

Modeling and Data-Driven Parameter Estimation for Woven Fabrics

David Clyde*
University of California, Los Angeles

Joseph Teran†
University of California, Los Angeles

Rasmus Tamstorf
Walt Disney Animation Studios



Figure 1: Photographs of five fabrics (canvas, cotton poplin, silk charmeuse, denim and wool coating) along with the corresponding simulation after parameter estimation (bottom row). Each piece of fabric is $1\text{m} \times 1\text{m}$, and the top corners are held 60cm apart. The real fabrics had multiple stable equilibria for the specified constraints, so exact matches cannot be expected.

ABSTRACT

Accurate estimation of mechanical parameters for simulation of woven fabrics is essential in many fields. To facilitate this we first present a new orthotropic hyperelastic constitutive model for woven fabrics. Next, we design an experimental protocol for characterizing real fabrics based on commercially available tests. Finally, we create a method for accurately fitting the material parameters to the experimental data. The last step is accomplished by solving inverse problems based on a Catmull-Clark subdivision finite element discretization of the Kirchhoff-Love equations for thin shells. Using this approach we are able to reproduce the fully nonlinear behavior corresponding to the captured data with a small number of parameters while maintaining all fundamental invariants from

*Also with Walt Disney Animation Studios.

†Also with Walt Disney Animation Studios.

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

SCA '17, July 28-30, 2017, Los Angeles, CA, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ISBN 978-1-4503-5091-4/17/07...\$15.00

<https://doi.org/10.1145/3099564.3099577>

continuum mechanics. The resulting constitutive model can be used with any discretization (e.g., simple triangle meshes) and not just subdivision finite elements. We illustrate the entire process with results for five types of fabric and compare photo reference of the real fabrics to the simulated equivalents.

CCS CONCEPTS

• Computing methodologies → Physical simulation;

KEYWORDS

constitutive modeling, orthotropy, cloth, data fitting

ACM Reference format:

David Clyde, Joseph Teran, and Rasmus Tamstorf. 2017. Modeling and Data-Driven Parameter Estimation for Woven Fabrics. In *Proceedings of SCA '17, Los Angeles, CA, USA, July 28-30, 2017*, 11 pages.

<https://doi.org/10.1145/3099564.3099577>

1 INTRODUCTION

Cloth simulation has broad applications ranging from the textile industry over entertainment to engineering applications such as fiber reinforcement in composites. In engineering and also for e-commerce applications such as virtual try-on, predictive power is obviously of paramount significance. By contrast, absolute accuracy is typically not a goal in entertainment. For graphics applications, it

is usually up to an artist to come up with a reasonable set of simulation parameters. Unfortunately, different artists often end up with different parameters for the same materials and since the approach is based on trial-and-error it can be difficult to reach satisfactory results. The desire for easier workflows when setting up simulations and also higher accuracy has motivated many researchers to look for experimental ways to determine simulation parameters. The appeal of such data-driven approaches is clear as it potentially allows for the creation of libraries of material parameters which can be used repeatedly and augmented as necessary.

In order for a data-driven approach to work, multiple components must be in place. First, a sufficiently versatile, accurate and preferably parsimonious cloth model must be established. Second, a set of experiments must be designed from which all model parameters can be estimated. From a practical point of view these experiments should be reproducible and since most end-users of cloth simulation are not experts in material testing, it is desirable that the tests are commercially available through a service bureau. Finally, a robust fitting method needs to be devised to fit the model parameters to the experimentally obtained data. In this paper we endeavor to provide all these components.

There is a plethora of different types of fabrics including knitted fabrics, woven fabrics, and a variety of specialty fabrics such as lace, 3d-fabrics, and non-wovens. A single model is unlikely to work well for all these different types of fabric, so in the following we will focus on woven fabrics. In order to fit a wide range of woven materials while minimizing mesh dependent behavior, we use a continuum assumption in our mechanical model. Since real clothing typically undergoes large deformations including large strains (especially in the shear component), it is essential to use a fully nonlinear model. To this end we begin by proposing a generalized orthotropic model characterized by a small number of parameters (§3). The orthotropy reflects the basic symmetry present in most woven fabrics, and by introducing nonlinearity in a way akin to Ogden’s constitutive model we are able to represent highly nonlinear behavior using a single model for the entire strain regime.

Much existing work related to characterization of fabrics is based on the Kawabata system, [Kawabata 1980]. However, these systems can be hard to find and are quite expensive. A more recent system, called FAST for “Fabric Assurance by Simple Testing” is much cheaper but limited to small strain deformations, [Minazio 1995]. In the graphics literature alternate methods have been proposed, but most of these require nonstandard equipment. The experimental protocol we propose, (§4), is largely based on existing ASTM standards¹ along with extensions which can be implemented using the same instruments as for the standard tests.

Data fitting is conceptually simple, but in practice fraught with peril. By using a fully implicit method for quasistatic simulation and by leveraging an infinitely smooth constitutive model, we are able to provide exact derivatives with respect to the model parameters in our optimization method based on the Broyden-Fletcher-Goldfarb-Shanno method (BFGS). Combined with good initial guesses based on bootstrap estimates this allows us to circumvent some of the data fitting challenges experienced in previous methods. We present our data fitting methodology in §5. All of

this is based on a Catmull-Clark subdivision finite element (FEM) discretization of the Kirchhoff-Love thin shell equations which is outlined in a supplemental document [Clyde et al. 2017]. Although the Kirchhoff-Love thin shell model is rather complex, its continuum basis is essential for providing parameters that depend minimally on mesh resolution and associated discrete anisotropy. Furthermore, we emphasize that the Kirchhoff-Love assumption, as well as the subdivision basis for FEM while useful for the fitting process are not necessary in practical simulation since the constitutive model can be used with simple linear strain triangles and commonly used graphics approaches for bending, e.g. [Bridson et al. 2003; Grinspun et al. 2003].

We summarize our novel contributions as:

- A hyperelastic constitutive model that separates the large strain from small strain behaviors in an intuitive way and thus allows for natural fitting to data.
- A BFGS optimization framework for fitting readily available ASTM data.
- Exact derivatives of objectives with respect to model parameters.
- Novel boot strapping heuristics for generating initial guesses in the optimization framework.
- A fully implicit Kirchhoff-Love thin shell model based on the OpenSubdiv library, as well as important observations for solving the associated linear systems.

2 RELATED WORK

Cloth simulation is a vast topic in both graphics and engineering, here we discuss just some of the work most related to ours. Recent work has proposed yarn-level simulations as a way to achieve accurate and very detailed results [Cirio et al. 2014; Kaldor et al. 2008]. These methods can create beautiful details, but require millions if not billions of degrees of freedom to represent typical garments. By adopting a continuum mechanics approach we aim to capture the cloth behavior at a tiny fraction of this cost. A middle-ground between the two approaches is the mesostructurally-based continuum model presented in [King et al. 2005]. This can capture yarn-level effects, but consequently also requires experimental data that characterizes the fabric at that level.

The estimation of clothing simulation parameters from real fabric deformation data has been investigated in recent years. Lubile and Magnenat-Thalmann [2008] compared the Kawabata system (KES) to the FAST system but found that they both have limitations. More recently [Power 2013] compared results from the FAST system to a newer Fabric Testing Kit (FTK) from Browzwear, but still found limitations. Despite these limitations the work by [Magnenat-Thalmann et al. 2007] uses the stress-strain curves from KES to drive an “accurate particle system”. This approach was further elaborated upon in [Volino et al. 2009]. Bhat et al [2003] developed an optimization procedure for estimating clothing simulation parameters directly from video data of moving fabrics. Beyond the already mentioned methods, existing methods from the textile industry for testing the mechanical properties of fabric are surveyed in [Wang et al. 2008].

A number of papers in graphics have proposed alternative and simple methods for measuring fabric properties. Wang et al. [2011]

¹Corresponding ISO standards also exist.

propose a sequence of tensile measurements using small loads (less than 15 N/m) for which they then fit a piecewise linear model. Miguel et al. [2012] use a more complicated setup for capturing data which include more complicated examples with shear buckling. However, they also primarily focus on the low strain domain (typically less than 200 N/m). By comparison KES tests up to 500 N/m. While the smallest strain regime is most important for typical cloth simulations, higher strains do occur even in everyday use of garments and are therefore also important to consider. Examples of higher strains occur especially near garment seams where parts of the fabric are physically constrained.

The experimental data of Wang et al. [2011] is based on a sparse set of measurements, and thus may not be suitable for fitting a more complex range of strain regimes. The experiments of Miguel et al. [2012] rely on complex, hand-constructed machinery, which requires careful manual labor by researchers. Our approach uses a wide range of experimental data from the readily available ASTM standard tests, and we examine our experimental error by repeating experiments with additional samples.

