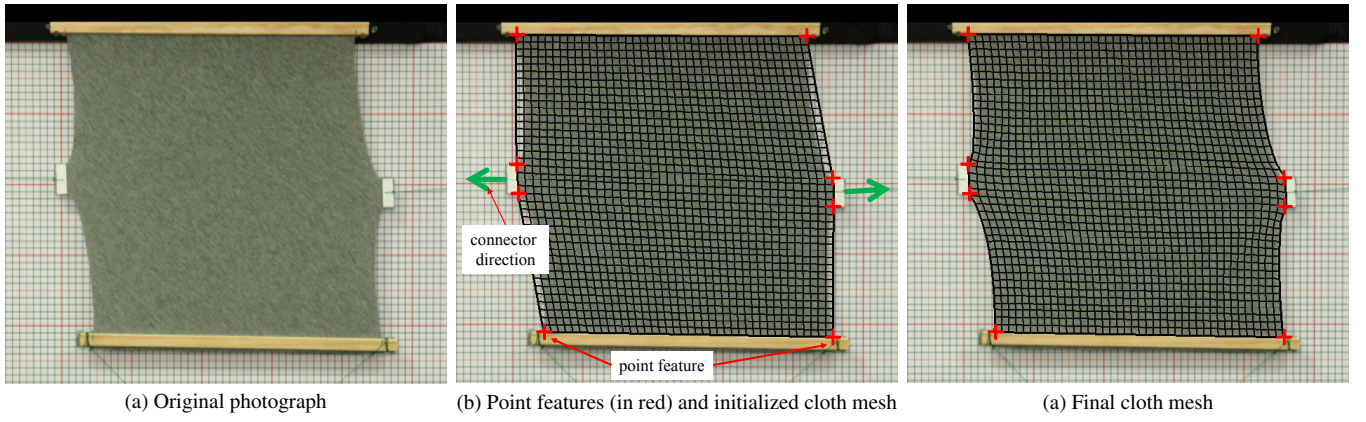


Figure 8: A 45° cloth sample shown in (a) is tested by hanging weights on its three edges. We use labeled point features (in red) to construct an initial cloth mesh in (b). Labeled connector directions (in green) are used as force directions in the simulator, so that the simulated cloth mesh finally matches with the observation as shown in (c). In (b) and (c) a grid of parameter lines from the simulated cloth has been superimposed on an image of the real cloth to allow comparison.

Images copyright Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi.

4.1.1 Experiment Setup

To fit the stretching parameters, we design experiments that demonstrate a sufficient set of cloth behaviors, and enable easy comparison of real and simulated data during optimization. Our method is inspired by the biaxial tensile method in the textile literature, which tests the cloth sample by stretching it simultaneously in both warp and weft directions. We configure a stretching tester as shown in Figure 7. For cloth materials with symmetric properties to their warp and weft directions, we create three 400mm×400mm cloth samples with bias angles 0°, 45° and 90° respectively. The bias angle is defined as the rotational angle from the warp-weft coordinate system to the sample's local coordinate system counterclockwise. Warp and weft directions can be recognized from thread directions in the weaving structure for most cloth materials.

To constrain cloth motion in a controlled way, the top and bottom edges are each sandwiched between a pair of wooden slats, and the left and right edges are attached to white cardboard rectangles in the middle of each edge. These locations are treated as boundary conditions and their positions can be easily measured using a calibrated camera. During each test, the top edge of the cloth sample is attached to the top of the testing board, while the other three edges still have freedom to move. Different weights are then applied on these three edges in order to drag the cloth sample into different shapes. The left and right sides are loaded with the same weights so that the sample does not lose its balance during the experiment. Each sample is typically tested by seven different weights at the bottom, going from 0g to 600g, and five weights on both sides, from 0g to 400g. In total, there are 35 tests for each sample and 105 tests for each cloth material. This test set covers the range of forces typically experienced by the cloth in clothing when it is worn. We use a calibrated DSLR camera to capture the cloth shape in each test. The camera is mounted approximately four meters away from the board with a long-focal-length lens to minimize perspective effects.

4.1.2 Parameter Optimization

Given our elastic model and a set of captured images, we formulate stiffness parameters as the solution to a minimization problem.

Problem Formulation: Let f_i^* be shape features captured from the i -th test, and $f_i(p_0, p_1, \dots, p_n)$ be corresponding features generated by cloth simulation using the given planar elastic model with 24 parameters p_0, p_1, \dots, p_n , as discussed in Section 3. Our goal is to find optimal parameters so that the difference between captured

features and simulated features can be minimized:

$$\{p_0, p_1, \dots, p_n\} = \arg \min_{\{p_0, p_1, \dots, p_n\}} \sum_{i=1}^T w_i \|f_i^* - f_i(p_0, p_1, \dots, p_n)\| \quad (6)$$

in which T is the number of tests. In order to prevent this sum from being dominated by greatly deformed shapes, we introduce a fall-off factor w_i to decrease their influence,

$$w_i = \min(\|f_{rest} - f_i\|^{-1}, 10^6) \quad (7)$$

in which f_{rest} are shape features of a resting cloth sample.

