Estimation of Yarn-Level Simulation Models for Production Fabrics

GEORG SPERL, Institute of Science and Technology Austria, Austria
ROSA M. SÁNCHEZ-BANDERAS, SEDDI, Spain
MANWEN LI, Under Armour, USA
CHRIS WOJTAN, Institute of Science and Technology Austria, Austria
MIGUEL A. OTADUY, Universidad Rey Juan Carlos, Spain

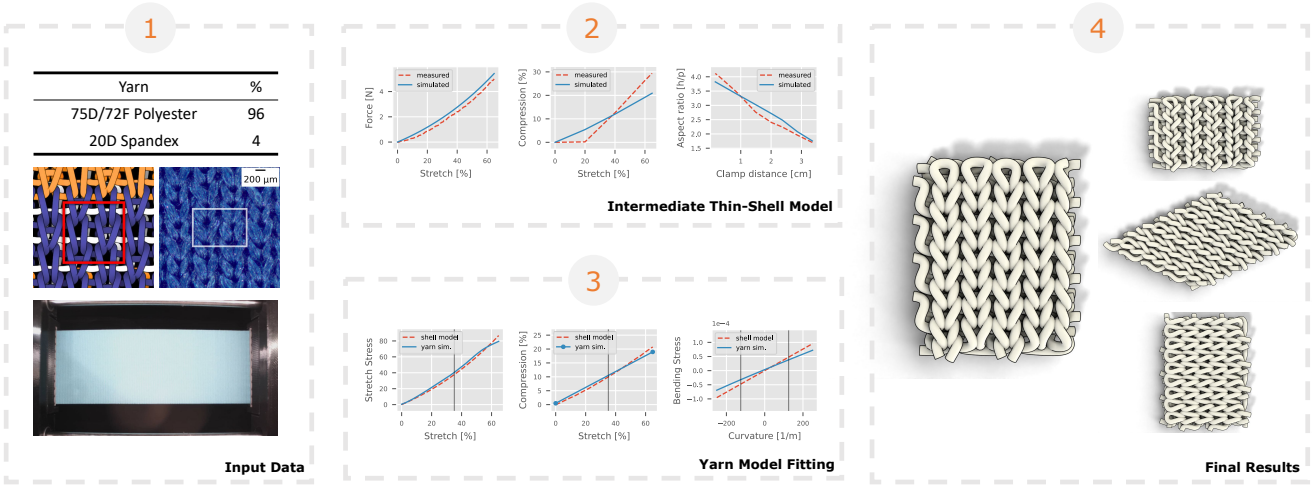

Fig. 1. The figure shows our pipeline to fit yarn-level mechanical parameters to real-world knits, applied to a double knit pique fabric (DKP). (1) We take as input the fabric composition, knit schematics, high-resolution photographs, and swatch-level physical tests. (2) We fit a thin-shell model to the non-uniform physical data, and we use it to generate target uniform data. (3) Then, we fit the yarn-level model to the uniform data, leveraging periodic simulations to reduce the computational cost. (4) The images show the yarn model for DKP, under uniform stretch on the weft, bias, and warp directions.

This paper introduces a methodology for inverse-modeling of yarn-level mechanics of cloth, based on the mechanical response of fabrics in the real world. We compiled a database from physical tests of several different knitted fabrics used in the textile industry. These data span different types of complex knit patterns, yarn compositions, and fabric finishes, and the results demonstrate diverse physical properties like stiffness, nonlinearity, and anisotropy.

We then develop a system for approximating these mechanical responses with yarn-level cloth simulation. To do so, we introduce an efficient pipeline for converting between fabric-level data and yarn-level simulation, including a novel swatch-level approximation for speeding up computation, and some small-but-necessary extensions to yarn-level models used in computer graphics. The dataset used for this paper can be found at http://mslab.es/projects/YarnLevelFabrics.

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

Authors' addresses: Georg Sperl, Institute of Science and Technology Austria, Klosterneuburg, Austria, georg.sperl@ist.ac.at; Rosa M. Sánchez-Banderas, SEDDI, Madrid, Spain, rosa.sanchez@seddi.com; Manwen Li, Under Armour, Baltimore, USA, manwenli9302@gmail.com; Chris Wojtan, Institute of Science and Technology Austria, Klosterneuburg, Austria, wojtan@ist.ac.at; Miguel A. Otaduy, Universidad Rey Juan Carlos, Madrid, Spain, miguel.otaduy@urjc.es.

Additional Key Words and Phrases: yarn-level cloth, knits, parameter estimation, simulation

ACM Reference Format:
Georg Sperl, Rosa M. Sánchez-Banderas, Manwen Li, Chris Wojtan, and Miguel A. Otaduy. 2022. Estimation of Yarn-Level Simulation Models for Production Fabrics. *ACM Trans. Graph.* 41, 4, Article 65 (July 2022), 15 pages. https://doi.org/10.1145/3528223.3530167

1 INTRODUCTION

The simulation of cloth at the yarn level has demonstrated outstanding capabilities, allowing results not possible with classic mesh-based simulation models. Thanks to yarn-level models, we can produce animation results with extreme detail, as well as mechanical behavior that exhibits structural nonlinearity [Cirio et al. 2014; Kaldor et al. 2008]. Yarn-level models also enable visual design of complex knit patterns [Leaf et al. 2018].

Previous works have shown that computer-graphics yarn-level models can provide a qualitative match to the mechanical response of real-world fabrics; however, they have not tried to match this mechanical response in a quantitative manner. In this project, we have documented, scanned, and tested a library of knit fabrics from real textile production, and we have fitted yarn-level simulation models that match their mechanical response.

In doing this, we have faced two major challenges. First, we demonstrate that real-world yarns exhibit complex deformation behavior that is not sufficiently captured by the Kirchhoff rod models typically used in computer animation and numerical simulation. Second, the best real-world data available is in the form of physical tests *at the swatch-level*, which is computationally prohibitive to simulate efficiently *at the level of individual yarns* — straightforward simulation-in-the-loop parameter estimations are intractable for this problem.

We introduce the first technique for the modeling and estimation of yarn-level fabric mechanics that succeeds to capture the macroscopic (swatch-level) response of textile-production knitted fabrics. We achieve this through three main contributions that address the challenges discussed above.

Data set of real-world fabrics. We compiled a data set of physical test data from 33 different knitted fabrics used by industry professionals in the production of casual and sports garments. The fabrics span different knit patterns (e.g., multiple layers), yarn compositions (e.g., plated yarns), and yarn finishes. On a macroscopic level, they show diverse stiffness, nonlinearity, and anisotropy. The data set consists of physical information about each yarn type, manually registered yarn geometry, high-resolution photographic scans, and physical measurements from experiments for each fabric type, and it will be made available to the research community.

Efficient fitting procedure. To estimate yarn-level parameters from swatch-level physical tests, we have designed a two-step procedure that circumvents the computational cost of simulating full fabric swatches at yarn level: we first fit a thin-shell model to swatch-level data [Miguel et al. 2012; Wang et al. 2011], then we generate analytical stress-strain data using the thin-shell model, and we finally fit a periodic version of the yarn-level model [Sperl et al. 2020].

Practical and versatile simulation models. Basic models for yarns and thin shells cannot capture the diversity of behaviors of real-world knitted fabrics, while complex models with a large number of parameters are vulnerable to overfitting. After experimenting with many of these models and observing how well they fit real-world data, we propose a few minimal extensions to typical models used in computer graphics to help strike the balance between simplicity and expressive power: an anisotropic area-preserving thin shell model, and a yarn model with two-phase stretching and contact energies.

In this work, we focus on the major aspects of macroscopic mechanical response, including nonlinearity and anisotropy of stretch, shear, and bending deformation. We leave for future work more complex aspects such as extreme nonlinearity, hysteresis, or curling. Under these limitations, we maximize parallelism between the data, parameterization, and estimation processes of thin-shell and yarn-level fitting; we do this to minimize the error introduced by using the thin-shell model as an intermediate representation, while circumventing the challenge of simulating full-swatch non-uniform deformations at yarn level.

After discussing related work, we give an overview of our approach in Section 3. We then describe the input data to the yarn-model estimation process in Section 4. We describe the thin-shell model, its estimation, and how it is used as an intermediate target for the yarn-level model in Section 5, and we continue with the description of the yarn-level model in Section 6 and the estimation of its parameters in Section 7. We discuss our results in Section 8 and end with a discussion and conclusion in Section 9 and Section 10.

2 RELATED WORK

2.1 Yarn-Level Cloth Simulation

Over two decades ago, Rémion et al. [1999] simulated knitted fabrics as deformable rods in contact, in contrast to finite-element modeling of yarn volumes, as commonly done in textile engineering [Liu et al. 2018]. Later, Kaldor et al. [2008] demonstrated the ability to simulate full garments at the yarn level. They showed that such models could reproduce qualitative macroscopic behavior of real-world knits. To accelerate computational speed, Kaldor et al. [2010] later introduced optimizations in the update of costly nonlinear contact terms. With the same goal of accelerating computations, Cirio et al. [2014] introduced an approach that avoids contact detection altogether, thanks to a reduced model of sliding yarns in persistent contact. This model was initially applied to woven fabrics, and later extended to knits [Cirio et al. 2017] and multi-layer fabrics [Sánchez-Banderas et al. 2020].

As discussed in the introduction, an important gap in yarn-level modeling of fabrics is their connection to real-world materials. Leaf et al. [2018] showed that varying the stiffness and rest-shape parameters of yarns could lead to the design of yarn patterns that matched the geometry of complex real-world knits. However, they did not validate the mechanical behavior of the resulting fabric models. In their work, they leveraged a key tool for computational efficiency, as they simulated periodic yarn patches under uniform deformations. We borrow this approach in our work. Schumacher et al. [2018] also used computations under uniform deformations with periodic boundary conditions, to characterize the mechanical response of two-dimensional microtextures. While these were not fabrics, they share similarities in the underlying deformation model. In contrast to their 'forward' characterization, in our work we look at inverse modeling of yarn-level fabrics. In textile engineering, recent works have looked at correctly initializing the yarn geometry for volumetric finite-element simulation [Wadekar et al. 2020]. In our experience, this leaves many open unknowns, such as the rest-shape geometry of the yarns and the contact model.

Another line of work has addressed the computational cost of full yarn-level simulations by bridging the gap with thin shell models. One approach is to enable hybrid simulations [Casafranca et al. 2020], focusing yarn-level computational effort only where needed. Another approach is to estimate thin-shell models that best-fit the mechanical response of yarn-level models [Sperl et al. 2020]. This approach uses simulations with periodic boundary conditions as discussed above, and applies the theory of homogenization to thin-shell deformation. This work has been later extended to incorporate yarn-level deformation variance in texture map representations for visual display [Sperl et al. 2021].

One of the main targets of yarn-level models in computer graphics has been the development of authoring tools. They include data structures to represent knit geometry on surface meshes [Wu et al. 2019; Yuksel et al. 2012], or algorithms to simplify the editing and fabrication of real-world knits [Kapllani et al. 2021; Narayanan et al. 2019] and even 3D woven fabrics [Wu et al. 2020].