Similar to our approach, Miguel et al. [2016] construct a single hyperelastic energy function for fabric. Unlike our approach they do not assume that cloth possesses orthotropic symmetry but instead allow for a more general anisotropic form, which they then fit against a subset of the data from [Miguel et al. 2012]. Other papers in graphics have considered orthotropic hyperelastic energies with a focus on how to make it intuitive for an artist to design and control orthotropic behavior, [Li and Barbič 2014; Xu et al. 2015]. We believe this is complementary to our work as the goal in animation is often plausibility rather than true reality.

In this work we limit ourselves to elastic deformations, but it is well-known that cloth exhibits significant amounts of hysteresis due to internal friction [Miguel et al. 2013; Williams 2010]. Additionally we do not consider rate-dependent behavior as they do in [Ryou et al. 2007], but we believe that leveraging these ideas is an important area of future work.

Our use of subdivision finite elements is greatly inspired by [Cirak and Ortiz 2001; Cirak et al. 2000]. That work is based on the Loop subdivision scheme, and was also the inspiration for the work related to cloth simulation by [Thomaszewski et al. 2006]. More recently the Catmull-Clark subdivision scheme has been considered for linearly elastic thin shells by [Wawrzinek et al. 2011]. We extend that work through the use of a fully implicit time integration scheme and a nonlinear constitutive model. A similar extension based on B-splines has been presented by [Kiendl et al. 2015].

3 CONSTITUTIVE MODEL

We provide extended details of our Kirchhoff-Love thin shell model in [Clyde et al. 2017]. With this approach, the cloth constitutive behavior is defined in terms of the potential energy density ψ . We assume that the constitutive model is orthotropic to accurately represent the anisotropy introduced by the warp and weft structure of woven cloth. Beyond that, the constitutive model must be sufficiently complex to capture the nonlinear stress-strain relationship exhibited by these materials, yet simple enough to facilitate our incremental fitting process. We design the model in terms of parameters that intuitively relate to specific deformation modes and strain

regimes. Mathematically, we write the energy as $\psi = \psi(\mathbf{E}, \mathbf{D}, \mathbf{k})$ where \mathbf{E} is the Green-Lagrange strain, $\mathbf{D} = [\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3]$ are the undeformed configuration warp/weft orthotropy ($\mathbf{d}_1, \mathbf{d}_2$) and out of plane (\mathbf{d}_3) directions, and $\mathbf{k} = [a_{11}, a_{12}, a_{22}, G_{12}]^T$ is a short list of material stiffnesses.

To capture the orthotropic components of the strain, we express \mathbf{E} in the basis $\tilde{\mathbf{E}} = \mathbf{D}^T \mathbf{E} \mathbf{D}$. Notably $\tilde{\mathbf{E}}$ exhibits a simple block structure due to the Kirchhoff-Love kinematic assumptions which prevent any stretches and shears in the out-of-plane direction \mathbf{d}_3 . This means that $\tilde{\mathbf{E}}$ has the form

$$\tilde{\mathbf{E}} = \begin{pmatrix} \tilde{E}_{11} & \tilde{E}_{12} & 0 \\ \tilde{E}_{12} & \tilde{E}_{22} & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (1)$$

To preserve the required orthotropic symmetry, we need the energy density ψ to be invariant under replacement of \mathbf{E} with $\mathbf{Q}^T \mathbf{E} \mathbf{Q}$ for any element \mathbf{Q} of the orthotropic symmetry group. One can check that any such replacement of \mathbf{E} will leave \tilde{E}_{11} and \tilde{E}_{22} unchanged, and may alter \tilde{E}_{12} only by reversing its sign. In other words, the orthotropy requirement will be satisfied for any energy of the form

$$\psi(\mathbf{E}, \mathbf{D}, \mathbf{k}) = \tilde{\psi}(\tilde{E}_{11}, \tilde{E}_{22}, |\tilde{E}_{12}|, \mathbf{k}).$$

In fact, this form retains full generality in the sense that any orthotropic function of \mathbf{E} can be written this way. This may be shown by examining the function basis for orthotropic functions specified in [Zheng 1993] and noting that each basis function can be expressed in terms of \tilde{E}_{11} , \tilde{E}_{22} , and $|\tilde{E}_{12}|$. Of the three invariants, \tilde{E}_{11} directly measures the strain due to warp stretch; likewise \tilde{E}_{22} and \tilde{E}_{12} measure the weft stretch and in-plane shearing strains, respectively.

Our model is defined in terms of four scalar functions $\eta_1, \eta_2, \eta_3, \eta_4$ to separate energy penalties for the distinct deformation modes

$$\psi = \frac{a_{11}}{2} \eta_1(\tilde{E}_{11}^2) + a_{12} \eta_2(\tilde{E}_{11} \tilde{E}_{22}) + \frac{a_{22}}{2} \eta_3(\tilde{E}_{22}^2) + G_{12} \eta_4(\tilde{E}_{12}^2) \quad (2)$$

The functions η_j are arbitrary except for the constraints $\eta_j(0) = 0$ and $\eta_j'(0) = 1$. The first constraint enforces a zero-energy, zero-stress rest configuration, while the second constraint allows a natural correspondence between the parameters a_{ij} and G_{12} and linear elasticity at infinitesimal strain. With this convention, the a_{ij} and G_{12} should be interpreted as describing the cloth's small-strain behavior while the curves η_j describe the nonlinear response to larger strains. Indeed, we can replicate the orthotropic St. Venant-Kirchhoff model [Başar et al. 2000] by simply choosing $\eta_j(x) = x$ for all j .

3.1 Parameterization

The scalar functions η_j in Eq. (2) must be sufficiently adjustable to fit the nonlinear stress-strain behavior of various material types. Simple spline representations are sufficiently general and easily adjustable by hand. However, we choose a form which is more suited to manipulation by an iterative inverse solver to facilitate subsequent data fitting. Inspired by [Itskov 2001] we use the notion of integer 'degree' $d_j \geq 1$ together with scalar parameters

Parameters	Meaning
ρ	Mass density
τ	Thickness
$a_{11}, a_{12}, a_{22}, G_{12}$	Infinitesimal strain parameters
$\mu_{1j}, \alpha_{1j}, 1 \leq j \leq d_1$	Nonlinear warp stretch response
$\mu_{2j}, \alpha_{2j}, 1 \leq j \leq d_2$	(Roughly) area preservation
$\mu_{3j}, \alpha_{3j}, 1 \leq j \leq d_3$	Nonlinear weft stretch response
$\mu_{4j}, \alpha_{4j}, 1 \leq j \leq d_4$	Nonlinear shear response
$\tilde{E}_{ij}^{\max}, \tilde{E}_{ii}^{\min}$	Transition point to extrapolation

Table 1: Complete list of parameters for determining cloth behavior in our model.

$\alpha_{j1}, \dots, \alpha_{jd_j}$ and $\mu_{j1}, \dots, \mu_{jd_j}$ to define η_j as

$$\eta_j(x) = \sum_{i=1}^{d_j} \frac{\mu_{ji}}{\alpha_{ji}} ((x+1)^{\alpha_{ji}} - 1).$$

This is akin to the constitutive model by Ogden, [Ogden 1972], but expressed in terms of the orthotropic integrity basis $\{\tilde{E}_{11}, \tilde{E}_{22}, \tilde{E}_{12}^2\}$ instead of the principal stretches. Satisfaction of the constraint $\eta_j(0) = 0$ is automatic, while the constraint $\eta_j'(0) = 1$ is equivalent to requiring $\sum_{i=1}^{d_j} \mu_{ji} = 1$.

3.2 Extrapolation

Our convention for η_j is effective at describing the experimental strain regime, but can give unpredictable results for strains outside the fitting dataset. Our data covers strains up to the breaking point of the cloth, but we still need a plausible extrapolation to regimes where the real-life cloth would have torn apart. Although such configurations should not be observed in practice, they can temporarily occur in the iterative solvers used for fitting and simulation.

To address this, for each entry \tilde{E}_{ij} of $\tilde{\mathbf{E}}$ we define corresponding ‘‘strain cutoffs’’, \tilde{E}_{ij}^{\min} and \tilde{E}_{ij}^{\max} , based on the boundaries of our collected data. If all three of the strain entries satisfy $\tilde{E}_{ij}^{\min} \leq \tilde{E}_{ij} \leq \tilde{E}_{ij}^{\max}$, then we proceed with the usual energy evaluation as described above. If some strain entries lie outside the valid regime, we instead evaluate our energy via a two term power series approximation expanded around the closest ‘valid’ strain. The choice of exactly two terms in the power series expansion is the simplest possible while still preserving the C^2 nature of the extrapolated energy density necessary for fitting and for implicit simulation. To maintain orthotropic symmetry, we must use $\tilde{E}_{12}^{\min} = -\tilde{E}_{12}^{\max}$. Table 1 provides the full list of parameters.