One important question here is: how are shape features defined? A direct way is to extract the silhouette of the cloth sample as a shape feature, since the sample only deforms in a plane. In practice, this is prone to having errors when sample edges are not cut straight and they can even be curly for certain cloth materials. So we manually label image locations of wood slats and white cardboard clips, and treat them as point features as shown in Figure 8b. Besides being used in the error metric as described in Equation 6, features can also be used to construct an initial cloth mesh for simulation using bilinear interpolation, also shown in Figure 8b. In particular, we treat the initial cloth shape with no loads as the resting cloth mesh. Each cloth mesh is represented by a triangle mesh over a 41×41 grid.

Our continuum-based cloth simulator uses the standard Finite Element Method (FEM) described by O'Brien and Hodgins [1999]. The simulator is conditioned in the same way as the test. For example, the top edge of the cloth mesh is always fixed, and its bottom edge is only allowed to move in a rigid way. To simplify the simulation, we fix feature orientations by only allowing features to translate. Once different forces are applied in labeled connector directions (shown with green arrows in Figure 8b), the cloth mesh will reach different equilibrium shapes in the cloth simulator. Simulated features are then used in Equation 6 to compute the error metric.

Optimization: Once both f_i and f_i^* are defined in Equation 6, we would like to optimize parameters p_0, p_1, \dots, p_n so that the error metric can be minimized. Like other optimization systems, this system also suffers from the local minima problem. Unlike other problems, however, we observe in our experiments that the local minima are often clustered in a small local region. Mathematically, this implies that the gradient vector provides a good clue to the convergence when the error is still large. So we use the BFGS extension of the Quasi-Newton method to handle the optimization

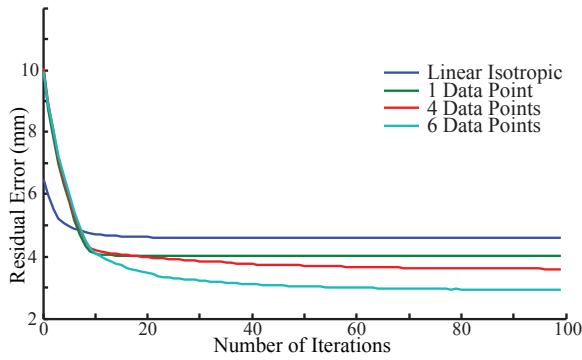


Figure 9: The convergence plot for modeling stretching behaviors, with data shown for Tango Red Jet Set. Using more data points, meaning more piecewise linear regions, helps the system converge to a lower minimum.




Figure 10: Residual errors of different schemes after 100 iterations for material Gray Interlock. The leftmost scheme uses a linear isotropic model with two parameters, so we treat it as a half-data-point scheme. While using more data points reduces the residual error during the optimization process, the validation shows that the overall fitting quality may become worse, because of the overfitting issue. This plot also shows that our model performs better in both optimization and validation than a model without considering any Poisson effect, which was used by Volino and his colleagues [2009].

process in the first 10 to 20 iterations. Once the residual error becomes smaller, we add random perturbations in order to jump away from local minima. We also restart Hessian matrix prediction in the Quasi-Newton method every few iterations, which causes the method to behave more like a gradient descent method. The reason we restart Hessian matrix prediction from time to time is because the system uses finite differencing to estimate the gradient vector, which is more likely to have errors when the system moves closer to local minima.

Although we cannot guarantee that the proposed optimization scheme will always converge to the global minimum, a local minimum can still be considered as an acceptable result if the resulting behavior is a good match to the real cloth. From our experiments, the relative difference between a local minimum and the expected global minimum is usually between 5% to 10%. We believe the local minimum clustering property is mainly caused by the cloth memory property, which potentially introduces ambiguity in the underlying elastic model and is not considered by our system yet.

Figure 9 shows a typical convergence plot, for the cloth material Tango Red Jet Set. The Quasi-Newton method quickly reduces the error in the first 10 to 20 iterations. Once the convergence is slowed

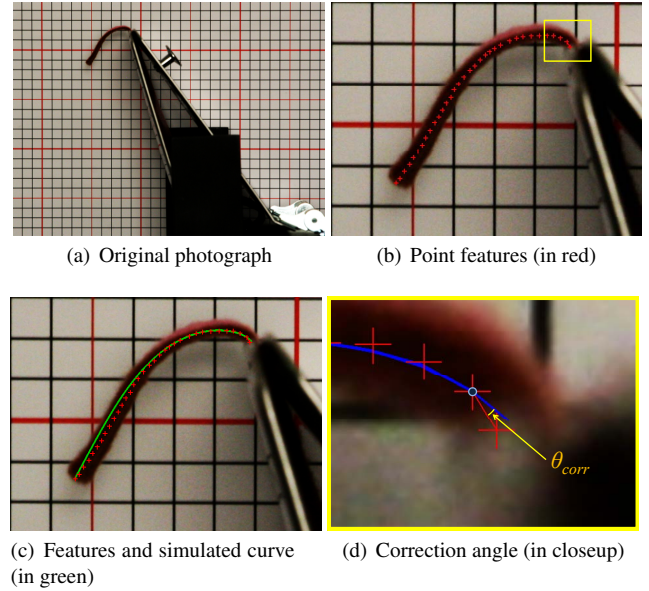


Figure 11: The experiment setup for bending tests (a). By using point features shown in (b), the optimization system finds a set of optimal parameters that matches the simulated mesh with the observed cloth shape (c). A correction angle θ_{corr} shown in (d) is used as an intrinsic parameter to account for labeling errors.