Our work also relates to the inverse design of simulated fibers [Derouet-Jourdan et al. 2010; Hadap 2006]. Aside from our application, these ideas are useful for hair animation [Derouet-Jourdan et al. 2013], the physical fabrication of curve networks [Pérez et al. 2015; Zehnder et al. 2016], and the design of materials which couple rods and elastic sheets [Pérez et al. 2017].

2.2 Estimation of Mechanical Parameters of Cloth

The estimation of the parameters of cloth models has addressed to date thin-shell models that describe fabric as a continuum. This line of work has a long history in the textile engineering field, where the approach was to design mechanical tests that could elicit mechanical properties in a separable way [Kawabata 1980; Peirce 1930]. When fabric is deformed, it is challenging to apply uniform stress; therefore, multiple mechanical properties lead to a complex interplay. We face a similar challenge when trying to circumvent the computational cost of yarn-level simulations by leveraging periodic deformations.

In computer graphics, we can distinguish two main directions in the estimation of cloth simulation models. One direction focuses on the accuracy of the estimation, in particular trying to address the nonlinear and anisotropic behavior of cloth. Works in this direction entail the design of physical tests to produce force-deformation examples, parameterization of the cloth model, and estimation algorithms [Miguel et al. 2012; Wang et al. 2011]. Some works have considered the hysteresis behavior in cloth deformations, by estimating models of internal friction [Miguel et al. 2013].

Another direction focuses on estimating simulation models from more casual data, such as video. The pioneering work of Bhat et al. [2003] used optimization methods to estimate mass, elasticity and damping parameters from videos of cloth motion. Bouman et al. [2013] used a machine-learning technique, where they learned a mapping from cloth model parameters to video features, and then inverted this mapping to fit parameters to new video footage. This approach has received major thrust with the explosion of deep learning methods [Runia et al. 2020; Yang et al. 2017]. Recent works look at the design of semi-controlled setups where particular mechanical properties are exposed, e.g., contact friction [Rasheed et al. 2020].

Modern machine learning methods have also posed the problem of efficient simulation-in-the-loop optimization. To this end, research on differentiable simulation answers how to compute gradients of (dynamic) equilibrium constraints for cloth simulation [Liang et al. 2019].

3 OVERVIEW

The main goal of our system is to efficiently solve for the yarn-level simulation parameters that will reproduce the large-scale in-plane responses measured by real-world knitted fabrics. The brute-force approach—simulating every yarn in a swatch of fabric, comparing the results to the measured data, and looping until the correct parameters are found—is computationally infeasible, as a single swatch can easily contain tens of thousands of knit loops. To avoid full-scale simulation, we approximate the fabric swatch as a periodically repeating pattern; this allows us to take advantage of recent technology for simulating knits with periodic boundary conditions [Leaf et al. 2018; Sperl et al. 2020, 2021].

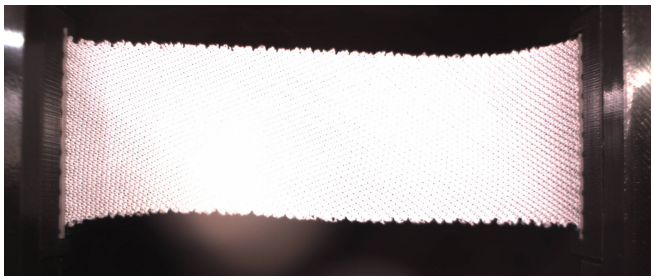

Fig. 2. This image shows an all-needle fabric (A2) stretched along the bias direction. Notice the spatially non-uniform shear and curved shape near the clamps.

However, simply tiling these periodic yarn physics over a patch the size of the physical knitted fabric swatch is insufficient to reproduce our physical test data. In particular, it cannot model the spatially non-uniform deformations that occur when fabrics are sheared (Figure 2). To ensure that our method performs well even in the presence of spatially non-uniform deformations, we introduce an intermediate thin-shell model and solve for a physical model that reproduces the test data. Once we have this model, we generate spatially uniform deformations that can finally be modeled with periodic yarn-level simulation.

Our parameter-fitting pipeline is illustrated in Figure 1. We first solve for a thin-shell model that reproduces the real world data samples. Next, we use this thin-shell model to generate spatially uniform data samples. We then fit yarn-level simulation parameters to these new uniform data.

We want our simulation models to be complex enough to model important features in the data, while simultaneously minimizing additional complexity to avoid overfitting. We discuss in Section 5 how a simple StVK thin shell model is insufficient to adequately capture array of behaviors observed in our physical tests, but we are able to reproduce the data well with the simple addition of anisotropy and an area-preservation term. Similarly, although we found the basic discrete elastic rod model insufficient for reproducing the behaviors of real-world textiles, Section 6 explains how we are able to reproduce the real-world data with the addition of two-phase models for contact and stretching model for capturing the behavior of plated yarns.

4 INPUT DATA

This section discusses the different types of data that can be gathered from real-world fabrics, and how the data can be used in practice for achieving our goals.

The first type of data we consider is the information used for fabricating each material. Examples of such data are the type(s) of

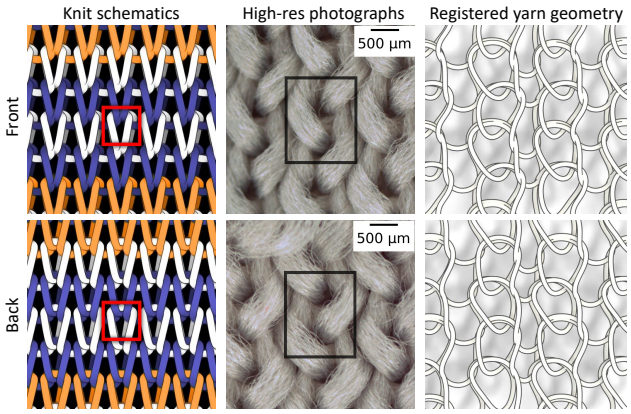

Fig. 3. Initialization of the yarn geometry for an all-needle fabric (A2). We take as input a schematic visualization that depicts the topology of the yarn construction (left) and high-res photographs that show the yarn geometry on the visible layers (middle). The squares highlight the repeat of the yarn pattern. We initialize the yarn geometry (right) by manually registering the yarn curves to the photographs, and smoothly interpolating to unobserved regions.

yarn used to create the fabric, and the knitting/weaving instructions. Another type of data is experimental measurements, such as yarn-level mechanical tests, swatch-level mechanical tests, and high-resolution images.

We could follow a number of approaches for using these data in simulation. A "bottom-up" modeling approach could take the fabrication settings and yarn-level mechanical tests, and simulate the mechanics of the knitting/weaving process that yield the final fabric pattern and its macroscopic response. A "top-down" approach, on the other hand, could take high-resolution images and swatch-level tests, and estimate both the pattern geometry and the yarn-level mechanical response. Unfortunately, both approaches suffer serious challenges. In the bottom-up approach, the fabrication process entails many unknowns, such as the forces applied by the needles or the plastic deformation of yarns, and there is no known procedure to measure the mechanical response of yarns in tight contact. In the top-down approach, the yarn pattern is not fully observable due to severe occlusion, especially in multi-layer fabrics, and the strains of individual yarns are unknown.

To work around these challenges, our project follows a hybrid approach, by combining fabrication settings and experimental measurements to design and fit yarn-level models. We use the yarn topology defined by knitting instructions, together with high-resolution photographs, to initialize the geometry of the yarn pattern. We also use swatch-level mechanical tests to estimate the mechanical parameters of the yarn-level model.

Section 4.1 describes the composition of the yarns used in our experiments; Section 4.2 describes the initialization of yarn geometry; and Section 4.3 describes the mechanical tests. Later in Section 9 we discuss additional yarn data that we considered but did not use in our system.

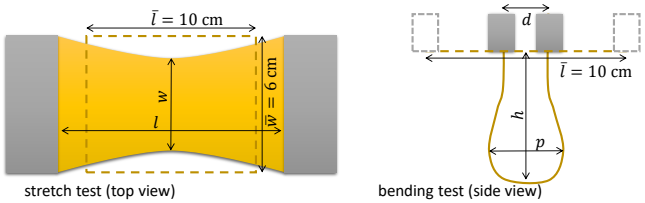

Fig. 4. Testing principles and fabric measurements for swatch-level stretch and bending tests. We have performed these tests along three directions (weft, bias, warp) to all 33 fabrics in our library.

4.1 Yarn Composition

For each fabric swatch in our data set, we list the associated yarn types, the number of filaments, the denier of each yarn type (i.e., mass in grams per 9,000 meters of unstretched yarn), the stiffness of each yarn type (measured in response to a force of 5cN), and the mass percentage of each yarn type. For example, 72D/75F polyester 90%, 20D spandex 10%, indicating that 90% of the fabric mass corresponds to a 75-filament 72-denier polyester, and 10% of the mass corresponds to 20-denier spandex. This information allows us to assign basic parameters like the density of each simulated yarn, and we have the option to derive additional information, as discussed in Section 9.

4.2 Initialization of Yarn Geometry

To define the yarn topology from knitting instructions, we leverage existing tools by knitting machine providers. Specifically, we have used the M1plus by STOLL, which outputs a schematic representation of the knit pattern (Figure 3-left). Note that this schematic defines topology, but it does not accurately define where yarn contacts occur. We also use this data to identify a representative repeating periodic tile of the pattern.

We obtain front and back high-resolution photographs of the knit pattern using a custom-built optical system, which captures photographs of 4912×3684 pixels at a resolution of 1.8 μm per pixel, with close-to-uniform illumination (Figure 3-middle). We first use these images to estimate a yarn radius, which we use for visualization and to help with initial yarn registration. Next, we identify the corresponding periodic tile on the photographs, and manually register the yarn topology to the visible layer on both front and back photographs. Manual registration took 1 to 10 hours per fabric, depending on pattern complexity. This step assigns 2D coordinates of the portions of the yarns visible in the photographs, so we still need to approximate the geometry of the occluded yarns, as well as the depth coordinates. We approximate this missing geometry by constraining the centerlines of yarns that overlap in the images to be two yarn radii apart in depth, and by smoothly interpolating the yarn geometry in unobserved regions. We include these manually registered 3D yarn geometries in our data set, and we use them as input to our yarn model parameter estimation (which is free to optimize the initial guess further, as in Section 7.2).

4.3 Physical Tests

To fit accurate yarn-level parameters, we seek physical test data that captures the force-deformation response of the fabric under stretch and bending in different directions, accounting also for nonlinearities and transverse behavior. We have developed a custom-built test rig capable of executing two state-of-the-art experiments: a clamped stretch test [Kawabata 1980] and a pear-loop bending test [Peirce 1930]. The principles of the tests and the size of the fabric swatches are shown in Figure 4. We chose the pear-loop test vs. the cantilever bending test [Wang et al. 2011] because in our early experiments it showed better sensitivity at low bending stiffness.

In the stretch test, a swatch of length $\bar{l} = 10$ cm (excluding the clamped region) and width $\bar{w} = 6$ cm is clamped on two ends and stretched horizontally under known forces. In practice, the stretch test is controlled by displacement, while force is measured. For each stretch force f_s, we compute the stretch $s = l/\bar{l} - 1$ (with l the deformed length) and the orthogonal compression $c = 1 - w/\bar{w}$ (with w the deformed width). In the bending test, the same swatch is bent by bringing the clamps together. For each inter-clamp distance d, we measure the aspect ratio $r = h/p$ of the pear-shaped loop (with h the height of the loop and p its width).