4 EXPERIMENTAL DESIGN

Our experimental approach is chosen to exercise all important deformation modes without sacrificing ease of obtainability. Standard, commercially-available procedures are strongly preferred over any test that requires physical construction of a complex apparatus. For this reason we focus on the ASTM testing standards. In section §7 we present results for canvas, cotton, silk, denim, and wool. The protocol for denim and wool is different than the remaining materials as noted below. To measure in-plane stretch and shear response,

we use slight modifications of the tensile test ASTM D5035. The modified test uses cloth strips cut along the warp (90°), weft (0°) and bias (45°) direction. The strips used are 75 mm long (in the main direction) and 50 mm wide. For the denim and wool the strips are 150 mm long. The strips are clamped at both ends and gradually stretched until the cloth tears apart, with the clamps moving apart at a rate of 300 mm/min (200 mm/min for denim and wool). The stretch distance and clamp forces are recorded at least once per 0.1 second throughout the test for canvas, cotton and silk and every 2 milliseconds for denim and wool. For future experiments we recommend the highest possible sampling rate. For comparison purposes, each direction is tested five times for canvas, cotton and silk using a new cloth strip each time. We additionally run this test with cloth strips cut at 22.5° and 67.5° relative to the warp direction for model validation. For denim and wool we only run two repeats of each test and omit the tests for 22.5° and 67.5° . Specifically for this paper, all of the above experiments were performed using an Instron 5569 tensile test machine.

Many different methods have been devised for measuring bend resistance of fabrics. One of the earliest methods based on a simple cantilever principle dates back to [Peirce 1930]. Yet, this is still the method being used in the ASTM D1388 method as well as the FAST system. The basic idea behind the Kawabata system was first published by [Livesey and Owen 1964], and both approaches are reviewed by [Ghosh and Zhou 2003]. We refer to the latter for diagrams illustrating these tests. To measure bending resistance, we use ASTM D1388. A 25 mm \times 200 mm cloth strip is slowly extended off the edge of a fixed plane and allowed to drape toward a second plane which is inclined at 41.5° . The recorded data is the overhang length of the cloth when it first makes contact with the inclined plane. This test is repeated for both warp- and weft-oriented cloth strips. We obtain the mass density of each material based on ASTM D3776, and we acquire a measured thickness for each material from ASTM D1777. However, the thickness of a woven fabric is not easy to rigorously define or measure; thus we treat the measured thickness as an initial guess as opposed to a guarantee of the correct value to be used in our model.

We remark that our dataset contains little or no information about response to negative stretch strains, i.e. in-plane compression. We do not know of a test procedure for that strain regime because of cloth’s typical tendency to buckle instead of undergoing any significant compression. It should also be noted that there is no standard for measuring the shear properties of fabrics. The in-plane stretch test along the bias direction is known as the bias extension test and is well-known [Cooper 1963], but is typically used with simplifying assumptions that are not valid for large strains, [Dridi et al. 2011]. However, due to our data fitting process we do not have to rely on those assumptions.

Lastly, we note that our dataset does not provide data to fit the term $a_{12}\eta_2(\tilde{E}_{11}\tilde{E}_{22})$ in our energy definition in Eq. (2). This term is related to the (pseudo) Poisson ratio for the fabric. The Poisson effect in fabrics is primarily due to crimp interchange which is most notable with unbalanced weaves, [Sun et al. 2005]. The proposed methods in the literature for capturing data for this term are based on digital image correlation (DIC), [Hursa et al. 2009]. This is commercially available in so called video extensometers, and

would conceptually be easy to add to our fitting method. However, we have found this term to have very little effect on the forces in our other experiments since the fabrics considered here have very little crimp to begin with. The more noticeable transversal contractions due to shearing, as seen in the bias extension test, are modeled accurately with our data.

5 DATA FITTING

In this section we present the methodology for fitting our model to the experimental data. First we construct a nonlinear objective function to minimize in §5.1. To facilitate the optimization we fit the parameters in multiple stages which we describe in §5.2. Due to the complexity of the objective function it is beneficial to start with good initial guesses, which we obtain through a bootstrap method (§5.3); this process also informs our choice of the degree parameters d_j . Finally we discuss the nonlinear solver in §5.5.

5.1 Optimization objective

The experimental observations consist of force-displacement data points for each of warp, weft, and bias stretch as well as bend test data for each of the warp and weft directions. Let \mathbf{k} denote the vector of all parameters of the constitutive model. For a given displacement with corresponding measured force \hat{f} , we can recreate the same stretch distance using hard constraints to simulate clamps in our finite element solver and compute the resulting equilibrium. This gives a simulated force $\hat{f} = \hat{f}(\mathbf{k})$. Similarly, the bend test data states that some specified cloth overhang length drapes to an inclination angle of exactly $\hat{\theta} = 41.5^\circ$. Again, we can recreate the test within our finite element solver by allowing a cloth strip of the same length to drape under gravity and measuring the equilibrium drape angle $\hat{\theta} = \hat{\theta}(\mathbf{k})$.

We denote list of measurements to be fit as $\{\check{p}_i : i \in I\}$ where I is an appropriate set of indices; thus each \check{p}_i is either a tensile force or a bend angle. We define the quality of a proposed fit via the following optimization objective function ε :

$$\varepsilon(\mathbf{k}) = \sum_{i \in I} \frac{(\hat{p}_i(\mathbf{k}) - \check{p}_i)^2}{\check{p}_i^2} \quad (3)$$

It is important to prioritize *relative* error in order to accurately match the shape of the stress-strain curve through all experimental regimes. Our datasets cover a wide variety of strain levels, so a method based on absolute error would over-prioritize accuracy in the highest-strain experiments without accurately fitting the small strain data.

5.2 Fitting stages

We design our fitting approach around the effects of the various parameters \mathbf{k} . In particular, many of the test data points are almost unaffected by certain parameters leading to decoupling (or weak coupling) of some parameters. In particular, the warp stretch results depend almost exclusively on the parameters τ , a_{11} , μ_{1j} , and α_{1j} (for all $1 \leq j \leq d$). Similarly, the weft stretch results depend almost exclusively on τ , a_{22} , μ_{3j} , and α_{3j} . The bias stretch forces depend on all parameters relevant for warp or weft stretch, plus the parameters G_{12} , μ_{4j} and α_{4j} . Finally, the bend tests depend on the thickness τ

together with the parameters relevant for a stretch along that same material direction.

In light of these dependencies, we design the following fitting strategy which relies on a series of minimizations of the objective function in Eq. (3) over subsets of the experimental data. The procedure is designed such that each optimization step will provide fits for some additional parameter(s), while having no significant effect on the simulated test results fitted in the preceding steps.

- (1) Select an initial guess parameter set \mathbf{k}^0 . The thickness guess is the experimentally measured thickness. The extrapolation cutoffs \tilde{E}_{ij}^{\max} , \tilde{E}_{ii}^{\min} are initially turned off completely. The remaining parameters' guesses are chosen as described in §5.3 below.
- (2) Minimize the objective summed over warp stretch data points only, by solving for a_{11} , μ_{1j} , and α_{1j} with all other parameters held fixed at their values from \mathbf{k}^0 .
- (3) Minimize the objective summed over weft stretch data points only, by solving for a_{22} , μ_{3j} , and α_{3j} with all other parameters held fixed at their values from \mathbf{k}^0 .
- (4) Choose the strain cutoffs \tilde{E}_{11}^{\max} and \tilde{E}_{22}^{\max} to be 95% of the largest strain present in the corresponding tensile tests. (Note that \tilde{E}_{11}^{\min} and \tilde{E}_{22}^{\min} should always be 0 for our datasets, since we have no experimental data measuring response to warp or weft compression.)
- (5) Minimize the objective summed over bias stretch data points only, by solving for G_{12} , μ_{4j} , and α_{4j} with all the preceding steps' parameters fixed at their solved-for values and all remaining parameters fixed at their values from \mathbf{k}^0 .
- (6) Choose the strain cutoff \tilde{E}_{12}^{\max} to equal 95% of the largest strain \tilde{E}_{12} present in the bias tensile test.
- (7) Finally, minimize the objective summed over the warp and weft bend data points for the thickness τ . As τ is updated in the iterative process we modify the other small strain parameters using $a_{ij}^{\text{new}} = \frac{\tau^0}{\tau} a_{ij}$ and $G_{12}^{\text{new}} = \frac{\tau^0}{\tau} G_{12}$. All tested tensile forces scale very nearly linearly with both τ and the small-strain parameter set $\{a_{11}, a_{12}, a_{22}, G_{12}\}$; thus the parameter modifications in this step do not impact the fits of tensile data.