Images copyright Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi.

down after that, random perturbation helps the system avoid local minima so the error can be gradually reduced in the next 40 to 80 iterations even further. The remaining error may be caused either by the cloth memory property, or cutting and labeling errors throughout the measurement process.

Figure 9 also compares error fitting curves of four different sampling schemes of the piecewise-linear elastic model sampled along the pattern shown in Figure 2. The first scheme (in blue) is the common linear isotropic model that uses two Lamé parameters. The second scheme (in green) that uses a single data point is essentially the linear model defined by the generalization of Hooke's law. The third scheme (in red) uses a data point at the origin and three data points at 0° , 45° and 90° respectively. The plot clearly indicates that using more data points (more piecewise linear regions) will produce a better fitting result. While continuing to add more data points will further reduce the error, we find it is more likely to overfit the data as Figure 10 shows. In this plot, schemes using 10 data points and 11 data points are built by adding more data points in the angle dimension and the radius dimension respectively. Because the error will become dominated by other issues rather than the model itself, we use 6 data points as a good tradeoff between accuracy and overfitting.

Intrinsic Parameters: Our experience shows that cutting and labeling errors are often inevitable, partially due to the memory property and edge-curling property of cloth materials. These errors cause inaccuracy in measurement and prevent the optimization system from further minimizing the error metric. So we introduce intrinsic parameters to account partially for these errors in our system. The first set of intrinsic parameters is used to provide a better rest configuration for the cloth mesh. The cloth sample may not be cut in a perfect $400\text{mm} \times 400\text{mm}$ shape so we use the initially reconstructed cloth mesh with no loads as the resting configuration. Although this mesh is close to the actual shape, it is still slightly deformed by gravity. Assuming that this candidate rest mesh is

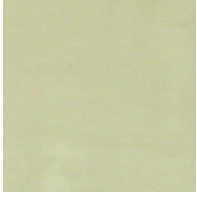

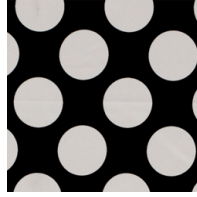





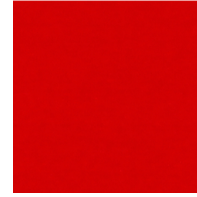
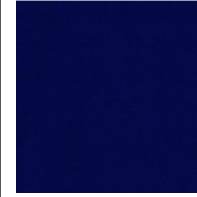
					
Color & Name	Ivory Rib Knit	Pink Ribbon Brown	White Dots on Black	Navy Sparkle Sweat	Camel Ponte Roma
Composition	95% Cotton 5% Spandex	100% Polyester	100% Polyester	96% Polyester 4% Spandex	60% Polyester 40% Rayon
Common Usage	Underwear	Blanket	Tablecloth	Sweater	Jacket
Density (kg/m ²)	0.276	0.228	0.128	0.224	0.284
Error (mm)	1.92	1.66	2.42	3.24	1.67
					
Color & Name	Gray Interlock	11oz Black Denim	White Swim Solid	Tango Red Jet Set	Royal target
Composition	60% Cotton 40% Polyester	99% Cotton 1% Spandex	87% Nylon 13% Spandex	100% Polyester	65% Cotton 35% Polyester
Common Usage	T-shirt	Jeans	Swimsuit	Fashion dress	Pants
Density (kg/m ²)	0.187	0.324	0.204	0.113	0.220
Error (mm)	2.58	2.30	1.57	2.06	0.89

Table 1: An elastic model database of ten different cloth materials. Each of them exhibits distinctive stretching and bending behaviors. The document in the supplemental material contains full sets of parameters.

only vertically stretched from the actual resting shape, we introduce a vertical scaling parameter for each sample to improve the mesh shape. The second set of intrinsic parameters are correction angles that are added to bias angles, because hand-cut samples may not have the intended bias angle. This is especially useful for the 45° sample, which is likely to have larger bias angle errors. In total, there are six intrinsic parameters for all three samples. They are optimized together with stiffness parameters using the optimization system.

4.2 Bending Measurement

We now describe the method for determining bending parameters. The optimization is largely similar to that for stretching, and we provide additional details specific to bending.

4.2.1 Experiment Setup

Figure 11a shows the setup we used for testing bending behaviors of cloth. The most important part of this setup is a pair of metal plates that clamp the cloth sample at the top of the apparatus. Because bending behaviors are also anisotropic, we create three 4cm-wide cloth samples with bias angles 0°, 45° and 90°, respectively. As the sample is incrementally advanced so that progressively more of it protrudes from the clamp, the sample drapes into different curves under its own weight. We capture these curves from a side view and manually label the trajectory of each curve using point features as shown in Figure 11b. Compared with other bending measurement devices that use sensors to measure curvatures and bending moments, our device is simple to build and easy to use.