One of the important decisions for the estimation of yarn-level models is the definition of a working range, which in turn affects the range of the physical test data. In garment design, fabric stretch is quantified for the warp direction, which is typically the stiff direction of the fabric, and is usually vertically aligned with the subject's body direction along the torso and limbs. Weft is often too compliant to be aligned with the vertical direction, as clothes would hang loose; instead it allows a comfortable fit along the body's circumference and flexibility for (un)dressing. Fashion ergonomics studies suggest a comfort stretch (i.e., stretch necessary in the warp direction for casual wear) of 5 to 30%, and a power stretch (i.e., stretch necessary for active wear) of 30 to 50% [Lyle 1977; Wang et al. 2008]. For the fabrics in our library, we have observed that the average warp stretch at 2 N stretch force is 31%, and at 5 N it is 52%. This suggests that 2 N and 5 N are rough upper bounds for comfort stretch and power stretch, respectively. Therefore, we have executed stretch tests up to 5 N. We have done this, together with bending tests, on three directions: weft (0 deg), bias (45 deg) and warp (90 deg). In the weft direction we often fall short of 5 N, as we reach the rig's stretch limit (160%).

In this project, we have not addressed the hysteresis of force-deformation tests, i.e., the force difference between loading and unloading regime, produced by inter-yarn friction [Miguel et al. 2013]. We leave this phenomenon as future work, which requires estimating a model of the inter-yarn friction forces. In the physical tests, we only consider stretch under loading conditions, from rest to 5 N.

4.4 Summary of Data

To summarize, our data set consists of the following information for each set of 33 different fabric samples:

- Physical characteristics of the yarn used in each fabric, including the material used, its density, its stiffness, and information about any special coating on the fibers
- Schematics describing the topology of each knitted pattern
- High resolution photographs of both sides of the fabric
- Initial yarn geometry and topology derived from knitting instructions and photographs
- Measurements resulting from mechanical tests of stretching and bending properties of the fabric

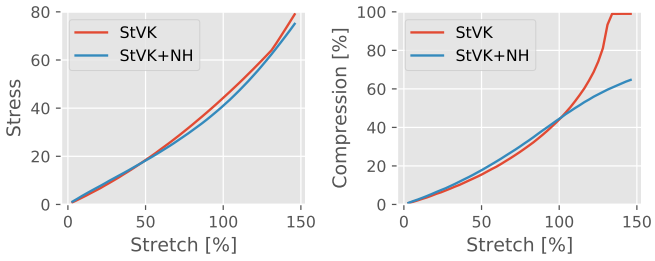

Fig. 5. Our thin-shell formulation augments the anisotropic StVK model with a Neo-Hookean area-preservation term [Smith et al. 2018]. This term does not affect the stress-stretch behavior (left), but it eliminates inversion problems (right). Both models, with and without the Neo-Hookean term, are fitted to a double-knit interlock fabric (DKIN1). Without the Neo-Hookean term, the fitted model is stable but suffers inversion (i.e., it reaches 100% compression).

5 INTERMEDIATE THIN-SHELL MODEL

As outlined in Section 3, we use a thin-shell model as an intermediary between mechanical tests on full-scale swatches (which can exhibit spatially non-uniform deformations) and periodic yarn-level simulations (which assume spatially uniform deformation). To solve for the material parameters in the thin shell model, we focus our efforts on in-plane scenarios where bending plays no role (Section 5.1). Additionally, although it is not the main result of our work, we offer a preliminary method for solving for bending parameters in Section 5.2.

5.1 In-Plane Deformation Model

We seek a deformation model that captures the anisotropic and nonlinear behavior of knitted fabrics, while minimizing the number of parameters. Note that the input data (Section 4.3) approximates uniaxial deformations, and lacks information about the fabric's response to biaxial deformations. Adding high-order strain dependency to the parameters, as done in previous works [Miguel et al. 2012; Wang et al. 2011], could over-parameterize the model and lead to overfitting. While limited, we opt for the robustness provided by a strain-independent parameterization.

A simple choice to represent both anisotropy and nonlinearity is the anisotropic Saint Venant-Kirchhoff (StVK) model [Volino et al. 2009]. However, we have observed that anisotropic StVK might exhibit a high directional Poisson's ratio (up to 2 in some cases), and inversion within the tested stretch range. Note that this does not make the model unstable, but it is of course unrealistic. To address this, we augment the anisotropic StVK model with a Neo-Hookean area-preservation term [Bonet and Wood 2008; Smith et al. 2018]. As shown in the plots in Figure 5 for a uniaxial stretch deformation with zero orthogonal stress, the stretch-aligned stress

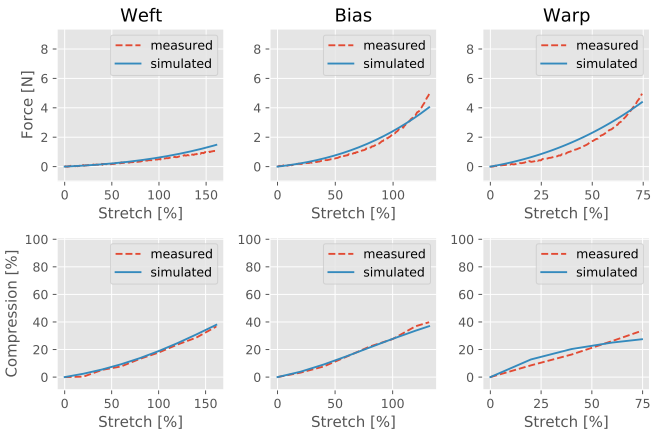

Fig. 6. Fitting of the in-plane thin-shell model to the physical test data, for an all-needle fabric (A2). Top: force vs. stretch fits. Bottom: orthogonal compression vs. stretch fits. Notice the extreme anisotropy of the fabric.

of the Neo-Hookean-augmented model remains the same as the original anisotropic StVK model, but inversion no longer occurs.

Formally, the Neo-Hookean-augmented anisotropic StVK model is formulated as follows. Given the deformation gradient $\mathbf{F}$, Green strain $\mathbf{E} = \frac{1}{2}\left(\mathbf{F}^T\mathbf{F} - \mathbf{I}\right)$ written in Voigt notation as $\boldsymbol{\varepsilon} = (E_{11}, E_{22}, 2E_{12})$, and area ratio $J = \det(\mathbf{F})$, the strain energy density is:

$$\Psi_{\text{inplane}} = \frac{1}{2}\boldsymbol{\varepsilon}^T \begin{pmatrix} k_{xx} & k_{xy} & 0 \\ k_{xy} & k_{yy} & 0 \\ 0 & 0 & k_{ss} \end{pmatrix} \boldsymbol{\varepsilon} + k_{n1}\,(J-1)^2 + k_{n2}\log^2 J. \tag{1}$$

To estimate the parameters ($k_{\text{inplane}} = \{k_{xx}, k_{xy}, k_{yy}, k_{ss}, k_{n1}, k_{n2}\}$ in (1)), we follow a simulation-in-the-loop optimization strategy [Miguel et al. 2012]. We search for the parameters that produce the best match to the stretch deformation data described in Section 4.3, subject to static equilibrium of the simulated cloth swatch. Formally, this is

$$k_{\text{inplane}} = \arg\min \sum_i w_s \left\| f_s(k_{\text{inplane}}, s_i) - f_{s,i} \right\|^2 + w_c \left\| c(k_{\text{inplane}}, s_i) - c_i \right\|^2. \tag{2}$$

Specifically, we use 6 target deformations for each weft, bias, and warp direction, distributed evenly along the stretch range. For each i target deformation, we apply the measured stretch s_i, simulate the fabric swatch to equilibrium, and evaluate the error with respect to measured stretch force $f_{s,i}$ and orthogonal compression c_i. The weights w_s and w_c normalize stretch force and compression error using the average measured values.

Figure 6 shows a representative fit of the in-plane thin-shell model for an all-needle fabric (A2). Notice the extreme anisotropy of the fabric. Results are discussed in more detail in Section 8, but the proposed model provides an accurate overall fit across all fabrics (avg. 17.59% ± 8.33% error for stretch force, and avg. 16.84% ± 8.11% error for orthogonal compression). We have noticed that the fabrics in our data set exhibit higher nonlinearity than the StVK model, but

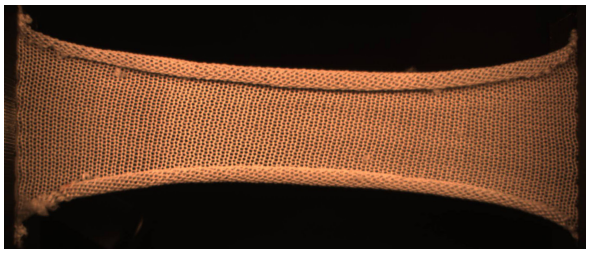

Fig. 7. This single-jersey fabric (SJ14) has a tendency to strongly curl out of plane, even during in-plane stretching tests. Such behaviors pose difficulties to the accurate measuring and modeling of bending properties.

the fit quality is sufficient to act as intermediate representation for the yarn-level model.

We use this thin shell model to generate target data for our yarn-level model by reproducing spatially uniform versions of the mechanical test scenarios in our data-set. Notably, the stretch tests clamp two ends of the fabric swatch but leave the other two sides free, which leads to a minimization of stress in the direction orthogonal to the stretch. To reproduce this behavior, we compute uniaxial stretch deformations with zero orthogonal stress of the strain energy density (1). Denoting the known applied stretch as s and the unknown orthogonal compression as c, the stretching direction can be defined by a unit vector $\mathbf{u}$ and the compression direction by an orthogonal vector $\mathbf{v}$, or alternatively by a rotation matrix $\mathbf{U} = (\mathbf{u}\ \mathbf{v})$. The resulting deformation gradient is $\mathbf{F} = \mathbf{U}\,\text{diag}(1+s, 1-c)\,\mathbf{U}^T = (1+s)\,\mathbf{u}\,\mathbf{u}^T + (1-c)\,\mathbf{v}\,\mathbf{v}^T$. We then compute the orthogonal compression as:

$$c = \arg\min \Psi_{\text{inplane}}(\mathbf{F}(s, c, \mathbf{U})). \tag{3}$$

Once the minimum-energy compression is known, we evaluate the stretch stress $\sigma_s = \frac{\partial \Psi_{\text{inplane}}}{\partial s} = \mathbf{u}^T \frac{\partial \Psi_{\text{inplane}}}{\partial \mathbf{F}}\mathbf{u}$, with $\frac{\partial \Psi_{\text{inplane}}}{\partial \mathbf{F}}$ the first Piola-Kirchhoff stress. We generate analytical stretch data $\{s, c, \sigma_s\}$ for weft, bias, and warp directions for each fabric, using the stretch range measured on the real fabric along each direction.

Because this thin shell model acts as a translator from the potentially non-uniform data to the perfectly uniform periodic yarn simulator, we should avoid encoding additional noise from numerical errors into the thin shell results. To verify its accuracy, we recomputed our results on meshes that were uniformly subdivided two times and found the average difference in output between the original and refined meshes to be only 2%.