The parameters μ_{ij} and α_{ij} are not necessarily restricted to be positive, and in several cases the solver's final fit does include negative parameters.

Our data does not exercise the energy term $a_{12}\eta_2(\tilde{E}_{11}\tilde{E}_{22})$ sufficiently to allow fitting. In the linearization at zero strain, a_{12} determines the orthotropic Poisson's ratios via $\nu_{12} = \frac{a_{12}}{a_{11}}$ and $\nu_{21} = \frac{a_{12}}{a_{22}}$. In light of this, we use the St. Venant-Kirchhoff form $\eta_2(x) = x$ and set $a_{12} = \frac{1}{2} \min\{a_{11}, a_{22}\}$ for simplicity.

5.3 Initial guesses

The minimization of Eq. (3) is generally difficult and benefits from good initial values in \mathbf{k}^0 . Fortunately, we can approximate the optimization problem for the warp and weft parameters by assuming very simple equilibrium configurations. In particular, we approximate the equilibrium by assuming that the entire cloth undergoes constant strain. Based on this assumption and an observed stretch deformation $\Delta\ell$, we can compute the total hyperelastic energy of the cloth strip, Ψ . Then the force component along the stretch

direction can be computed as $\frac{\partial \Psi}{\partial (\Delta \ell)}$. This approximation of the equilibrium forces gives rise to a simplified version of the optimization objective ε , by replacing the simulated equilibrium forces \hat{p}_i with their newly approximated counterparts. We then solve this simpler problem and use the resulting parameter set as the initial guess for the original problem in the main solver.

These approximate forces are relatively easy to differentiate allowing the use of a full Newton optimization solve instead of the quasi-Newton methods in our main solver. Further, the quicker function evaluations make it feasible to run thousands of iterations until convergence. The forces computed using the approximate equilibrium configuration typically come within 1% of the forces obtained using the full quasistatic solve, with the worst error rising up near 3%. This means our approximate optimization problem is a very good model for the full solve, and therefore the initial guesses \mathbf{k}^0 we obtain are very close to the final fitted parameters \mathbf{k} . This generally saves a great deal of runtime in the main solver.

Ideally we would use similar approximate equilibria for all experiments in the dataset. However, we are not aware of any appropriate simplifying assumptions for the bias stretch or bending tests. For tensile tests in the bias direction, it is well known that an approximate equilibrium can be computed for small strains using the piecewise-constant strain assumption, [Dridi et al. 2011]. However, this approximating assumption proved unreliable in our testing, and in any case cannot be applied at larger strains.

5.4 Choice of degree parameters

In the list of fitting stages provided in §5.2, steps 2, 3, and 5 require fitting the curves η_1 , η_3 , and η_4 respectively. An iterative method is suitable for determining the parameters α_{jk} and μ_{jk} , but the integer degree parameters d_j must be determined by other means. The choice of d_j strongly affects the difficulty of the fitting process. If d_j is too small, η_j may not be sufficiently expressive to model the materials' nonlinear behavior under large deformations. If d_j is too large, we introduce near-nullspaces in the fitting process, which will complicate the optimization problem. In practice, the ideal d_j depends on the specific material being fitted, and thus we choose d_j on a per-solve basis.

To find the simplest possible model that fits the data, each solve stage is attempted first with a degree 1 curve. The degree is then incremented until a satisfactory fit is obtained. In warp or weft fitting, we compute initial guesses for each new curve degree based on the procedure of §5.3. For bias fits with degree 1 curves, the optimization is simple enough to be reliably solved regardless of the initial guess chosen. When running higher degree bias fits, the previous lower-degree curve is used as an initial guess.

5.5 Optimization method

For the final optimizations steps in §5.2 we explored using both BFGS and Gauss-Newton methods, together with a line search protocol which guarantees satisfaction of the strong Wolfe conditions. Here we discuss the steps needed in computing the functional $\varepsilon(\mathbf{k})$ and its derivatives needed in these approaches.

5.5.1 Equilibrium configuration. Given a mesh configuration \mathbf{q} , the corresponding clamp forces (for tensile tests) or drape angle (for bend test) can be computed as a simple function of \mathbf{q} . However, our

optimization objective function $\varepsilon = \varepsilon(\mathbf{k})$ depends on various simulated *equilibrium* measurements $\hat{p}_i(\mathbf{k})$. In other words, to evaluate the terms $\hat{p}_i(\mathbf{k})$ appearing in the definition of ε (e.g. at each new solver step or line search iteration) we must compute the associated mesh equilibrium configuration.

We use $\mathbf{q}^e = \mathbf{q}^e(\mathbf{k})$ to denote the equilibrium configuration determined by the parameters \mathbf{k} . These are obtained by solving the force balance in the absence of inertia (see supplementary technical document [Clyde et al. 2017] for derivation) while satisfying the linear boundary condition constraints

$$\begin{aligned} \mathbf{f}(\mathbf{q}^e) + \mathbf{B}^T \boldsymbol{\lambda} &= 0 \\ \mathbf{B} \mathbf{q}^e &= \mathbf{b}. \end{aligned} \quad (4)$$

As outlined in [Clyde et al. 2017], $\mathbf{f}(\mathbf{q})$ denotes the combination of elastic and gravity forces and \mathbf{B} expresses the boundary conditions. For a tensile experiment, these boundary conditions describe the clamped edges held fixed at the appropriate stretch distance. For a bend experiment, the constraints keep the cloth draping over the edge of a fixed plane.

5.5.2 Differentiating the equilibrium configuration. To minimize $\varepsilon(\mathbf{k})$, both BFGS and Gauss-Newton require differentiation of each term in the sum in Eq. (3). That is, for a given experimental measurement \check{p} , we must compute

$$\frac{\partial}{\partial \mathbf{k}} \left(\frac{(\hat{p}(\mathbf{k}) - \check{p})^2}{\check{p}^2} \right). \quad (5)$$

Each simulated force $\hat{f}(\mathbf{k})$ is a linear function of the corresponding equilibrium positions \mathbf{q}^e . Simulated angles $\hat{\theta}(\mathbf{k})$ additionally require composition with an inverse trigonometric function. In either case, the main difficulty of evaluating Eq. (5) lies in computation of $\frac{\partial \mathbf{q}^e}{\partial \mathbf{k}}$.

The function $\mathbf{q}^e(\mathbf{k})$ is defined implicitly via the system in Eq. (4). Thus to compute its derivative we must differentiate the entire equilibrium system. Note that the term $\mathbf{f}(\mathbf{q}^e)$ depends on \mathbf{k} both through $\mathbf{q}^e = \mathbf{q}^e(\mathbf{k})$ and through the usual dependence of elastic forces on the model parameters \mathbf{k} . The differentiated system is:

$$\begin{aligned} \frac{\partial \mathbf{f}}{\partial \mathbf{q}} \frac{\partial \mathbf{q}^e}{\partial \mathbf{k}} + \frac{\partial \mathbf{f}}{\partial \mathbf{k}} + \mathbf{B}^T \frac{\partial \boldsymbol{\lambda}}{\partial \mathbf{k}} &= 0 \\ \mathbf{B} \frac{\partial \mathbf{q}^e}{\partial \mathbf{k}} &= 0, \end{aligned}$$

Here $\frac{\partial \mathbf{f}}{\partial \mathbf{q}} = -\mathbf{K}$ is the negative Hessian of the elastic potential V , and $\frac{\partial \mathbf{f}}{\partial \mathbf{k}}$ can be computed directly. Then the unknowns $\frac{\partial \mathbf{q}^e}{\partial \mathbf{k}}$ and $\frac{\partial \boldsymbol{\lambda}}{\partial \mathbf{k}}$ can be obtained by solving the linear system

$$\begin{pmatrix} \mathbf{K} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \frac{\partial \mathbf{q}^e}{\partial \mathbf{k}} \\ -\frac{\partial \boldsymbol{\lambda}}{\partial \mathbf{k}} \end{pmatrix} = \begin{pmatrix} \frac{\partial \mathbf{f}}{\partial \mathbf{k}} \\ \mathbf{0} \end{pmatrix}. \quad (6)$$

We remark that the matrices $\frac{\partial \mathbf{q}^e}{\partial \mathbf{k}}$ and $\frac{\partial \mathbf{f}}{\partial \mathbf{k}}$ each have $|\mathbf{k}|$ columns. Thus Eq. (6) really represents $|\mathbf{k}|$ linear systems, each with the same system matrix but a different right-hand side. As an alternative to the method presented here, one could consider the adjoint method [Giles and Pierce 2000]. However, since the number of constraints due to boundary conditions is almost always larger than $|\mathbf{k}|$ this would be more expensive.