4.2.2 Parameter Optimization

Unlike the planar elastic model, the bending elastic model is defined separately for each bias angle. Therefore, the whole optimization problem can be solved separately for each sample. We use point features to formulate the bending error metric as discussed in Section 4.1.2. Its solution gives a set of bending stiffness parameters for each tested cloth material as before. Because the simulation is sensitive to the initial segment of the curve and the labeling process may have errors, we adjust the initial segment of each curve by introducing a correction angle as an intrinsic parameter for each test, as shown in Figure 6d. We then use the optimization algorithm described in Section 4.1.2 to find optimal stiffness parameters. Once we obtain stiffness parameters for all three bias angles, the whole bending elastic model is defined by linear interpolation in the polar space spanned by the bias angle and α .

5 Database

We have used the measurement apparatus described above to build a database of the elastic behavior of ten different cloth materials, shown in Table 1. Each cloth material typically needs 24 hours for parameter estimation, including two hours to cut cloth samples, one hour for data capture, two hours to label features, and 19 hours for optimizing material parameters. For optimization we used an Intel-based machine with four X7560 2.27GHz Xeon octo-core CPUs.

The ten materials in our database are made of different fabric compositions and different woven or knitted patterns. They exhibit distinct elastic behaviors, which are representative of a variety of cloth materials. For example, Poisson effects are minimal with the Ivory Rib Knit material, likely due to its vertically ridged structure. It is stretchy in the horizontal direction and is commonly used to make

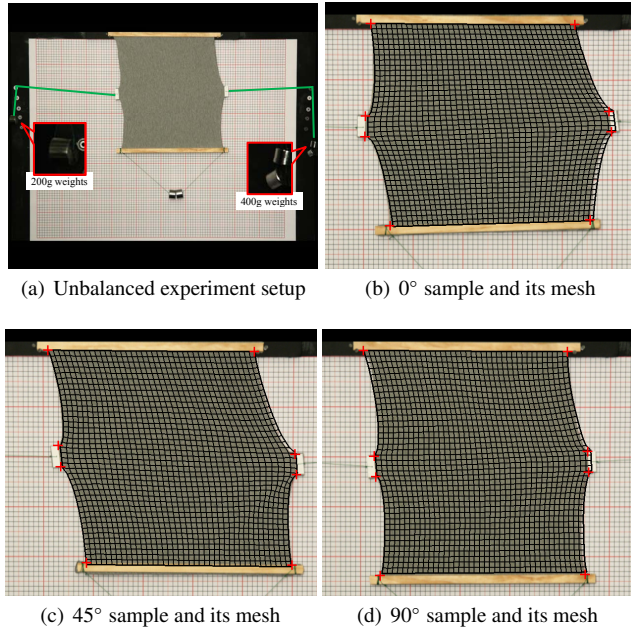


Figure 12: *Unbalanced biaxial tensile tests for validation purposes. Simulated cloth meshes are rendered as a grid over the image to show the difference between their shapes and captured cloth sample shapes.*

Images copyright Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi.

tight fitting T-shirts and undergarments. In contrast, the Gray Interlock material has more obvious anisotropic behaviors and it is more compliant when it bends, so it is usually used to make common T-shirts. Some materials have stiffness in stretching, such as the Pink Ribbon Brown material commonly used for making blankets, and the 11oz Black Denim material often used for making jeans, shown in Figures 1b and d respectively. Some materials exhibit behaviors that are not well captured by our model, such as edge-curling in the Navy Sparkle Sweat material.

Each residual error reported in Table 1 is the weighted mean distance from simulated features to labeled point features for stretching measurement. The error value typically falls between 1mm and 3mm. Compared with the scale of the 400mm×400mm cloth samples we used for the experiment, these errors are sufficiently small that they indicate a good match to the real data (see also the validations in the next section). Compliant materials, such as Navy Sparkle Sweat and Gray Interlock, tend to have larger errors because their cloth samples are more deformed during the experiment.

While a low numerical error is desirable, it does not always correlate to a perceptual notion of similarity, especially for new configurations. For example, in Figure 14, a simulated cloth sheet made of Gray Interlock looks (in the authors' opinion) more similar to its corresponding photograph than does the White Swim Solid, even though the second material has a smaller residual error.

We believe this database¹ to be a useful resource for the graphics community. The results can be easily included in existing cloth simulators to produce realistic behaviors of many different types of materials in clothing animations, and validate further research on measurement and modeling.

¹The database, which includes both the raw image measurements, the fit parameters, and image textures can be downloaded at: <http://graphics.berkeley.edu/papers/Wang-DDE-2011-08>.

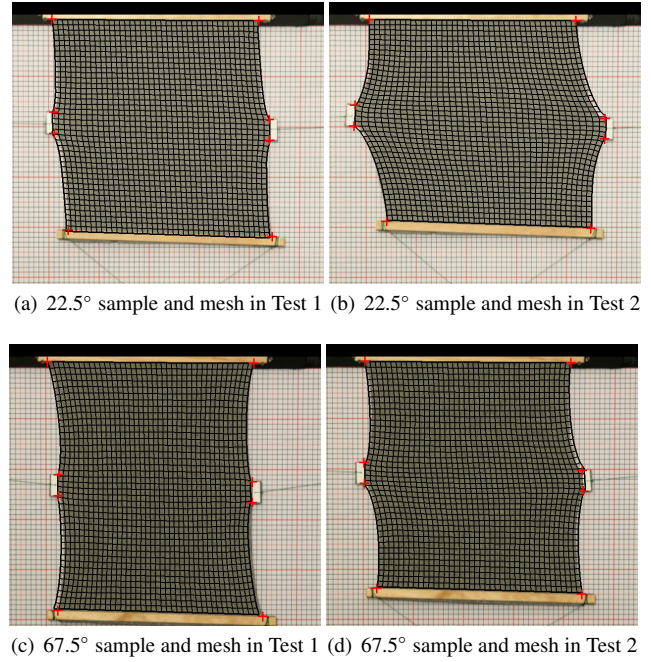


Figure 13: *Off-angle tensile tests for validation purposes. The first row shows two examples of the 22.5° sample and the second row shows two examples of the 67.5° sample.*

Images copyright Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi.