5.2 Bending Model

The fabrics in our database exhibit a wide range of complex behaviors when subject to our physical bending test; for example, some fabrics curl out of plane when stretched (Figure 7) or break symmetry during bending tests. These complications make it challenging to isolate simple bending relationships from our data and to accurately reproduce these results in simulation. Nevertheless, we document here our first efforts toward fitting the bending behavior of the materials in our data set.

We first note that, although the bending behavior of our fabrics appears more isotropic than the in-plane behavior, some fabrics (like

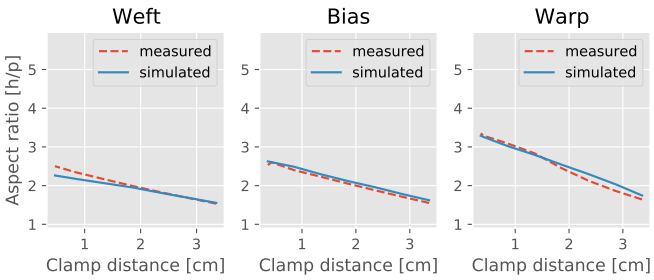

Fig. 8. Fitting of the thin-shell bending model to the physical test data, for a links fabric (L1). The plots show the aspect ratio of the pear loop vs. inter-clamp distance.

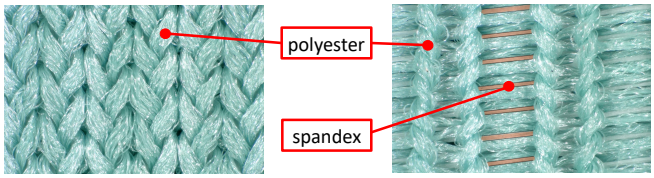

Fig. 9. High-res photographs of a plated double-knit interlock fabric (DKIN8) at 20% warp stretch (left) and 150% weft stretch (right). A multi-filament stiff yarn, polyester, provides texture and stiff response under high forces. A single-filament flexible yarn, spandex (partially highlighted), provides flexible response under low forces. The flexible yarn is stretched during knitting, and then it compresses the stiff yarn as it retracts and relaxes into the stitch structure.

the links fabric in Figure 8) exhibit noticeable bending anisotropy. To ensure our model generalizes to these scenarios, we fit these behaviors with a discrete-shell bending model with anisotropic stiffness.

In the computer graphics literature, there are multiple choices for discrete curvature models [Grinspun et al. 2003; Wardetzky et al. 2007]. We opt for an edge-based curvature metric, as this allows simple parameterization of anisotropy using the rest-shape orientation of mesh edges. Given an edge with bend angle θ, and incident triangles with average altitude from base to vertex H, we define the edge curvature as $\kappa = 3\theta/H$. This curvature metric converges to the mean curvature of a cylinder with edges aligned with the cylinder axis. This is particularly important for estimating the yarn-level model using analytical deformation data, as periodic yarn deformations will be designed following cylindrical bending (see Section 7.1).

Based on the curvature metric above, the bending energy density is:

$$\Psi_{\text{bending}} = k_\theta \, \kappa^2. \tag{4}$$

We multiply this energy density by the area of the incident triangles to obtain the discrete edge energy. We define bending stiffness values k_θ for the weft, bias, and warp directions, and interpolate linearly between them. The bending parameters are then $k_{\text{bending}} = \{k_{\theta,\text{weft}}, k_{\theta,\text{bias}}, k_{\theta,\text{warp}}\}$.

Similar to the estimation of in-plane model parameters, we use a simulation-in-the-loop optimization approach to estimate the bending parameters. We search for the parameters that produce the best match to the bending deformation data described in Section 4.3, subject to static equilibrium of the simulated cloth swatch. Formally, this is

$$k_{\text{bending}} = \arg\min \sum_i \| r(k_{\text{bending}}, d_i) - r_i \|^2.$$

Specifically, we use 7 target deformations for each weft, bias, and warp direction, distributed evenly between inter-clamp distances of 0.3 and 3.4 cm. For each i target deformation, we impose the inter-clamp distance d_i, simulate the fabric swatch to equilibrium, and measure the aspect ratio r of the bending loop.

Figure 8 shows a representative fit of the bending thin-shell model for a links fabric (L1). Again, results are discussed in more detail in Section 8, but the proposed model provides an accurate overall fit across all fabrics (avg. 5.61% ± 2.21% error).

To generate analytical target data for yarn-level estimation, we simply evaluate the bending stress $\sigma_\kappa = \frac{\partial \Psi_{\text{bending}}}{\partial \kappa} = 2k_\theta \, \kappa$. We obtain data $\{\kappa, \sigma_\kappa\}$ for weft, bias, and warp directions for each fabric, using the curvature range observed on the real fabric along each direction.

6 YARN MODEL PARAMETERIZATION

Our next goal is to find the parameters of a periodic yarn-level simulation so that it reproduces the same response to deformation as the thin-shell model described in the previous section (and thus, the fabric-level deformation tests in our data set).

We seek a yarn model with a minimal number of parameters, and which inherently captures the complexity of the coarse-scale behavior. In this regard, we start with rod models used previously for yarn-level cloth simulation [Bergou et al. 2008; Kaldor et al. 2008], and we introduce the minimal extensions necessary to capture the behavior of real fabrics with potential complications like plated yarns made of multiple materials. We describe in turn the models we use for yarn *stretch*, *bending*, and *contact*. We observe in our use case that *twist* forces are small and do not affect the overall mechanical response. Our current model does not account for inter-yarn friction, and we leave this to future work.

6.1 Stretch

Many fabrics blend yarns of different types to achieve complex mechanical and/or aesthetic behavior. One common example in our fabric library is plating, where two or more yarns are knitted side by side. Figure 9 shows a plated fabric consisting of both a flexible spandex yarn and a stiff polyester yarn. At rest, the flexible yarn is stretched, while the stiff yarn is compressed. As a result, the fabric is flexible under low forces, and then it turns stiff under high forces, once the stiffer polyester yarn is stretched. The complex interplay of plated yarns cannot be captured by modeling each yarn type separately.

Motivated by this complex stretch behavior of multi-yarn fabrics, we have designed a yarn stretch model that represents the combined response of multiple yarns. The force profile includes three linear regimes: one for low stretch with stiffness k_{s1}; another one for stretch larger than $\bar{\epsilon}_s$ with stiffness k_{s2}; and the compression regime again with stiffness k_{s2}. The

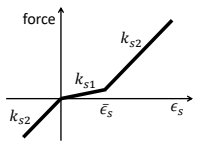

inset shows the force profile as a function of yarn stretch ε_s. Roughly, k_{s2} represents the stiffness of the stiffer yarn in a plated fabric, and $\bar{\varepsilon}_s$ the onset of stretch for this yarn. k_{s1} represents the combined response at low stretch. The stretch energy of a yarn segment with rest length $\bar{L}$ is formally:

$$W_s = \begin{cases} \frac{1}{2}\,\bar{L}\,k_{s2}\,{\varepsilon_s}^2 & \varepsilon_s \le 0 \\ \frac{1}{2}\,\bar{L}\,k_{s1}\,{\varepsilon_s}^2 & 0 \le \varepsilon_s \le \bar{\varepsilon}_s \\ \frac{1}{2}\,\bar{L}\left(k_{s2}\,(\varepsilon_s - \bar{\varepsilon}_s)^2 + k_{s1}\,\bar{\varepsilon}_s\,(2\,\varepsilon_s - \bar{\varepsilon}_s)\right) & \bar{\varepsilon}_s \le \varepsilon_s \end{cases} \quad (5)$$

Surprisingly, we found that this proposed nonlinear stretching behavior is important even for simulating fabrics made only of a single stiff yarn (e.g., polyester). We speculate that this could be due to uncertainty in rest lengths and/or friction state, as well as inherent stretch nonlinearity under low forces (e.g., due to filament realignment). For this reason, we use the nonlinear stretch model for all fabrics in the project.

6.2 Bending

We choose a yarn bending model following the formulation of Bergou et al. [Bergou et al. 2010, 2008], but we disconnect the bending stiffness and stretch stiffness of the yarn, making them independent model parameters. In addition, under the uncertainty about the yarn's cross-section shape, we choose an isotropic bending model in our project, and we leave the design of richer bending models to future work.

Following the discrete curvature vector $\boldsymbol{\kappa}$ proposed by Bergou et al. [2008], we define the bending energy at a yarn vertex as

$$W_b = \frac{1}{2}\,\frac{2}{\bar{L}_a + \bar{L}_b}\,k_b\,\|\boldsymbol{\kappa} - \bar{\boldsymbol{\kappa}}\|^2, \quad (6)$$

where $\bar{\boldsymbol{\kappa}}$ is the rest curvature, and $\bar{L}_a$ and $\bar{L}_b$ are the rest lengths of the incident yarn edges.

6.3 Contact

The cross-section of yarns deforms in a complex and nonlinear way when yarns are compressed into contact. For a multi-filament yarn, fibers may be loose and voluminous under low stretch, and they may realign anisotropically under combined stretch and contact. This phenomenon is even more complex for plated yarns (Figure 9) when viewed as a single composite yarn.

For these reasons, we choose to model both moderately soft contacts over a large distance and stiff contacts over a small distance. This models both the forgiving collision response of knit loops gently touching when the fabric is relaxed, as well as a strong resistance when the space within braided fibers collapses and there is no more room to compress. We model this two-phase contact force by combining two barrier potentials, one modeling softer large-radius contact, and another one modeling stiffer small-radius contact, as shown in the inset. Each barrier potential is parameterized by its radius R_i and stiffness k_{ci}, $i \in 1, 2$.

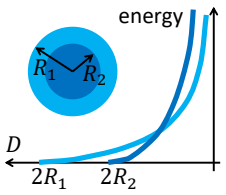

We build on the yarn-yarn contact model of Kaldor et al. [2008], but we substitute their barrier term with the one proposed by Li et al. [2020]. However, we slightly modify it to use relative distances, as this improved the scale of the stiffness for parameter optimization. Given a yarn-yarn centerline distance D, contact radius R_i, contact stiffness k_{ci}, and a barrier function $f(x) = -(x-1)^2 \log x$, the contact energy is

$$W_c = k_{ci}\,\bar{L}_a\,\bar{L}_b \int_0^1 \int_0^1 f\left(\min\left(\frac{D}{2\,R_i}, 1\right)^2\right) da\,db, \quad (7)$$

with $\bar{L}_a$ and $\bar{L}_b$ the rest lengths of the two colliding yarn edges, and the double integral over colliding edges a and b is evaluated with Simpson's rule.

7 YARN MODEL ESTIMATION

We now discuss how to estimate yarn-level parameters in order to best fit the data produced by the thin-shell model. Section 7.1 summarizes the simulation of yarn-level fabrics under periodic boundary conditions, which are key to compare to the thin-shell data.

The yarn-level model contains two types of unknown parameters: yarn rest shapes and mechanical parameters. We follow an optimization procedure that alternates the estimation of these two parameter subsets. Section 7.2 describes the estimation of yarn rest shapes, and Section 7.3 the optimization of mechanical parameters.