5.5.3 *BFGS vs. Gauss-Newton.* The complexity of the fitting solves depends heavily on the choice of the degree d_j for each function η_j in the energy definition. For degree 1, either optimization method is reliable. However, Gauss-Newton frequently stalls in fitting curves of degree 2 or larger, repeatedly choosing search directions that allow virtually no progress. BFGS is robust in solving degree 2 problems, and handles degree 3 problems well when using the initial guesses specified in §5.3. We generally find degree 2 or 3 to be the minimum requirement for accurate fitting of our datasets. Thus the improved performance of BFGS as compared to Gauss-Newton is quite important.

5.5.4 *Multiresolution approach.* Cloth simulation based on subdivision surface finite elements has the advantage of higher-order convergence under mesh refinement as compared to simpler finite elements or approximations thereof. We use this feature to speed up our fitting process. In particular, we run each fitting solve using a comparatively coarse mesh resolution to describe each simulated experiment. The resulting parameters are then used as an initial guess for a higher resolution fitting solve. This process may be repeated as many times as desired to obtain satisfactory convergence of high-resolution fitting solves more quickly than would be possible without such a means of generating initial guesses. In practice, we typically use this technique only for the solves involving tensile data. For canvas, cotton and silk, we use an 8×12 grid of finite elements for our coarsest meshes. Then we progress upward through resolutions 16×24 and finally 32×48 to refine the fit. For denim and wool we use 8×24 , 16×48 , and 32×96 grids.

6 IMPLEMENTATION

6.1 Quasistatic solver

As discussed in §5.5.1, our fitting process requires solving many quasistatics problems of the form in Eq. (4) for each BFGS iteration. The nonlinear quasistatic problems are solved using Newton’s method with line search to minimize the total of hyperelastic and potential cloth energy subject to the given linear constraints. We note that the hyperelastic nature of our constitutive model is essential for this. Unfortunately, the resulting linear systems often exhibit condition numbers on the order of 10^{10} , which easily can lead to inaccurate solutions. As a result the Newton step can sometimes be an ascent direction for the energy or otherwise fail to produce significant progress, even when the current iterate is visibly not a local energy minimizer. The solver escapes from such situations by taking several backward Euler time steps before returning to the Newton iteration process. This use of backward Euler simulation as a failsafe is triggered anytime three consecutive Newton steps give a combined energy reduction of 0.1% or less, or immediately if Newton chooses an ascent direction for the cloth energy. A potential alternative is to leverage the method in [Volino and Magnenat-Thalmann 2007].

6.2 Linear solver

Our fitting solver and simulator both rely heavily on solution of symmetric indefinite linear systems. For this we use the PARDISO solver in Intel’s MKL, [Schen and Gärtner 2006]. PARDISO operates by explicit matrix factorization followed by back-substitution. This

Fabric	Density [g/m ²]	Thickness [mm]	Thread count [ends/” × picks/”]	Weave
Canvas	294.0	0.53	108 × 56	Basket
Cotton	103.6	0.18	140 × 86	Plain
Silk	83.0	0.18	653 × 123	Satin
Wool	480.6	1.28	32 × 28	Twill
Denim	400.0	0.66	70 × 47	Twill

Table 2: Measured properties of the test materials.

means consecutive solves of the same system are much faster. For example, our computation of the derivative $\frac{\partial q^e}{\partial k}$ in §5.5.2 requires solution of the same system matrix for $|k|$ different right-hand sides, and thus benefits from a factorization-based linear solver.

6.3 Precision

As mentioned above, cloth quasistatics problems tend to be extremely ill-conditioned. Thus it is critical to retain as much precision as possible. We address this in three steps. First, we upgrade OpenSubdiv from single to double precision arithmetic. Second, we use iterative refinement in our linear solves, using quad precision to store the residual. Finally, we are careful to choose a sufficient number of Gaussian quadrature points. We find 4×4 quadrature points to be sufficient. Despite this, there still exist fairly simple quasistatics problems on which our solver fails to make progress. However, our solver is robust enough to handle all problems which arise during the course of our model fitting procedure.

6.4 Interpretation of data from repeated tests

Each tensile experiment in our datasets is repeated multiple times, using a freshly cut cloth strip for each iteration. Comparing the resulting force-displacement data frequently shows relative differences of up to 20% for a fixed elongation. The cloth strip lengths and force measurements in our tensile tests are very precise due to the nature of the machinery used. We believe the major source of experimental error is imprecision in lining up the target warp-vs-stretch-direction angles. Targeting a certain stretch angle can be a very difficult task, especially for extremely light or stretchy materials. This task is in general easiest for warp and weft directions; we generally expect the angle error to be larger for the off-angle stretches in the 45° (bias), 22.5° , and 67.5° directions. For warp or weft experiments, a strip stretched along the target direction will produce larger forces than a strip stretched along a direction offset by a few degrees. Thus, given data from multiple test executions we choose the curve with *largest* forces for fitting purposes, assuming that the corresponding experiment came closest to using the target stretch direction. There is no simple way to determine which bias cloth sample was most accurately aligned; thus we randomly select one sample for fitting.

7 RESULTS

The material model and fitting procedure are demonstrated by fitting parameters for five materials: canvas, cotton poplin, silk charmeuse, wool coating, and denim; see Table 2. Our target criterion for a successful fit is for the magnitude of fitting error to be