6 Results

The parameters from our model can be directly used in a cloth simulator and used to generate clothing animations. Animated versions of figures 1, 14 and 15 appear in the supplementary video.

6.1 Validation

Optimizing the error metric described in Equation 6 generates a set of stiffness parameters that can minimize the difference between image features and simulated features in a test set. However, it is still not clear whether the result matches with the actual elastic properties, because the selected features only sparsely represent a subset of cloth shapes in the whole deformation space. So we use extra tests for validation that are not included in the optimization.

Validation I - Unbalanced Tensile Tests: We use unbalanced stretching tests for 0°, 45°, and 90° samples of the same cloth material, by putting a 200g weight on the left side and a 400g weight on the right side as shown in Figure 12a. Unbalanced tests are not used in parameter optimization. Figures 12b-d show cloth shapes from all three unbalanced validation tests, and simulated cloth meshes are rendered as a dark overlay grid for comparison. The weighted mean feature differences are 4.32mm, 3.17mm, and 4.90mm, for samples with bias angles 0°, 45° and 90° respectively. This experiment indicates that the model closely approximates the strain-stress relationship of these cloth samples even in unbalanced cases. We believe the main reason causing these errors to be slightly larger than the original optimization error, is because the cloth memory property will behave differently in a single unbalanced test than in a sequence of balanced tests.

Validation II - Different Bias Angles: In order to test the performance of the elastic model with other bias angles, we cut two extra cloth samples of the Gray Interlock material, with 22.5° and

67.5° bias angles respectively. They were tested by balanced bi-axial tensile tests in the same way as other cloth samples but they are not used to optimize stiffness parameters. Because we do not know intrinsic parameters for these two samples, we run an extra optimization procedure solely to obtain their intrinsic parameters. We can then use the cloth simulator to find the difference between the simulated mesh and the captured cloth shape. Some examples of these two samples are shown in Figure 13. The weighted mean feature differences of the two samples are 4.50mm and 4.39mm respectively. Interestingly, we noticed from the intrinsic parameter that the actual bias angle of the nominal 22.5° cloth sample should be 28.5° instead. We verified this observation by inspecting the actual cloth texture. This example demonstrates that a large error in creating the sample is sometimes unavoidable.

Validation III - Hanging Cloth: Figure 14 shows a hanging example that compares simulated cloth meshes with real photographs. (Please also refer to the accompanying video for animations of the models.) Each cloth piece is assumed to be 1m×1m and wrinkles are formed by separating two holding points 84cm apart. The difference between 0° and 90° results demonstrates that Gray Interlock, Tango Red Jet Set and Navy Sparkle Sweat have more obvious anisotropic behaviors, while the other materials are more isotropic. We also compare simulation results of using the linear isotropic model with our model in Figure 14a and 14b. Although the differences appear subtle, they can substantially affect the perception of a material and they show that the linear isotropic model is not sufficient to realistically represent real cloth elastic behaviors.

Because we do not capture cloth BRDFs nor lighting conditions (we simply approximated the overall diffuse color and texture for rendering our examples), cloth visual appearance is somewhat different from the captured images. Cloth silhouettes are also slightly different, because the rendering view does not perfectly match with the actual camera and some cloth materials tend to have curly edges. While simulated hanging cloth matches well with the real photograph for most cloth materials, we also notice that highly compliant materials like Navy Sparkle Sweat and White Swim Solid match somewhat less well. We think this is due to the cloth memory and edge-curling effects.

6.2 Animations

We include animations (please see the supplementary video) of our models being used in cloth simulators. The simulator runs on an Intel i7 930 2.80GHz CPU. The simulation time step ranges from 0.1 milliseconds to 0.07 milliseconds, and each rendered frame typically takes 20 to 30 seconds to generate. While stiffness parameters are chosen from the database to generate elastic behaviors of different cloth materials, we manually select damping and friction parameters to produce desired dynamic effects.

Cloth Draping: Figure 15 shows cloth sheets made of different materials draping over spheres underneath. (Please also refer to the accompanying video for whole animations and more materials.) It can be clearly seen that Royal Target has distinctively larger wrinkles compared with other materials. Meanwhile, the first three materials also demonstrate different wrinkle patterns, even though they have some similarities. For example, White Swim Solid and Camel Ponte Roma are more stretchy and the wrinkles cover a larger region than Ivory Rib Knit. White Swim Solid is especially likely to have small and sharp wrinkles, because it is compliant in bending.