7.1 Periodic Yarn-Level Simulations

We build on previous work on periodic simulation of yarns and rods [Leaf et al. 2018; Schumacher et al. 2018], and we follow in particular the formulation by Sperl et al. [2020]. We refer the reader to their paper for implementation details, and here we summarize the connection between yarn-level and uniform coarse quantities.

We separate two sets of degrees of freedom on a periodic yarn tile: a uniform coarse deformation $\mathbf{q} = (s, c, \kappa, \mathbf{U})$, which gathers stretch s, orthogonal compression c, bending curvature κ, and a 2D rotation matrix $\mathbf{U}$ which defines the direction of stretch and/or bending; and a vector of nodal yarn displacements $\mathbf{u}$, expressed relative to the coarse deformation. Together, these degrees of freedom define full nodal yarn positions $\mathbf{x}(\mathbf{q}, \mathbf{u})$. To make the overall behavior of the yarn simulation match the prescribed coarse scale behavior, the nodal yarn displacements must satisfy periodic boundary conditions and must be absent of yarn sliding, yarn twist, and net rigid motion. We express all these constraints together as $\mathbf{c}(\mathbf{u}) = 0$.

The various deformation components described in Section 6 (stretch, bending, and contact) compile a set of discrete energy elements $\{W_i\}$ over a periodic tile of rest area $\bar{A}$. We compute the overall energy density of the tile as

$$\Psi_{\text{yarns}}(\mathbf{x}(\mathbf{q}, \mathbf{u})) = \frac{1}{\bar{A}} \sum_i W_i(\mathbf{x}(\mathbf{q}, \mathbf{u})). \quad (8)$$

Given a coarse deformation $\mathbf{q}$, the yarn-level deformation can be obtained as the minimum-energy configuration (i.e., equilibrium) that satisfies the constraints:

$$\mathbf{u} = \arg\min \Psi_{\text{yarns}}(\mathbf{x}(\mathbf{q}, \mathbf{u})), \quad \text{s.t. } \mathbf{c}(\mathbf{u}) = 0. \quad (9)$$

We solve this optimization using the constrained Newton approach of Sperl et al. [2020]. As each yarn tile contains only tens to hundreds of yarn nodes, the optimizations are fast in practice.

To compare yarn simulations to the thin-shell analytical target data, we need two additional ingredients. First, we need to compute stretch deformations under minimum-energy orthogonal compression. To this end, given a stretch s, we find the compression c that minimizes Ψ_{yarns} in (8), subject to equilibrium conditions (9) on the yarn deformation. We implement this optimization using the COBYLA method [Powell 1994b], solving (9) before every energy evaluation.

Second, we need to compute homogenized coarse stress, in particular the stretch stress $\sigma_s = \frac{\partial \Psi_{\text{yarns}}}{\partial s}$. Applying the chain rule to (8) for a coarse stretch s yields:

$$\frac{\mathrm{d}\Psi_{\text{yarns}}}{\mathrm{d}s} = \frac{\partial \Psi_{\text{yarns}}}{\partial \mathbf{x}} \left(\frac{\partial \mathbf{x}}{\partial s} + \frac{\partial \mathbf{x}}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial s} \right), \tag{10}$$

which simplifies to $\frac{\mathrm{d}\Psi_{\text{yarns}}}{\mathrm{d}s} = \frac{\partial \Psi_{\text{yarns}}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial s}$ under equilibrium, as the force on yarn displacements $\frac{\partial \Psi_{\text{yarns}}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{u}}$ is orthogonal to the gradient of yarn displacements $\frac{\partial \mathbf{u}}{\partial s}$. Moreover, again due to equilibrium, net yarn forces $\frac{\partial \Psi_{\text{yarns}}}{\partial \mathbf{x}}$ are zero in the interior of the tile; therefore, the computation of the coarse stress simplifies to a gathering of tile boundary forces [Schumacher et al. 2018]. Note that we compute stretching stress analytically and compute bending stress by finite differencing Ψ_{yarns}.

7.2 Yarn Rest-Shape Estimation

The rest shape of individual yarns, i.e., their rest length and rest curvature, is unknown, as it is heavily influenced by plasticity during the knitting process [Sperl et al. 2020]. One way to address this is to estimate both the yarn rest shape and mechanical parameters together, to best fit the thin-shell mechanical data. However, we have seen that the error function with respect to mechanical parameters alone is plagued with local minima, requiring the use of global optimization methods. Global optimization of both the rest shape and mechanical parameters together appears intractable; therefore, we have devised a different procedure for rest-shape estimation, motivated by the stability of the fabric's coarse-level rest state.

At the fabric's rest state, net yarn forces are zero, due to equilibrium between all force components. However, the input yarn geometry (Section 4.2) is not at rest, due to unbalanced inter-yarn contact. If we let the yarns relax to reach equilibrium under no coarse deformation, the fabric suffers a non-zero coarse stress $\frac{\partial \Psi_{\text{yarns}}}{\partial \mathbf{q}}$, i.e., the fabric's expected rest state is actually not stable and wants to deform (e.g. contract).

Based on these observations, we separate the full parameter estimation into two problems. We let yarn rest shapes ensure stability of the fabric's rest state, and we let mechanical parameters fit the coarse mechanical response. In our experience, finding a stable but slightly pre-tensed rest state (i.e., with non-zero though balanced yarn forces) was key for obtaining good fits of the mechanical response with a small mechanical parameter set.

We pose the problem of yarn rest-shape estimation as follows: We seek rest shapes such that the coarse stress at the fabric's expected rest state is small, and the equilibrium configuration of the yarns deviates little from the input yarn geometry. To solve this problem, we follow a heuristic approach, alternating equilibrium solves (9) with resetting of yarn rest shapes at the current configuration. While doing this, and to bound the deviation from the initial geometry, we bound the change in yarn rest lengths to 20%. Every time rest shapes are reset, contact forces push the yarns away. At initial steps, these contact forces may be very strong and too localized, therefore we run only a few simulation steps before resetting rest shapes. At later steps, contacts become smooth, and we let simulations run toward convergence. See Figure 10 for an example comparing the real-world fabric, the hand-registered initial yarn geometry, and the result of our rest-shape optimazation.

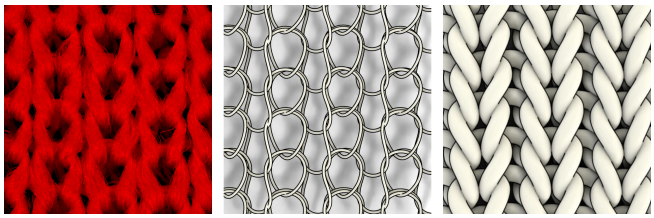

Fig. 10. High-resolution photograph (left), hand-registered yarn geometry (center), and simulated yarn rest-shapes (right) for an all-needle fabric (A1).

Yarn rest-shape estimation must be executed after every change to the mechanical parameters, as the yarn equilibrium configuration is changed. For this reason, we alternate rest-shape estimation and the mechanical parameter optimization described next.

7.3 Mechanical Parameter Estimation

To estimate the mechanical parameters of the yarn model, we pose and solve a numerical optimization problem. The remainder of this section discusses the objective function, the optimized parameters, and the solvers we use.

7.3.1 Objective function. We have designed various error metrics between the yarn model and the target thin-shell data. The stretch error component err_{stretch} measures error in stretch stress. The default error, based on the first Piola-Kirchhoff stress σ_s, ramps up at high stretches, due to the high nonlinearity of the stress function. Instead, we measure error in the second Piola-Kirchhoff stress, obtained as $\frac{\sigma_s}{1+s}$. The compression error component err_{compress} measures simply the difference in orthogonal compression. To generate the stretch and compression values, we computed zero-orthogonal-stress deformations on both the thin-shell and yarn models, sampled over the stretch range of the physical test data, for all three directions weft, bias, and warp. The yarn-level zero-orthogonal-stress computation is a slowly-varying function that is expensive to evaluate on-demand but well-approximated by simple interpolation; therefore, we compute the optimal orthogonal compression only on both ends of the stretch range, and linearly interpolate in between.

For the estimation of the yarn model, we have used the following objective function, with stretch and compression error components:

$$g = \sum_i \alpha_i \, {err_{\text{stretch},i}}^2 + \sum_j \beta_j \, {err_{\text{compress},j}}^2. \tag{11}$$

To set the weights $\{\alpha_i, \beta_j\}$ of the error components, we follow these heuristics. First, we normalize each error component based on the maximum target value in the comfort stretch range. Second,

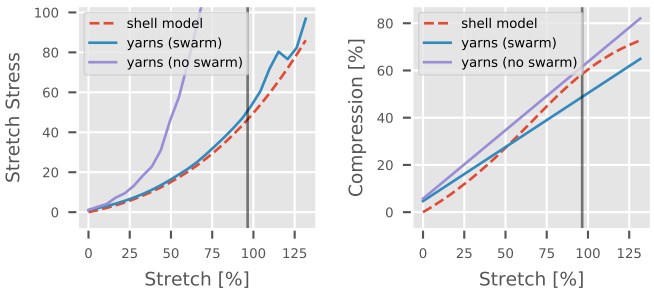

Fig. 11. The performance of our system with a local non-linear solver (purple) compared to a particle swarm optimization (blue), on an all-needle fabric (A2). Without the swarm optimization, the fitting gets stuck in a local minimum, causing the stretching error to blow up. The black bar denotes the transition from the "comfort" to the "power" range of stretches.

we apply a decaying weight $v_{\text{comfort}}/\max(v_{\text{comfort}}, v_{\text{target}})$ on both stretch and compression error components to favor high-quality fitting of the comfort stretch range vs. the power stretch range (see Section 4.3 for the definitions). v_{comfort} is the target value at the maximum comfort stretch. As such, this weight is 1 within the comfort range, and decays under further stretching.

We found it hard to robustly incorporate bending energy into the objective function without hurting the quality of in-plane fitting. Therefore, we opted to fit stretch and compression only, and use bending data for post-hoc test of the results. This strategy gave us good qualitative fitting of bending, but the design of a good error metric for quantitative fitting appeared challenging. We provide a full discussion in Section 9.1.

7.3.2 Optimized parameters. The mechanical parameters of the yarn model are (Section 6): low-stretch stiffness k_{s1}, high-stretch stiffness k_{s2}, high-stretch onset $\bar{\varepsilon}_s$, bending stiffness k_b, outer contact radius R_1, outer contact stiffness k_{c1}, inner contact radius R_2, and inner contact stiffness k_{c2}.

Three of our parameters can be set according to a general heuristic and removed from the optimization procedure. First, we set the high-stretch stiffness k_{s2} as the yarn-stretch stiffness (which is provided as input) of the stiffest yarn in the fabric. Next, the inner contact acts as a non-penetration constraint, and we found that the fabric's mechanical response is barely sensitive to its actual parameter values. Therefore, we fix the inner contact stiffness k_{c2} to 1e3, and the inner contact radius R_2 to be 25% of a base radius R_{est}. We define R_{est} geometrically as the minimum required radius such that all yarn segments are in contact in the registered geometry.