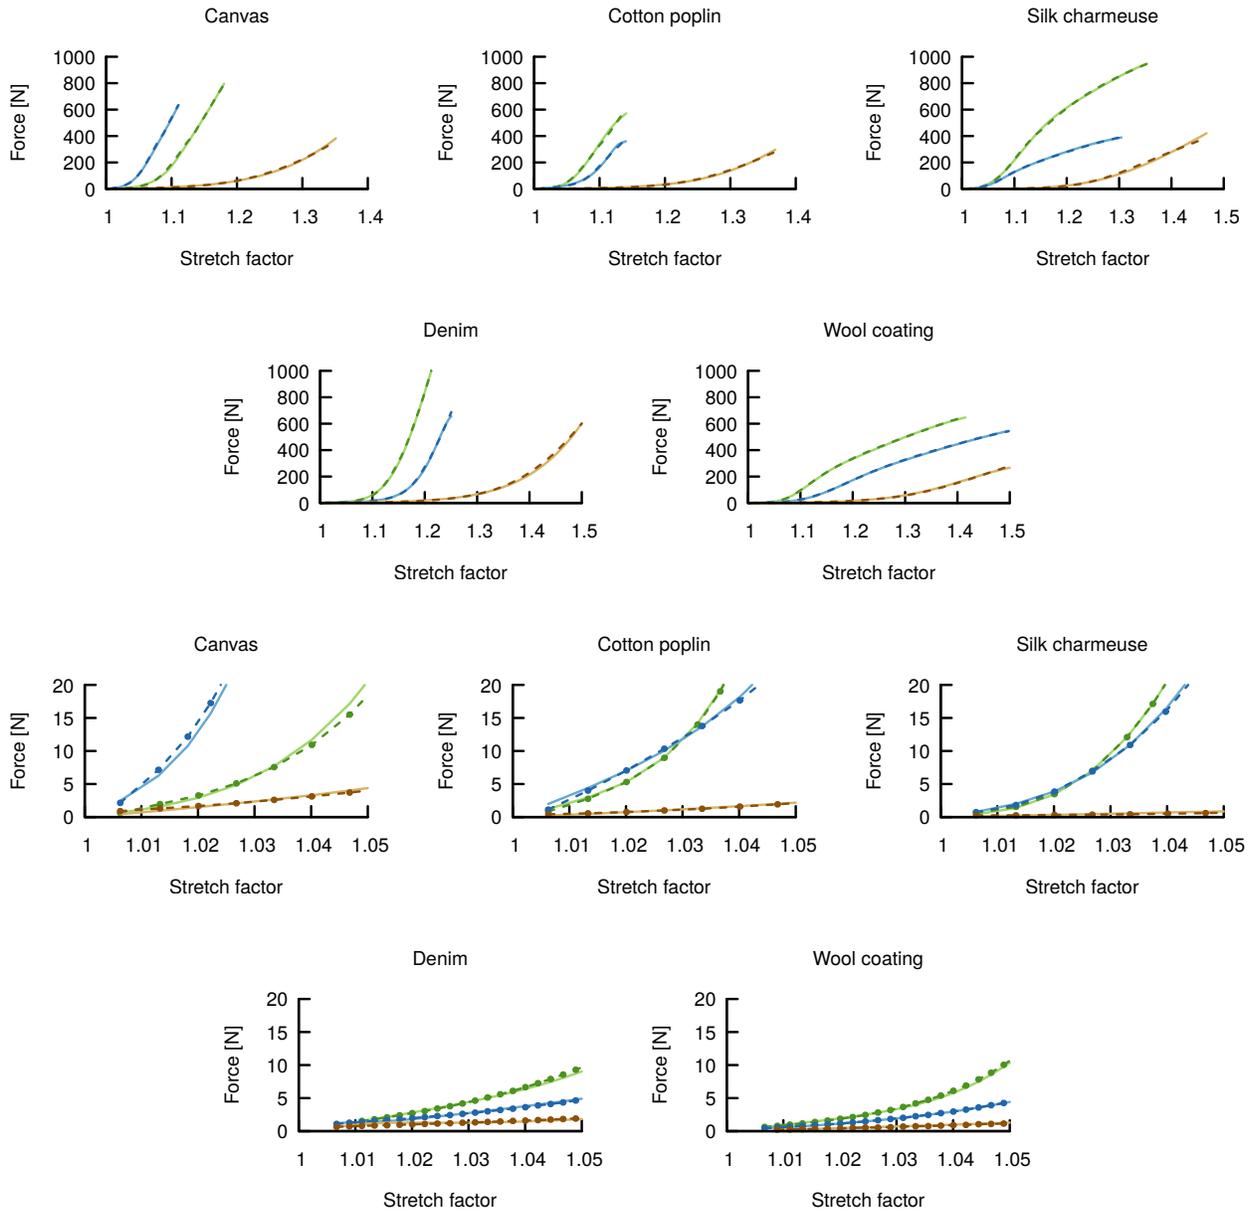


Figure 2: Force-elongation curves for 5 cm wide strips of canvas, cotton, silk, denim, and wool. Measured data is shown with dashed lines while fitted curves are shown in solid. The top set of plots shows the full strain domain while the bottom set shows a closeup for the small-strain regime. Each plot shows curves for warp (green), weft (blue) and bias (brown). Dots represent data points but are only shown in the small strain plots for legibility.

smaller than the unavoidable error sources inherent in our framework. The most significant of these is experimental error: the relative difference between 5 executions of the same tensile experiment is often 15-20% for a fixed elongation. Smaller errors are introduced through the numerical discretization, nonlinear solver residual, Kirchhoff-Love hypotheses, and interpretation of data from slowly-moving tensile experiments as representing quasistatic equilibria.

We provide plots comparing experimental vs modeled forces in Figure 2. Quantitatively, the accuracy of fitted curves is judged by computing the average relative error over the domain. This is just a rescaling of L^1 error to view the result as a percentage. The fitting process produces average relative error under 5% in most cases; see Table 3. Thus our fits are well within the level of experimental error in the data.

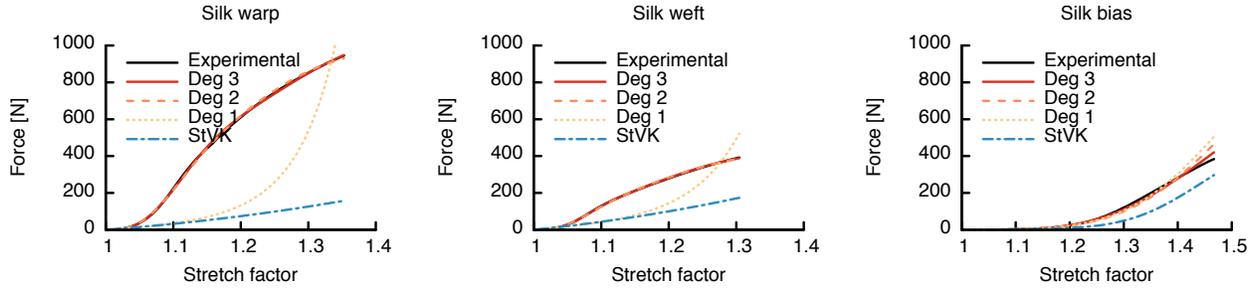


Figure 3: Demonstration of the improvement of fit quality as the model complexity increases for silk. For warp and weft, degree 2 is sufficient to fit the experimental data well (average relative error under 2%). For bias, the transition from degree 2 to 3 still makes significant progress (average relative error reduces from 6% to 4.5%). In all cases, an orthotropic St. Venant-Kirchhoff model provides a poor fit. Due to our minimization of relative error, the small strain regime is prioritized even when using models too simple to match the full curve. For the remaining fabrics, the final curve degrees for warp, weft, and bias respectively were 3, 3, and 2 (cotton); 2, 3 and 3 (canvas); 5, 5, and 4 (denim); and 5, 5, and 3 (wool).

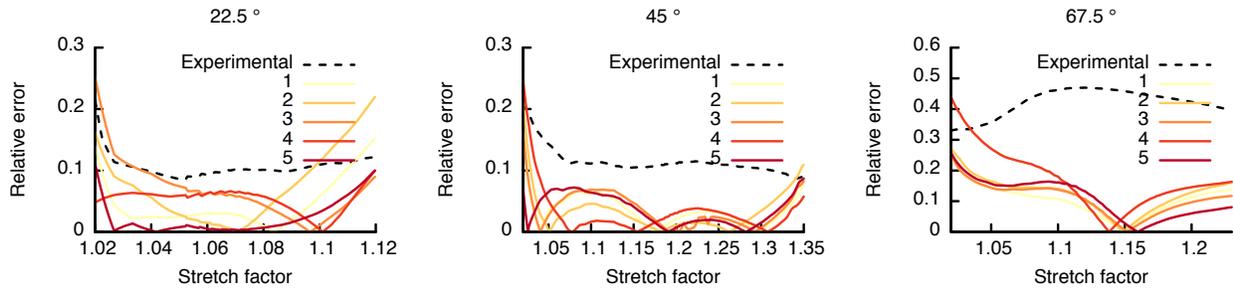


Figure 4: Relative error plots for tensile tests at 22.5°, 45°, and 67.5° after adjusting for accidental cloth sample misalignments. The black curve shows the relative difference between the largest and the smallest unadjusted force measurements. Each of the remaining curves show the relative error of the model fit compared to the experimental data for one cloth specimen. With very few exceptions the fitting error is well below the experimental error.

Fabric	Warp	Weft	Bias
Canvas	1.9 %	2.0 %	5.1 %
Cotton poplin	2.9 %	3.2 %	2.6 %
Silk charmeuse	1.9 %	1.7 %	4.5 %
Wool coating	1.1 %	0.8 %	3.4 %
Denim	1.6 %	2.3 %	4.7 %

Table 3: Average relative error of the fitted model vs. the experimental data.

Figure 3 demonstrates the progression of fit quality as the degrees d_j are increased and the resulting choice of degrees to fit the data accurately while minimizing model complexity. The bending tests match less accurately. Recall that by design of the bending experiment, all cloth samples should drape downward at an angle of 41.5°. However, our fitted models produce a wide range of drape angles for some materials; see Table 4. A more sophisticated bend test like the one used in the Kawabata system could be beneficial. Alternatively, the Kirchhoff-Love assumptions might need modification to capture the bending response for some types of cloth. That said, the popular discrete shells bending model as used in

[Miguel et al. 2016] and slightly modified for [Wang et al. 2011] can also be derived as a discretization of the Kirchhoff-Love kinematic assumptions.

Fabric	Warp drape	Weft drape
Canvas	54.6°	23.7°
Cotton poplin	43.5°	39.2°
Silk charmeuse	52.8°	26.2°
Wool coating	42.2°	40.7°
Denim	50.6°	30.0°

Table 4: Bend test drape angles for the fitted models. To match experimental data, all drape angles should equal 41.5°.

For validation, the modeled forces are compared against tensile experiments at stretch directions 22.5° and 67.5° from weft. Each validation experiment is executed 5 times. To allow for the difficulty of accurately cutting strips at specific angles, we solve for an unknown angle deviation $\Delta\theta$ in each experimental run; thus the model forces used for comparison assume the strip is aligned at 22.5 + $\Delta\theta$ or 67.5 + $\Delta\theta$ degrees. With this adjustment, the relative

errors between model and experimental forces are generally much smaller than experimental error. The relative errors from the canvas validation are shown in Figure 4 as a function of the elongation. The relative errors for all the fabrics are summarized in Table 5. While these errors are larger than the fitting errors, they are consistently below the variation in the experimental data.