Walking Mannequin: Figure 1 shows a wood mannequin model wearing shirts made of different cloth materials. This example demonstrates distinctive wrinkles and folds caused by different cloth materials, even when all shirts follow the same garment pattern. As an example, Pink Ribbon Brown tends to form a few large

wrinkles as shown in Figure 1b. For this reason, this material is typically used for blankets rather than shirts. Meanwhile, Navy Sparkle Sweat shown in Figure 1c is compliant in both stretching and bending behaviors; therefore wrinkles and folds are likely to be smoothed out. Figure 1a shows a shirt made of Gray Interlock, which is commonly used to make T-shirts. Although 11oz Black Denim (Figure 1d) and White Dots on Black (Figure 1e) are both stiff in stretching, the latter one is more compliant in bending and more likely to form small wrinkles.

7 Limitations

Although our experiments show that the proposed optimization framework successfully models elastic behaviors of various cloth materials, it also has several limitations that need to be addressed in the future. First, the piecewise linear elastic model only provides an approximation to the actual strain-stress relationship when the cloth is slightly deformed. If larger forces are applied to the cloth, this approximation will become less accurate. Adding more data points into the model would help reduce this problem, but doing so would also introduce more computational cost and more local minima so the result quality is not guaranteed to be the same as before. Second, our current model does not consider the cloth memory property. This form of plasticity can be a severe problem for certain cloth materials, because it causes ambiguities in the underlying elastic model and confuses the optimization process by introducing more local minima. Finally, cutting and labeling errors inevitably affect the performance of our system, especially for cloth materials with edge-curl effects. Using more cloth samples and more tests can help reduce this problem, but that also requires more computational cost and human labor. Finally, we have only looked at measuring static parameters. Dynamic parameters, such as internal damping, should also be measured.

8 Conclusion and Future Work

We present a comprehensive study on data-driven elastic models for cloth, including measurement of a database from real-world cloth samples using inexpensive devices, an optimization-based modeling approach that finds optimal parameters under a piecewise linear model, and various simulations that validate the model and demonstrate its ability to distinguish cloth materials. Elastic models in the database can be easily incorporated into existing continuum-based cloth simulators to generate distinctive wrinkles, folds and drapes, which are often difficult to model by manually setting parameters in linear isotropic models.

Looking into the future, we plan to solve existing limitations in our current system as described in Section 7. We are interested in incorporating more cloth materials into the database, constructing data-driven friction and damping models for the dynamic motion of cloth, and even extending the optimization framework to handle solid elastic models for three-dimensional deformable bodies.

In summary, data-driven models have found wide acceptance in computer graphics, in areas such as motion capture for animation, and measured BRDFs for reflectance. In this work, we have explored the new domain of data-driven elastic models for cloth. Because realistic clothing is a key element of virtual characters, we believe our database, measurement setup, and piecewise linear elastic model will find broad applicability.

Acknowledgments

We thank Jingchao Li for helping us prepare cloth samples and capture measurement data. We also thank the members of the Berke-