After pruning these parameters, the final set of 5 optimized parameters is $\mathbf{p} = \{k_{s1}, \bar{\varepsilon}_s, k_b, R_1, k_{c1}\}$.

We also use R_{est} to define the optimization range of the outer contact radius, which we optimize in $R_1 \in [0.5\,R_{\text{est}}, 1.5\,R_{\text{est}}]$. Similarly, we define a base bending stiffness as the default stiffness in discrete rod simulation $k_{b,\text{base}} = k_{s2}\frac{R_{\text{est}}^2}{4}$, to optimize $k_b \in [10^{-3}k_{b,\text{base}}, 10k_{b,\text{base}}]$. We further fit $k_{s1} \in [10^{-3}k_{s2}, k_{s2}]$, $\bar{\varepsilon}_s \in [0.0, 0.15]$, $k_{c1} \in [10^{-2}, 10^2]$. k_{s1}, k_b, and k_{c1} are fit as log-space parameters. Finally, we specifically allow $\bar{\varepsilon}_s$ up to 0.20 for DKIN10, to mitigate extreme stiffening under stretching.

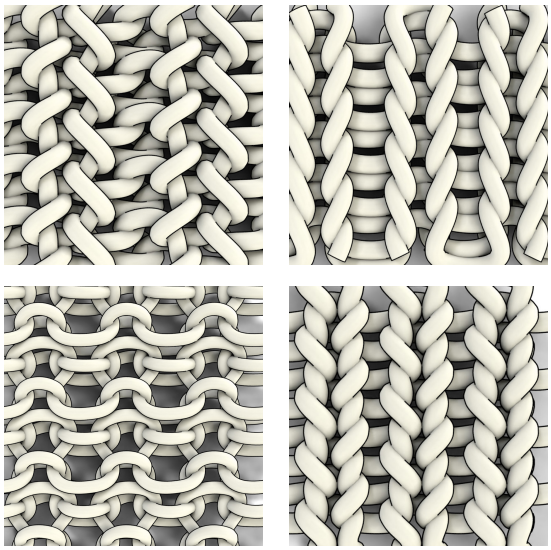

Fig. 12. Some examples of the diversity of fabrics in our data set. The fabrics are (clockwise from top-left): all-needle (A1), double-knit interlock (DKIN9), single jersey (SJ9), and links (L3). Each is stretched in the weft direction for better visibility.

7.3.3 Optimization solvers. We found that our optimization landscape features numerous local minima. Therefore, we combine global optimization for initialization and local optimization for refining of optimized material parameters. Figure 11 gives an example of how local optimization can get stuck and result in strong errors.

For the global step, we found naïve grid sampling to be infeasible. Instead, we use 10 steps of swarm optimization [Bonyadi and Michalewicz 2017] with 64 particles. Each particle corresponds to a set of candidate material parameters and will compute stress and compression error using 7 samples per direction. For performance, we compute everything as parallel as possible; first, the rest-shape heuristic (Section 7.2) per particle, then the optimal compression optimizations (Section 7.3.1) per particle and direction for interpolation, and finally the individual stress samples to evaluate the errors. Our implementation uses a decaying inertia weight from 0.9 to 0.1 for each particle with heuristic repulsion enabled.

For local optimization we use COBYLA [Powell 1994a], starting from the best parameters found by the swarm optimization. Here, we first compute the two optimal compressions (for the purpose of quick evaluation via linear interpolation as described in Section 7.3.1) and then 20 simulations per direction for error evaluation.

8 RESULTS

We collected results for 33 diverse fabrics, exhibiting a range of different knits and yarn compositions used in real-world industrial applications. The supplementary material accompanying this paper details the following fabric samples: all-needle fabrics made of polyester fiber (labeled A1 - A3); double-knit interlock fabric made with high-gauge polyester fiber (DKIN1 - DKIN7), spandex/polyester plated yarn (DKIN8), and low-gauge polyester fiber (DKIN9 - DKIN11); double-knit pique fabric made of spandex/polyester plated yarn

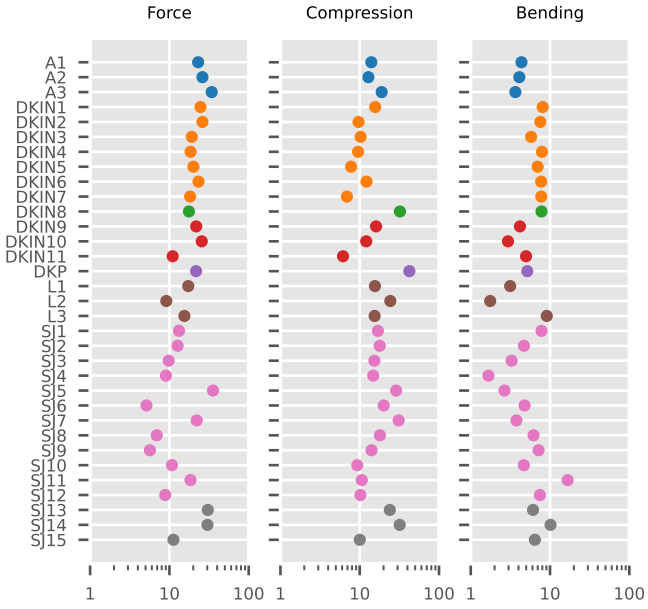

Fig. 13. Overview over our thin-shell fitting results. The x-axis is error percentage, and the y-axis lists specific types of fabric in our database. Specifically, the error is relative to the average ground-truth measured stretching, compression, or bending datum per pattern as discussed in Section 5.1.

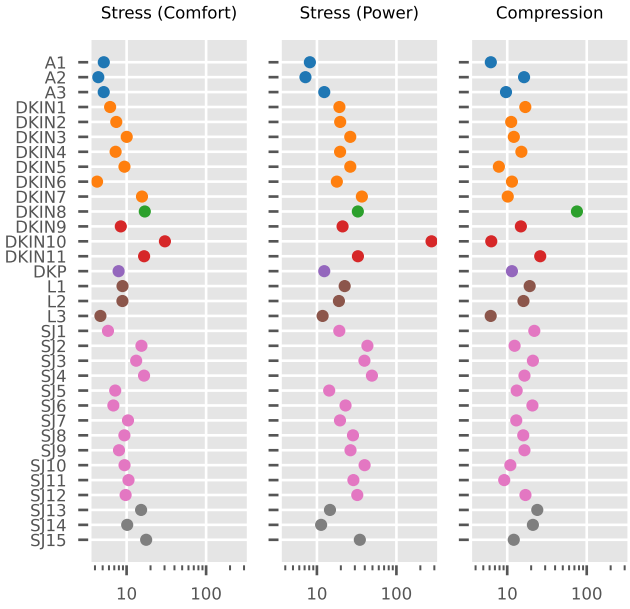

Fig. 15. Overall ability of our yarn-level solver to reproduce the corresponding real-world behaviors of materials in our database. The x-axis is error percentage, and the y-axis lists specific types of fabric in our database. The error is relative to the maximum ground-truth datum in the comfort stretch per pattern (see Section 7.3.1).

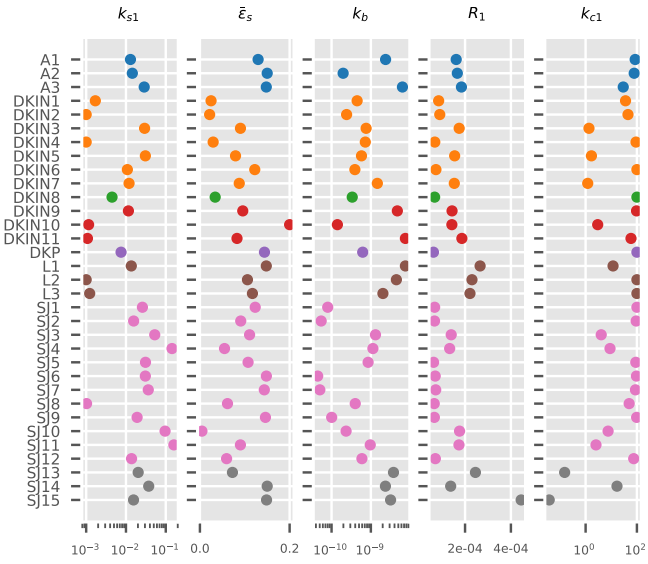

Fig. 14. Estimated parameter values of the yarn-level model for all 33 fabrics in the test database. A linear (resp. log) scale is used when the parameters are estimated in linear (resp. log) scale.

(DKP); links fabric made of spandex/polyester plated yarn (L1 - L3); and single jersey fabric made of high-gauge multi-yarn (SJ1 - SJ12) and low-gauge polyester fiber (SJ13 - SJ15). Within each family of fabrics, the samples vary in yarn composition, gauge, and fabric finish. Figure 12 displays some rendered examples. This data set offers a unique level of access to industrial quality yarns and knitted fabrics; we consider the collection and publication of this data to be an important contribution of our work.

Figure 13 shows how well our thin shell model reproduces the behavior of the fabrics in our database (which are color-coded based on knit family) after we optimize for its parameters. We display results for all 33 fabrics, measuring the percentage error relative to the average measured data in stretching force, orthogonal compression, and the pear-loop ratio for quantifying bending discussed in Section 4. We test against a wide range of fabrics with different material behaviors and find that, although some fabric families have a wider standard deviation of error (SJ vs A, for example), our shell model fits each fabric family with roughly the same magnitude of error.

The main result of our work is a system for producing a periodic yarn-level solver with the exact same topology as the original fabric, but with yarn-level parameters chosen such that the yarn-level simulation approximately matches the physical tests of the real fabric. Figure 14 shows the range of the estimated parameters for all 33 fabrics in the database. Figure 15 shows the fitting errors, indicating how well our yarn-level simulation is able to reproduce

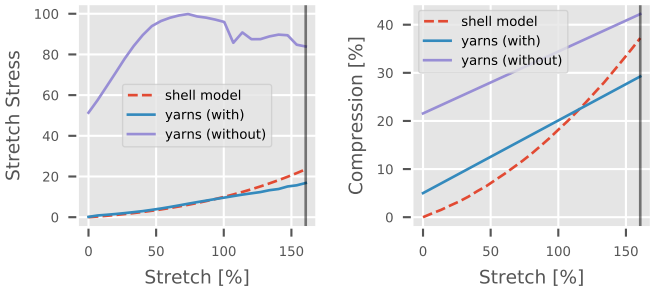

Fig. 16. The performance of our optimization with and without our bi-phasic yarn-stretching model, on an all-needle fabric (A2). Omitting the more complex yarn model massively increases the overall stretching error and adds to the compression error (purple line).