Fabric	Model		Experiment	
	22.5°	67.5°	22.5°	67.5°
Canvas	5.8 %	12.8 %	11.0 %	42.0 %
Cotton poplin	7.3 %	11.1 %	22.4 %	30.6 %
Silk charmeuse	9.8 %	5.6 %	36.2 %	22.3 %

Table 5: The average relative error of the model vs. the experimental data for the validation tests. Each number is the average over 5 experiments.

The qualitative look of the fitted materials is demonstrated in a series of simulations. The drape test (Figure 1) uses $1\text{m} \times 1\text{m}$ squares of each material suspended from two corners held 60cm apart. This is similar to the setup used by [Wang et al. 2011] but with a different distance between the corners. The simulated results may be compared against photographs of the same test run with real fabric samples. The picture frame and corner pull tests (Figure 5) show deformations of a $10\text{cm} \times 10\text{cm}$ cloth square. This latter is similar to the setup used by [Miguel et al. 2012]. In the picture frame test, the cloth is held along all edges and the constraints are moved to induce large shearing deformation. In the corner pull test, the cloth is stretched diagonally using constraints at the corners.

All data fitting and simulations were run on an Intel Xeon E5-2698 v4 CPU with 64 GB RAM. Fitting each new material currently requires approximately 4 hours of compute time, together with some operator intervention to decide which degrees d_j produce the desired level of fitting accuracy.

8 DISCUSSION

Under the Kirchhoff-Love hypotheses, bending resistance is determined by the response to small strains ($< 3\%$ in our examples). The current tensile datasets exhibit large relative error between repeated tests in this regime. We believe this causes the discrepancies in our fitting of bend data; thus the bend fits could be improved using data from a tensile experiment which provides increased accuracy in small strain.

Our current optimization objective in Eq. (3) penalizes relative error in all captured data points with equal weight. However, the tensile test machinery provides slightly varying sample density across various regimes of the strain domain. Further, it is unclear what distribution of sampled strains is desirable. This can be addressed by some choice of constant weights for each term in Eq. (3), or by interpolating an experimental curve and then sampling points according to some chosen target measurement density (e.g., uniform).

We do not attempt to model hysteresis from internal friction in the cloth. However, this is an important factor for the behavior of real cloth. Previously, this has been investigated by [Miguel et al.

2013] and we hope that our fitting method can be extended to also estimate the internal friction parameters. Similar to hysteresis, we have ignored any dependence in the deformation behavior on strain rate. Additional experimentation might be needed to determine whether this is reasonable.

When the ultimate goal is to create pleasing and/or accurate cloth simulations, it should be noted that the interaction between the fabric and its environment is just as important as the behavior of the cloth itself. Most obvious in this context is correct contact handling. However, especially for light weight fabrics (e.g., chiffon) the coupling between the cloth and the fluid flow of the air around it is also essential.

The Catmull-Clark subdivision surface-based method used for fitting the experimental data allows use of the Kirchhoff-Love hypotheses, but the resulting non-interpolating finite element basis means that standard triangle mesh cloth collision algorithms cannot be used. Extending our work to implement a collision scheme would enable simulation of more everyday situations involving cloth at high strains, which would allow greater benefit from our accurately fit hyperelastic response in that regime. However, we note that the results obtained from our fitting process can be used with any discretization; in particular they can be used with triangle meshes for which existing contact handling algorithms apply.

9 ACKNOWLEDGMENTS

Thanks to Karen Tamstorf for help preparing the cloth samples for the reference photos, and to Colin Eckart and Norman Joseph for help rendering. The mechanical testing was performed by Vartest Laboratories. The work is supported by NSF CCF-1422795, ONR (N000141110719, N000141210834), DOD (W81XWH-15-1-0147) as well as a gift from Disney Research.

REFERENCES

- Y. Başar, M. Itskov, and A. Eckstein. 2000. Composite laminates: nonlinear interlaminar stress analysis by multi-layer shell elements. *Computer Methods in Applied Mechanics and Engineering* 185, 2-4 (2000), 367–397. [https://doi.org/10.1016/S0045-7825\(99\)00267-4](https://doi.org/10.1016/S0045-7825(99)00267-4)
- Kiran S. Bhat, Christopher D. Twigg, Jessica K. Hodgins, Pradeep K. Khosla, Zoran Popović, and Steven M. Seitz. 2003. Estimating Cloth Simulation Parameters from Video. In *Proc 2003 ACM SIGGRAPH/Eurographics Symp Comp Anim (SCA '03)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 37–51.
- Robert Bridson, Sebastian Marino, and Ronald Fedkiw. 2003. Simulation of clothing with folds and wrinkles. In *SCA '03: Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. Eurographics Association, 28–36.
- Fehmi Cirak and Michael Ortiz. 2001. Fully C^1 -conforming subdivision elements for finite deformation thin-shell analysis. *Internat. J. Numer. Methods Engrg.* 51, 7 (2001), 813–833. <https://doi.org/10.1002/nme.182>
- Fehmi Cirak, Michael Ortiz, and Peter Schröder. 2000. Subdivision surfaces: A new paradigm for thin-shell finite-element analysis. *Internat. J. Numer. Methods Engrg.* 47, 12 (2000), 2039–2072. [https://doi.org/10.1002/\(SICI\)1097-0207\(20000430\)47:12<2039::AID-NME872>3.0.CO;2-1](https://doi.org/10.1002/(SICI)1097-0207(20000430)47:12<2039::AID-NME872>3.0.CO;2-1)
- Gabriel Cirio, Jorge Lopez-Moreno, David Miraut, and Miguel A. Otaduy. 2014. Yarn-level Simulation of Woven Cloth. *ACM Trans. Graph.* 33, 6, Article 207 (Nov. 2014), 11 pages. <https://doi.org/10.1145/2661229.2661279>
- David Clyde, Joseph Teran, and Rasmus Tamstorf. 2017. Simulation of nonlinear Kirchhoff-Love thin shells using subdivision finite elements. (2017). Supplemental technical document.
- D. N. E. Cooper. 1963. A Bias Extension Test. *Textile Research Journal* 33, 4 (1963), 315–317. <https://doi.org/10.1177/004051756303300408>
- Samia Dridi, Abdelwaheb Dogui, and Philipp Boisse. 2011. Finite element analysis of bias extension test using an orthotropic hyperelastic continuum model for woven fabric. *The Journal of The Textile Institute* 102, 9 (2011), 781–789. <https://doi.org/10.1080/00405000.2010.522048>
- Tushar K. Ghosh and Naiyue Zhou. 2003. Characterization of fabric bending behavior: A review of measurement principles. *Indian Journal of Fibre & Textile Research* 28,

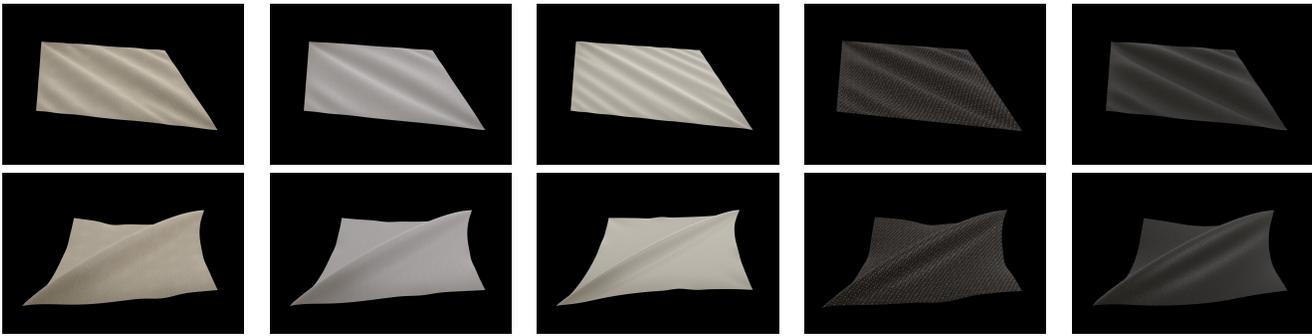


Figure 5: Fitted materials are demonstrated in a simulated picture frame test (first row) and corner pull test (second row). The materials from left to right are canvas, cotton, silk, denim, and wool.