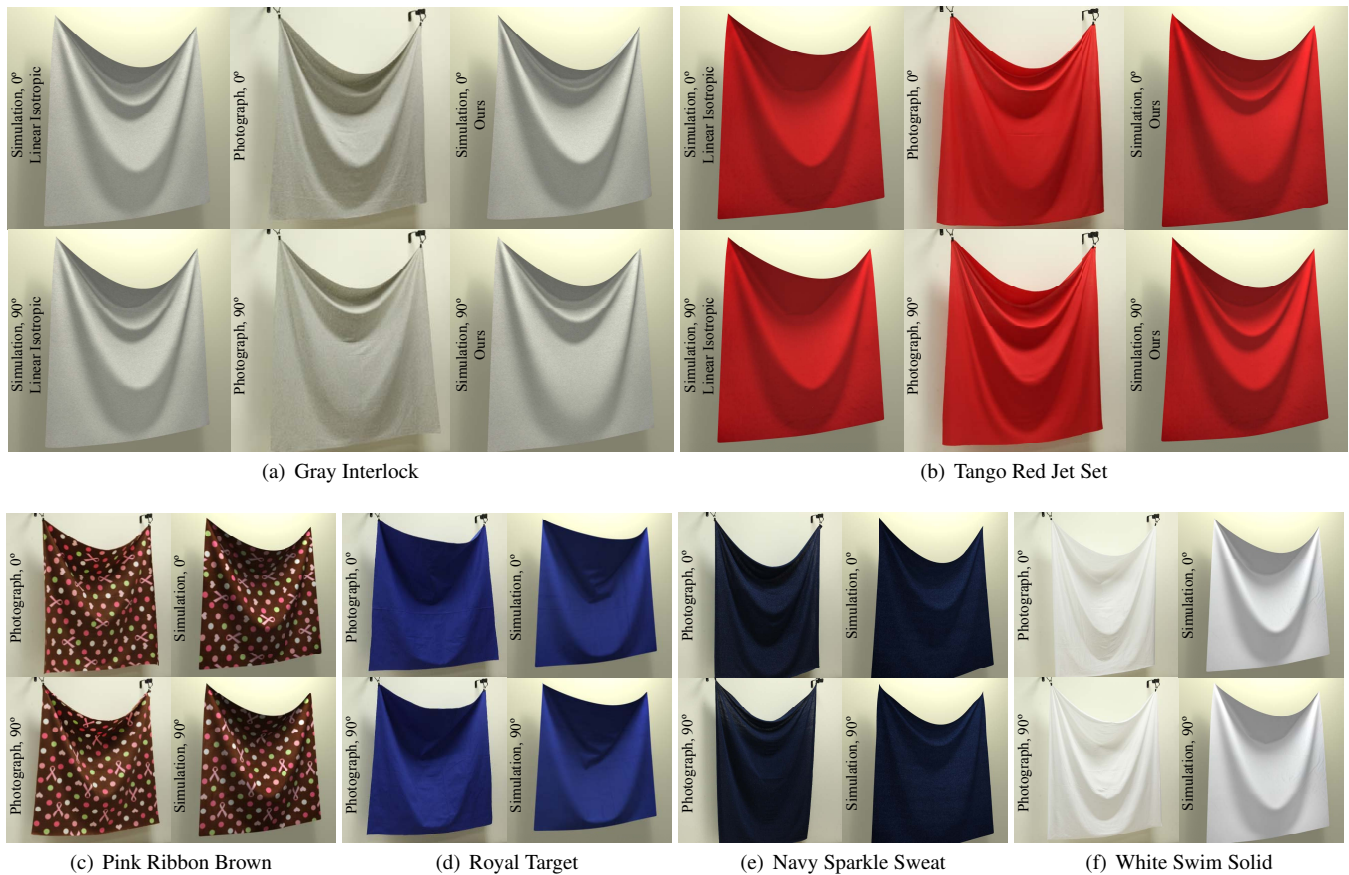


Figure 14: A comparison of real and simulated hanging sheets. By choosing various cloth materials, a $1m \times 1m$ simulated cloth sheet is deformed into different shapes under gravity and compared to photographs of real sheets.

Images copyright Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi.

ley Graphics Group for their support and helpful suggestions. This work was supported in part by NSF grants HCC-1011832 and IIS-0915462, UC Lab Fees Research grant 09-LR-01-118889-OBRIJ, and the Intel Science and Technology Center for Visual Computing. We also thank Intel, NVIDIA, Adobe, and Pixar for additional support through equipment and funding.

References

- ALSAWAF, F. 1985. *A model for textile tensile curves*. PhD thesis, University of Manchester Institute of Science and Technology.
- BARAFF, D., AND WITKIN, A. 1998. Large steps in cloth simulation. In *Proc. of SIGGRAPH 98*, E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 43–54.
- BHAT, K. S., TWIGG, C. D., HODGINS, J. K., KHOSLA, P. K., POPOVIĆ, Z., AND SEITZ, S. M. 2003. Estimating cloth simulation parameters from video. In *Proc. of the ACM SIGGRAPH/Eurographics symposium on Computer animation*, 37–51.
- BICKEL, B., BÄCHER, M., OTADUY, M. A., MATUSIK, W., PFISTER, H., AND GROSS, M. 2009. Capture and modeling of non-linear heterogeneous soft tissue. *ACM Transactions on Graphics (SIGGRAPH 2009)* 28, 3 (July), 89:1–89:9.
- BOISSE, P., GASSER, A., AND HIVET, G. 2001. Analyses of fabric tensile behavior: determination of the biaxial tension-strain surfaces and their use in forming simulations. *Composites: Part A*, 32, 1395–1414.
- BRADLEY, D., POPA, T., SHEFFER, A., HEIDRICH, W., AND BOUBEKEUR, T. 2008. Markerless garment capture. *ACM Transactions on Graphics (SIGGRAPH 2008)* 27, 3 (August), 99:1–99:9.
- BREEN, D. E., HOUSE, D. H., AND WOZNY, M. J. 1994. A particle-based model for simulating the draping behavior of woven cloth. *Textile Research Journal* 64, 11, 663–685.
- BREEN, D. E., HOUSE, D. H., AND WOZNY, M. J. 1994. Predicting the drape of woven cloth using interacting particles. In *Proc. of SIGGRAPH 94*, E. Fiume, Ed., Computer Graphics Proceedings, Annual Conference Series, 365–372.
- BRIDSON, R., FEDKIW, R., AND ANDERSON, J. 2002. Robust treatment of collisions, contact and friction for cloth animation. In *Proc. of ACM SIGGRAPH 2002*, E. Fiume, Ed., vol. 21 of *Computer Graphics Proceedings, Annual Conference Series*, 594–603.
- BRIDSON, R., MARINO, S., AND FEDKIW, R. 2003. Simulation of clothing with folds and wrinkles. In *Proc. of the ACM SIGGRAPH/Eurographics symposium on Computer animation*, 28–36.
- CHOI, K.-J., AND KO, H.-S. 2002. Stable but responsive cloth. *ACM Transactions on Graphics (SIGGRAPH 2002)* 21, 3 (July), 604–611.
- DEBEN, 2011. Biaxial Tensile Stage for Textiles and Polymers, Jan. <http://www.deben.co.uk/details.php?id=11>.




Figure 15: Cloth draping over a sphere. The first three materials shown in (a)-(c) are all compliant so they exhibit similar wrinkle patterns (specific differences are discussed in the text), while the Royal Target material is significantly stiffer and has larger wrinkles.

Images copyright Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi.