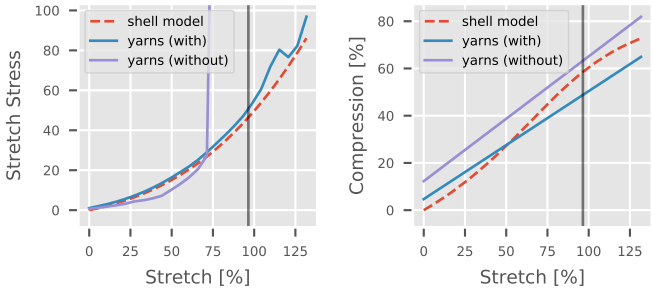

Fig. 17. The performance of our optimization with and without our rest-shape optimization, on an all-needle fabric (A2). Omitting the rest-shape optimization causes the stretching error to blow up and adds to the compression error (purple line).

the fabrics. The leftmost plot shows the ability of our method to accurately model the behavior of materials within a "comfortable" range of stretching motions considered in fashion ergonomics, as defined in Section 4.3. We plot here the percentage differences in stress, so smaller errors imply a more accurate modeling of the material. As discussed in Section 7.3, our optimization emphasizes the accuracy within this "comfort" range more heavily, due to its importance in industrial applications. The next adjacent plot shows the accuracy of our yarn-level solver over the "power stretch" range of motion expected of an athletic garment, but normalized using the maximum target value in the comfort stretch range. Although many fabrics still produce remarkably accurate results for the full range (particularly the "A1-3" fabrics), the accuracy is lower overall due to our solver's intentional bias toward accuracy within the "comfort" region. The rightmost plot in Figure 15 shows how accurately our yarn-level model matches the orthogonal compression experienced by garments within the stretch tests. We discuss our yarn simulation's ability to model fabric bending in Section 9. Overall, out of our data set of 33 fabrics, 24 of our yarn-level simulations are accurate to within 10% of the target data in the most important "comfort stretch" range of forces. The error tends to increase as we enter highly non-linear behaviors with larger stretches with very large stretches, with one optimization (DKIN10) failing to converge.

We tested our optimization (with a periodic yarn simulation in the loop) with both a local gradient-based solver and an optimizer based on particle swarms on a number of machines, the most powerful of which was a AMD EPYC 7662 server with 256 cores and 1TB of RAM; all optimizations were run simultaneously in parallel on different cores of the same machine. The local solver averaged 11m49s per fabric, with a minimum of 2m15s and maximumum of 46m11s. The swarm-based solver took about twice as long, averaging 25m19s per solve, with a minimum time of 5m18s and maximum of 1h15m25s.

Figure 16 shows the effect of our bi-phasic yarn stiffness model discussed in Section 6.2. The naive elastic rod stiffness model works reasonably well for most fabrics, but it can fail to find suitable parameters, especially in cases with plated yarns made from multiple types of fabric. Our new two-phase yarn stiffness model, in contrast succeeds to map the behavior without a blow-up in fitting error. Figure 17 illustrates the importance of the yarn rest-shape estimation discussed in Section 7.2. The black bar in these figures denotes the transition from the "comfort" to the "power" range of stretches.

The thin shell simulator obviously approximates fabric at a completely different level than the periodic yarn-based model, so we should expect some differences in fitting error between the two. The thin shell model is in a sense "closer to the data" in that we treat both the model and the real-world samples as geometric surfaces. The fabric samples are also much larger than individual stitches, and they occasionally exhibit features (like non-uniform deformation) that are not possible to model with our yarn-level solver. For these reasons, it is reasonable to expect that the thin shell model might provide a more accurate fit than the yarn level model. We also note that the thin shell model has almost twice as many degrees of freedom as our yarn-level solver (9 vs 5 DOFs), so it should also have more representation power.

9 DISCUSSION AND FUTURE WORK

9.1 Bending

Our database includes data from mechanical stretching tests as well as bending tests. Although, the main effort of our work is to accurately reproduce the stretching tests, we can also consider matching the bending data. Unfortunately, matching the bending data is challenging or a number of practical and theoretical reasons. The bending data itself is difficult to capture using the processes we proposed — by fitting the fabric to a single curve and measuring curvature information from it. This single-curve assumption breaks down when the fabrics curl dramatically (Figure 7), or when sheared fabrics asymmetrically bulge out of plane. (Note that the tendency of fabrics to curl makes it difficult to measure orthogonal compression as well.)

We started our investigation with a yarn-level model that was optimized to match stretching data for a given fabric, as explained in the previous section. We wondered whether such a model could reproduce the *bending* behavior of that fabric, even though it was not trained on its bending data. Perhaps unsurprisingly, the fits to the bending data were not nearly as precise as the fits to the stretching data, but we did find that the bending behaviors actually match fairly well *qualitatively.* Essentially, our solver correctly exhibits

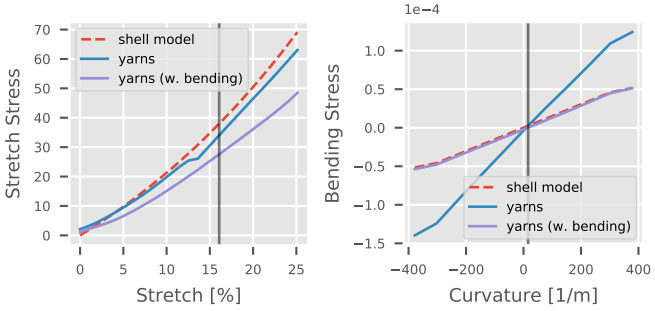

Fig. 18. The performance of our yarn-level model with (blue) and without (purple) training on bending data in the objective function, on a double-knit interlock fabric (DKIN1). Including bending data makes the stretching error worse (left) but dramatically improves the bending error (right).

weaker bending stiffness for weak fabrics, and stronger bending stiffness for strong fabrics.

To quantify bending error, we considered a bending error component err_{bend} that measures the difference in bending stiffness k_θ between the thin-shell and yarn models. On the yarn model, we compute this stiffness through central-difference approximation of the second derivative of the yarn energy Ψ_{yarns} at ±20% of the curvature range in the physical test data. We considered measuring error on stiffness, and not on stress, because the thin-shell bending model does not account for curling effects, and it assumes a flat rest state. Some of the fabrics, e.g., jersey knits, showed evident curling on the yarn model, which manifests as an offset bending stress that cannot (and probably should not) be removed by the parameter optimization.

If we actually add the bending data into the objective function (11), our yarn-level solver matches bending behavior much better, but compromises stretching quality. (Figure 18 shows a representative example, and Table 1 compiles the total effect on all fabrics.) This trade-off behavior is somewhat unsurprising given the weighted least-squares form of our objective function, but we also wonder whether the bending metric can be improved. The bending stiffnesses of the fabrics in our database span several orders of magnitude, so it may be the case that the absolute difference from the input data is the wrong metric to use in the future.

9.2 Other Yarn Data

We use yarn density and average stiffness data measured from physical tests to inform our simulated yarn-models, as discussed in Section 4.1. We also considered using yarn mass and yarn-level stretch tests for improving fidelity; by combining the yarn's physical mass with the fabric's mass density, it is possible to estimate each yarn's rest length. However, when we compared the estimated values to the actual yarn lengths measured from the yarn geometry (Section 4.2), we found that the resulting pre-stretch values would require unreasonable stretch forces for stiff yarns such as polyester.

We also obtained individual yarns of all the tested fabrics, and we performed yarn-level stretch tests using a commercial device. With this data, we are able to estimate the stretching stiffness of individual yarns. Specifically, we linearly approximate the yarn stretch response in the 5 cN range, which is an upper bound of per-yarn forces under 5 N of swatch-level force. However, we do not use this yarn stretch stiffness alone in the simulation model (only as large-stretch and compression stiffness k_{s2}) for multiple reasons: (i) the fitting procedure is sensitive to experimental noise in the low-stretch regime, yielding an unrealistically large stiffness fit in the 5 cN range; (ii) many of the tested fabrics combine multiple yarn types, and the composite stiffness is a complex combination of the stiffness of individual yarn types; and (iii) production fabrics undergo finishes that could change the mechanical response of yarns.

Table 1. Effect on various fitting errors for all fabrics if we omit bending energy ("No bending" column) or include it ("Bending" column) in the yarn-level parameter optimization.

	No bending		Bending	
Regime	avg err	std dev	avg err	std dev
Stress Comfort	10.40%	±5.27%	15.05%	±5.12%
Stress Full	31.69%	±44.51%	41.41%	±50.39%
Compression	16.28%	±11.58%	24.32%	±15.03%
Bending	97.78%	±98.26%	64.03%	±39.37%

9.3 Future Work

We presented two novel extensions to our yarn-level simulation: two-phase stiffness (for modeling plated yarns) and two-phase contact modeling. During this project, we also considered a number of other phenomenologically plausible extensions to the yarn level model. We could model the yarn's anisotropic cross section [Montazeri et al. 2019], and extend that to an anisotropic bending model. Our current model ignores friction and hysteresis in models, but we can consider this in the future. We note that our yarn simulations are well conditioned without friction because they are subject to periodic boundary conditions, but in a non-periodic simulation, a fabric full of frictionless fibers may unravel. Adding friction or even cohesion to model "fuzz" may also increase the realism of our yarn contact and fabric modeling.

As noted in Section 9.1, some fabrics have a tendency to curl. We do not yet model this curly, non-flat rest shape in our thin shell model. We could do this in the future by adding non-zero rest angles to the edges and solving for these additional degrees of freedom.

When fitting real-world materials with computer models, validation is an important task [Oberkampf and Roy 2010]. As we use the thin-shell model as an intermediate representation, the final estimated yarn-level model is not validated against the measured ground-truth data, and the final results may suffer higher error than the one reported. However, comparing full, non-periodic yarn-level simulations to the swatch-level mechanical tests comes with challenges. We would need to model the free edges of the swatches, including critical items such as yarn-yarn contact friction mentioned above. The error added by these modeling aspects would introduce high uncertainty to the validation. Romero et al. [2021] recently developed a protocol for validating rod simulation models, but it does not support the complex contact interactions of knitted yarns.

Finally, we would like to consider more complicated fabrics in the future, notably those composed of multiple layers, or those fabricated by 3D weaving [Wu et al. 2020] or warp knitting. These fabrics would challenge our geometry initialization procedure, because we would need to do more work to estimate the location of the yarns at the start of the simulation.

10 CONCLUSION

This work marks the first demonstration that yarn-level simulation is capable of approximating the mechanical stretching response of real-world fabrics. We compiled a database from physical tests of several different knitted fabrics used in the textile industry, which spans several complex knit patterns, yarn compositions, and yarn coatings, resulting in diverse physical properties like stiffness, nonlinearity, and anisotropy.

We developed a system for optimizing yarn-level parameters in order to match these real-world data, and we offer a few novel extensions to make yarn-level simulation models more capable of replicating the bi-phasic stiffness behavior of plated yarns and real-world contact scenarios. We are releasing our data set to the public research community, in hopes that it inspires future work and acts as a potential benchmark for yarn-level cloth research. In particular, we hope that future scholars use our data set and results as an inspiration for potentially finding a reliable connection between yarn-level parameters and large-scale material behavior. Finding such a connection will address a long-standing problem in material-science and directly aid in the fabrication of novel fabrics.