- 4 (2003), 471–476. <http://hdl.handle.net/123456789/22691>
- Michael B. Giles and Niles A. Pierce. 2000. An Introduction to the Adjoint Approach to Design. *Flow, Turbulence and Combustion* 65, 3 (2000), 393–415. <https://doi.org/10.1023/A:1011430410075>
- Eitan Grinspun, Anil N. Hirani, Mathieu Desbrun, and Peter Schröder. 2003. Discrete shells. In *SCA '03: Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. Eurographics Association, 62–67.
- A. Hursa, T. Rolich, and S. Ercegović Ražić. 2009. Determining Pseudo Poisson's Ratio of Woven Fabric with a Digital Image Correlation Method. *Textile Research Journal* 79, 17 (2009), 1588–1598. <https://doi.org/10.1177/0040517509104316>
- Mikhail Itskov. 2001. A generalized orthotropic hyperelastic material model with application to incompressible shells. *Internat. J. Numer. Methods Engrg.* 50, 8 (2001), 1777–1799. <https://doi.org/10.1002/nme.86>
- Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2008. Simulating Knitted Cloth at the Yarn Level. *ACM Trans. Graph.* 27, 3, Article 65 (Aug. 2008), 9 pages. <https://doi.org/10.1145/1360612.1360664>
- Sueo Kawabata. 1980. *The Standardization and Analysis of Hand Evaluation* (2nd edition ed.). The Textile Machinery Society of Japan, Osaka, Japan.
- Josef Kiendl, Ming-Chen Hsu, Michael C.H. Wu, and Alessandro Reali. 2015. Isogeometric Kirchhoff-Love shell formulations for general hyperelastic materials. *Computer Methods in Applied Mechanics and Engineering* 291 (2015), 280–303. <https://doi.org/10.1016/j.cma.2015.03.010>
- M. J. King, P. Jearanaisilawong, and S. Socrate. 2005. A continuum constitutive model for the mechanical behavior of woven fabrics. *International Journal of Solids and Structures* 42, 13 (2005), 3867–3896. <https://doi.org/10.1016/j.ijsolstr.2004.10.030>
- Yijing Li and Jernej Barbič. 2014. Stable Orthotropic Materials. In *ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA '14)*, Vladlen Koltun and Eftychios Sifakis (Eds.). The Eurographics Association, 41–46. <https://doi.org/10.2312/sca.20141121>
- R. G. Livesey and J. D. Owen. 1964. Cloth Stiffness and Hysteresis in Bending. *Journal of the Textile Institute Transactions* 55, 10 (1964), T516–T530. <https://doi.org/10.1080/19447026408662430>
- Christiane Luible and Nadia Magnenat-Thalmann. 2008. The Simulation of Cloth Using Accurate Physical Parameters. In *Proceedings of the Tenth IASTED International Conference on Computer Graphics and Imaging (CGIM '08)*. ACTA Press, Anaheim, CA, USA, 123–128. <http://www.miralab.ch/bibliography/431.pdf>
- Nadia Magnenat-Thalmann, Christiane Luible, Pascal Volino, and Etienne Lyard. 2007. From Measured Fabric to the Simulation of Cloth. In *Computer-Aided Design and Computer Graphics, 2007 10th IEEE International Conference on*. 7–18. <https://doi.org/10.1109/CADCG.2007.4407832>
- E. Miguel, D. Bradley, B. Thomaszewski, B. Bickel, W. Matusik, M. A. Otaduy, and S. Marschner. 2012. Data-Driven Estimation of Cloth Simulation Models. *Computer Graphics Forum* 31, 2pt2 (May 2012), 519–528. <https://doi.org/10.1111/j.1467-8659.2012.03031.x>
- Eder Miguel, David Miraut, and Miguel A. Otaduy. 2016. Modeling and Estimation of Energy-Based Hyperelastic Objects. *Computer Graphics Forum* 35, 2 (2016). <https://doi.org/10.1111/cgf.12840>
- Eder Miguel, Rasmus Tamstorf, Derek Bradley, Sara C. Schwartzman, Bernhard Thomaszewski, Bernd Bickel, Wojciech Matusik, Steve Marschner, and Miguel A. Otaduy. 2013. Modeling and Estimation of Internal Friction in Cloth. *ACM Trans. Graph.* 32, 6, Article 212 (Nov. 2013), 10 pages. <https://doi.org/10.1145/2508363.2508389>
- Pier Giorgio Minazio. 1995. FAST - Fabric Assurance by Simple Testing. *International Journal of Clothing Science and Technology* 7, 2/3 (1995), 43–48. <https://doi.org/10.1108/09556229510087146>
- R. W. Ogden. 1972. Large Deformation Isotropic Elasticity - On the Correlation of Theory and Experiment for Incompressible Rubberlike Solids. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 326, 1567 (1972), 565–584. <https://doi.org/10.1098/rspa.1972.0026>
- F. T. Peirce. 1930. The "Handle" of Cloth as a Measurable Quantity. *Journal of the Textile Institute Transactions* 21, 9 (1930), T377–T416. <https://doi.org/10.1080/19447023008661529>
- Jess Power. 2013. Fabric objective measurements for commercial 3D virtual garment simulation. *International Journal of Clothing Science and Technology* 25, 6 (2013), 423–439. <https://doi.org/10.1108/IJCST-12-2012-0080>
- Hansun Ryou, Kwansoo Chung, and Woong-Ryeol Yu. 2007. Constitutive modeling of woven composites considering asymmetric/anisotropic, rate dependent, and nonlinear behavior. *Composites Part A: Applied Science and Manufacturing* 38, 12 (2007), 2500–2510. <https://doi.org/10.1016/j.compositesa.2007.08.003>
- Olaf Schenk and Klaus Gärtner. 2006. On fast factorization pivoting methods for sparse symmetric indefinite systems. *ETNA. Electronic Transactions on Numerical Analysis [electronic only]* 23 (2006), 158–179. <http://eudml.org/doc/127439>
- Huiyu Sun, Ning Pan, and Ron Postle. 2005. On the Poisson's ratios of a woven fabric. *Composite Structures* 68, 4 (2005), 505–510. <https://doi.org/10.1016/j.compstruct.2004.05.017>
- Bernhard Thomaszewski, Markus Wacker, and Wolfgang Straßer. 2006. A Consistent Bending Model for Cloth Simulation with Corotational Subdivision Finite Elements. In *Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '06)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 107–116. <http://dl.acm.org/citation.cfm?id=1218064.1218079>
- Pascal Volino and Nadia Magnenat-Thalmann. 2007. Stop-and-go cloth draping. *The Visual Computer* 23, 9 (2007), 669–677. <https://doi.org/10.1007/s00371-007-0152-5>
- Pascal Volino, Nadia Magnenat-Thalmann, and Francois Faure. 2009. A simple approach to nonlinear tensile stiffness for accurate cloth simulation. *ACM Trans. Graph.* 28, Article 105 (September 2009), 16 pages. Issue 4. <https://doi.org/10.1145/1559755.1559762>
- Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi. 2011. Data-Driven Elastic Models for Cloth: Modeling and Measurement. *ACM Transactions on Graphics (SIGGRAPH 2011)* 30, 4, Article 71 (Aug. 2011), 12 pages. <https://doi.org/10.1145/2010324.1964966>
- X. Wang, X. Liu, and C. Hurren Deakin. 2008. Physical and mechanical testing of textiles. In *Fabric Testing*, Jinlian Hu (Ed.). Woodhead Publishing, Chapter 4, 90–124. <https://doi.org/10.1533/9781845695064.90>
- Anna Wawrzinek, Klaus Hildebrandt, and Konrad Polthier. 2011. Koiter's Thin Shells on Catmull-Clark Limit Surfaces. In *Vision, Modeling, and Visualization (2011)*, Peter Eisert, Joachim Hornegger, and Konrad Polthier (Eds.). The Eurographics Association. <https://doi.org/10.2312/PE/VMV/VMV11/113-120>
- Robert W. Williams. 2010. *Measuring and modeling the anisotropic, nonlinear and hysteretic behavior of woven fabrics*. Ph.D. Dissertation. University of Iowa. <http://ir.uiowa.edu/etd/907>
- Hongyi Xu, Funshing Sin, Yufeng Zhu, and Jernej Barbič. 2015. Nonlinear Material Design Using Principal Stretches. *ACM Trans. Graph.* 34, 4, Article 75 (July 2015), 11 pages. <https://doi.org/10.1145/2766917>
- Q.-S. Zheng. 1993. On transversely isotropic, orthotropic and relative isotropic functions of symmetric tensors, skew-symmetric tensors and vectors. Part I: Two dimensional orthotropic and relative isotropic functions and three dimensional relative isotropic functions. *International Journal of Engineering Science* 31, 10 (1993), 1399–1409. [https://doi.org/10.1016/0020-7225\(93\)90005-F](https://doi.org/10.1016/0020-7225(93)90005-F)