- ENGLISH, E., AND BRIDSON, R. 2008. Animating developable surfaces using nonconforming elements. *ACM Transactions on Graphics (SIGGRAPH 2008)* 27, 3 (August), 66:1–66:5.
- GOLDENTHAL, R., HARMON, D., FATTAL, R., BERCOVIER, M., AND GRINSUN, E. 2007. Efficient simulation of inextensible cloth. *ACM Transactions on Graphics (2007)* 26, 3 (July).
- GRINSUN, E., HIRANI, A. N., DESBRUN, M., AND SCHRÖDER, P. 2003. Discrete shells. In *Proc. of the ACM SIGGRAPH/Eurographics symposium on Computer animation*, 62–67.
- KALDOR, J. M., JAMES, D. L., AND MARSCHNER, S. 2008. Simulating knitted cloth at the yarn level. *ACM Transactions on Graphics (SIGGRAPH 2008)* 27, 3 (August), 65:1–65:9.
- KALDOR, J. M., JAMES, D. L., AND MARSCHNER, S. 2010. Efficient yarn-based cloth with adaptive contact linearization. *ACM Transactions on Graphics (SIGGRAPH 2010)* 29, 4 (July), 105:1–105:10.
- KATO TECH CO., L., 2011. KES-FB2-AUTO-A Pure Bending Tester, Jan. http://english.keskato.co.jp/products/kes_fb2.html.
- KATO TECH CO., L., 2011. KES-G2 Strip Biaxial Tensile Tester, Jan. http://english.keskato.co.jp/products/kes_g2.html.
- KAWABATA, S., POSTLE, R., AND NIWA, M. 1982. *Objective specification of fabric quality, mechanical properties and performance*. Textile Machinery Society of Japan.
- KING, M., JEANANAILAWONG, P., AND SOCRATE, S. 2005. A continuum constitutive model for the mechanical behavior of woven fabrics. *International Journal of Solids and Structures* 42, 13 (June), 3867–3896.
- KUNITOMO, S., NAKAMURA, S., AND MORISHIMA, S. 2010. Optimization of cloth simulation parameters by considering static and dynamic features. In *ACM SIGGRAPH 2010 Posters*, 15:1–15:1.
- MCCARTNEY, J., HINDS, B. K., AND KELLY, D. 2005. Modelling of anisotropic performance fabrics. *Journal of Materials Processing Technology* 159, 2, 181–190.
- MÜLLER, M. 2008. Hierarchical position based dynamics. In *Proc. of Virtual Reality Interactions and Physical Simulations*.
- NG, S., NG, R., AND YU, W. 2005. Bilinear approximation of anisotropic stress-strain. *Research Journal of Textile and Apparel* 12, 4 (Nov).
- O'BRIEN, J. F., AND HODGINS, J. K. 1999. Graphical modeling and animation of brittle fracture. In *Proc. of SIGGRAPH 99, Computer Graphics Proceedings, Annual Conference Series*, 137–146.
- PENG, X., AND CAO, J. 2005. A continuum mechanics-based non-orthogonal constitutive model for woven composite fabrics. *Composites: Part A*, 36, 859–874.
- PROVOT, X. 1995. Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Graphics Interface 95*, 147.
- SLAUGHTER, W. S. 2002. *The Linearized Theory of Elasticity*. Birkhauser.
- THOMASZEWSKI, B., PABST, S., AND STRASSER, W. 2009. Continuum-based strain limiting. In *Proc. of Eurographics 2009*, vol. 28, 569–576.
- VOLINO, P., MAGNENAT-THALMANN, N., AND FAURE, F. 2009. A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Transactions on Graphics* 28, 4 (September), 105:1–105:16.
- WANG, H., O'BRIEN, J., AND RAMAMOORTHY, R. 2010. Multi-resolution isotropic strain limiting. *ACM Transactions on Graphics (SIGGRAPH 98)* 29, 6 (December), 156:1–156:10.
- WHITE, R., CRANE, K., AND FORSYTH, D. A. 2007. Capturing and animating occluded cloth. *ACM Transactions on Graphics (SIGGRAPH 2007)* 26, 3 (July).
- WOJTAN, C., MUCHA, P. J., AND TURK, G. 2006. Keyframe control of complex particle systems using the adjoint method. In *Proc. of the ACM SIGGRAPH/Eurographics symposium on Computer animation*, 15–23.
- YEOMANA, M. S., REDDYB, D., BOWLESC, H. C., BEZUIDENHOUTD, D., ZILLAD, P., AND FRANZ, T. 2010. A constitutive model for the warp-weft coupled non-linear behavior of knitted biomedical textiles. *Biomaterials* 31 (November), 8484–8493.
- ZHANG, W., LEONARD, J. L., AND ACCORSI, M. L. 2005. Analysis of geometrically nonlinear anisotropic membranes: theory and verification. *Finite Elements in Analysis and Design* 41 (May), 963–988.
- ZHOU, N., AND GHOSH, T. 1998. On-line measurement of fabric bending behavior. *Textile Research Journal*, 7 (July), 533–542.
- ZWICK ROELL, G., 2011. Biaxial Tensile Test, Jan. <http://www.zwick.com/en/applications/metals/stripsheet/biaxial-tensile-test.html>.