ACKNOWLEDGMENTS

We wish to thank the anonymous reviewers for their helpful comments. To develop this project, we were helped by many people both at Under Armour (Clay Dean, Randall Harward, Kyle Blakely, Craig Simile, Michael Seiz, Brooke Malone, Brittainy McFarland, Emilie Phan, Lindsey Kern, Courtney Oswald, Haley Barkley, Bob Chin, Adam Bayer, Connie Kwok, Marielle Newman, Nick Pence, Allison Hicks, Allison White, Candace Rubenstein, Jeremy Stangland, Fred Fagergren, Michael Mazzoleni, Nathaniel Berry, Manuel Frank) and SEDDI (Gabriel Cirio, Alejandro Rodríguez, Sofía Dominguez, Alicia Nicas, Elena Garcés, Daniel Rodríguez, David Pascual, Manuel Godoy, Sergio Suja, Sergio Ruiz, Roberto Condori, Alberto Martín, Graham Sullivan). We also thank the members of the Visual Computing Group at IST Austria and the Multimodal Simulation Lab at URJC for their feedback. This research was supported by the Scientific Service Units (SSU) of IST Austria through resources provided by Scientific Computing, and it was funded in part by the European Research Council (ERC Consolidator Grant 772738 *TouchDesign*).

REFERENCES

Miklós Bergou, Basile Audoly, Etienne Vouga, Max Wardetzky, and Eitan Grinspun. 2010. Discrete Viscous Threads. *ACM Trans. Graph.* 29, 4, Article 116 (2010), 116:1–116:10 pages.

Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. 2008. Discrete Elastic Rods. *ACM Transactions on Graphics (SIGGRAPH)* 27, 3 (2008), 63:1–63:12.

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 *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (San Diego, California) *(SCA '03)*. Eurographics Association, Goslar, DEU, 37–51.

Javier Bonet and Richard D. Wood. 2008. *Nonlinear Continuum Mechanics for Finite Element Analysis* (2 ed.). Cambridge University Press.

Mohammad Reza Bonyadi and Zbigniew Michalewicz. 2017. Particle swarm optimization for single objective continuous space problems: a review. *Evolutionary computation* 25, 1 (2017), 1–54.

Katherine L. Bouman, Bei Xiao, Peter Battaglia, and William T. Freeman. 2013. Estimating the Material Properties of Fabric from Video. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Juan J. Casafranca, Gabriel Cirio, Alejandro Rodríguez, Eder Miguel, and Miguel A. Otaduy. 2020. Mixing Yarns and Triangles in Cloth Simulation. *Computer Graphics Forum* 39, 2 (2020).

Gabriel Cirio, Jorge Lopez-Moreno, David Miraut, and Miguel A. Otaduy. 2014. Yarn-Level Simulation of Woven Cloth. *ACM Trans. on Graphics (Proc. of ACM SIGGRAPH Asia)* 33, 6 (2014).

G. Cirio, J. Lopez-Moreno, and M. A. Otaduy. 2017. Yarn-Level Cloth Simulation with Sliding Persistent Contacts. *IEEE Transactions on Visualization and Computer Graphics* 23, 2 (2017), 1152–1162.

Alexandre Derouet-Jourdan, Florence Bertails-Descoubes, Gilles Daviet, and Joëlle Thollot. 2013. Inverse dynamic hair modeling with frictional contact. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–10.

Alexandre Derouet-Jourdan, Florence Bertails-Descoubes, and Joëlle Thollot. 2010. Stable inverse dynamic curves. *ACM Transactions on Graphics (TOG)* 29, 6 (2010), 1–10.

Eitan Grinspun, Anil N. Hirani, Mathieu Desbrun, and Peter Schröder. 2003. Discrete Shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (San Diego, California) *(SCA '03)*. Eurographics Association, Goslar, DEU, 62–67.

Sunil Hadap. 2006. Oriented strands: dynamics of stiff multi-body system. In *Proceedings of the 2006 ACM SIGGRAPH/Eurographics symposium on Computer animation*. 91–100.

Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2008. Simulating Knitted Cloth at the Yarn Level. *ACM Trans. Graph.* 27, 3, Article 65 (2008), 65:1–65:9 pages.

Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2010. Efficient Yarn-based Cloth with Adaptive Contact Linearization. *ACM Transactions on Graphics* 29, 4 (July 2010), 105:1–105:10.

Levi Kapllani, Chelsea Amanatides, Genevieve Dion, Vadim Shapiro, and David E. Breen. 2021. TopoKnit: A Process-Oriented Representation for Modeling the Topology of Yarns in Weft-Knitted Textiles. *Graphical Models* 118 (2021), 101114.

S. Kawabata. 1980. *The standardization and analysis of hand evaluation.* Textile Machinery Soc. Japan.

Jonathan Leaf, Rundong Wu, Eston Schweickart, Doug L. James, and Steve Marschner. 2018. Interactive Design of Periodic Yarn-Level Cloth Patterns. *ACM Trans. Graph.* 37, 6, Article 202 (2018).

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental Potential Contact: Intersection-and Inversion-Free, Large-Deformation Dynamics. *ACM Trans. Graph.* 39, 4, Article 49 (jul 2020).

Junbang Liang, Ming C. Lin, and Vladlen Koltun. 2019. Differentiable Cloth Simulation for Inverse Problems. In *Conference on Neural Information Processing Systems (NeurIPS)*.

D. Liu, S. Koric, and A. Kontsos. 2018. Parallelized Finite Element Analysis of Knitted Textile Mechanical Behavior. *Journal of Engineering Materials and Technology* 141, 2 (12 2018).

D. S. Lyle. 1977. *Performance of Textiles.* John Wiley & Sons, New York.

E. Miguel, D. Bradley, B. Thomaszewski, B. Bickel, W. Matusik, M. A. Otaduy, and S. Marschner. 2012. Data-Driven Estimation of Cloth Simulation Models. *Comput. Graph. Forum* 31, 2pt2 (2012), 519–528.

Eder Miguel, Rasmus Tamstorf, Derek Bradley, Sara C. Schvartzman, 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 (2013).

Zahra Montazeri, Chang Xiao, Yun Fei, Changxi Zheng, and Shuang Zhao. 2019. Mechanics-aware modeling of cloth appearance. *IEEE transactions on visualization and computer graphics* 27, 1 (2019), 137–150.

Vidya Narayanan, Kui Wu, Cem Yuksel, and James McCann. 2019. Visual Knitting Machine Programming. *ACM Trans. Graph.* 38, 4, Article 63 (2019).

W Oberkampf and Christopher Roy. 2010. *Verification and Validation in Scientific Computing.* Cambridge University Press.

F. T. Peirce. 1930. The "Handle" of Cloth as a Measurable Quantity. *Journal of the Textile Institute Transactions* 21, 9 (1930), T377–T416.

Jesús Pérez, Miguel A Otaduy, and Bernhard Thomaszewski. 2017. Computational design and automated fabrication of kirchhoff-plateau surfaces. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.

Jesús Pérez, Bernhard Thomaszewski, Stelian Coros, Bernd Bickel, José A Canabal, Robert Sumner, and Miguel A Otaduy. 2015. Design and fabrication of flexible rod

meshes. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–12.

Michael JD Powell. 1994a. A direct search optimization method that models the objective and constraint functions by linear interpolation. In *Advances in optimization and numerical analysis*. Springer, 51–67.

M. J. D. Powell. 1994b. *A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation.* Springer Netherlands, Dordrecht, 51–67.

Abdullah Haroon Rasheed, Victor Romero, Florence Bertails-Descoubes, Stefanie Wuhrer, Jean-Sebastien Franco, and Arnaud Lazarus. 2020. Learning to Measure the Static Friction Coefficient in Cloth Contact. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Yannick Remion, Jean-Michel Nourrit, and Didier Gillard. 1999. Dynamic Animation Of Spline Like Objects. In *Proc. of WSCG*.

Victor Romero, Mickaël Ly, Abdullah Haroon Rasheed, Raphaël Charrondière, Arnaud Lazarus, Sébastien Neukirch, and Florence Bertails-Descoubes. 2021. Physical Validation of Simulators in Computer Graphics: A New Framework Dedicated to Slender Elastic Structures and Frictional Contact. *ACM Trans. Graph.* 40, 4, Article 66 (2021).

Tom F. H. Runia, Kirill Gavrilyuk, Cees G. M. Snoek, and Arnold W. M. Smeulders. 2020. Cloth in the Wind: A Case Study of Physical Measurement Through Simulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. 2018. Mechanical Characterization of Structured Sheet Materials. *ACM Trans. Graph.* 37, 4, Article 148 (2018).

Breannan Smith, Fernando De Goes, and Theodore Kim. 2018. Stable Neo-Hookean Flesh Simulation. *ACM Trans. Graph.* 37, 2, Article 12 (2018).

Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized Yarn-Level Cloth. *ACM Transactions on Graphics (TOG)* 39, 4 (2020).

Georg Sperl, Rahul Narain, and Chris Wojtan. 2021. Mechanics-Aware Deformation of Yarn Pattern Geometry. *ACM Transactions on Graphics (TOG)* 40, 4 (2021).

Rosa M. Sánchez-Banderas, Alejandro Rodríguez, Héctor Barreiro, and Miguel A. Otaduy. 2020. Robust Eulerian-on-Lagrangian Rods. *ACM Trans. Graph.* 39, 4, Article 59 (2020).

Pascal Volino, Nadia Magnenat-Thalmann, and Francois Faure. 2009. A Simple Approach to Nonlinear Tensile Stiffness for Accurate Cloth Simulation. *ACM Trans. Graph.* 28, 4, Article 105 (2009).

Paras Wadekar, Vignesh Perumal, Genevieve Dion, Antonios Kontsos, and David Breen. 2020. An optimized yarn-level geometric model for Finite Element Analysis of weft-knitted fabrics. *Computer Aided Geometric Design* 80 (2020), 101883.

Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi. 2011. Data-Driven Elastic Models for Cloth: Modeling and Measurement. In *ACM SIGGRAPH 2011 Papers* (Vancouver, British Columbia, Canada) *(SIGGRAPH '11)*. Association for Computing Machinery, New York, NY, USA, Article 71.

X. Wang, X. Liu, and C.Hurren Deakin. 2008. Physical and mechanical testing of textiles. In *Fabric Testing*, Jinlian Hu (Ed.). Woodhead Publishing, 90–124.

Max Wardetzky, Miklós Bergou, David Harmon, Denis Zorin, and Eitan Grinspun. 2007. Discrete Quadratic Curvature Energies. *Comput. Aided Geom. Des.* 24, 8–9 (2007), 499–518.

Kui Wu, Hannah Swan, and Cem Yuksel. 2019. Knittable Stitch Meshes. *ACM Trans. Graph.* 38, 1, Article 10 (2019).

Rundong Wu, Joy Xiaoji Zhang, Jonathan Leaf, Xinru Hua, Ante Qu, Claire Harvey, Emily Holtzman, Joy Ko, Brooks Hagan, Doug James, François Guimbretière, and Steve Marschner. 2020. Weavecraft: An Interactive Design and Simulation Tool for 3D Weaving. *ACM Trans. Graph.* 39, 6, Article 210 (2020).

Shan Yang, Junbang Liang, and Ming C. Lin. 2017. Learning-Based Cloth Material Recovery From Video. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Cem Yuksel, Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2012. Stitch Meshes for Modeling Knitted Clothing with Yarn-level Detail. *ACM Trans. Graph.* 31, 4 (2012), 37:1–37:12.

Jonas Zehnder, Stelian Coros, and Bernhard Thomaszewski. 2016. Designing structurally-sound ornamental curve networks. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–10.