

Neighbor-Aware Data-Driven Relaxation of Stitch Mesh Models for Knits

YURA HWANG, University of Utah, USA

JENNY HAN LIN, University of Utah, USA

JERRY HSU, University of Utah, USA

BENJAMIN MASTRIPOLITO, University of Utah, USA

JAMES MCCANN, Carnegie Mellon University, USA

CEM YUKSEL, University of Utah, USA

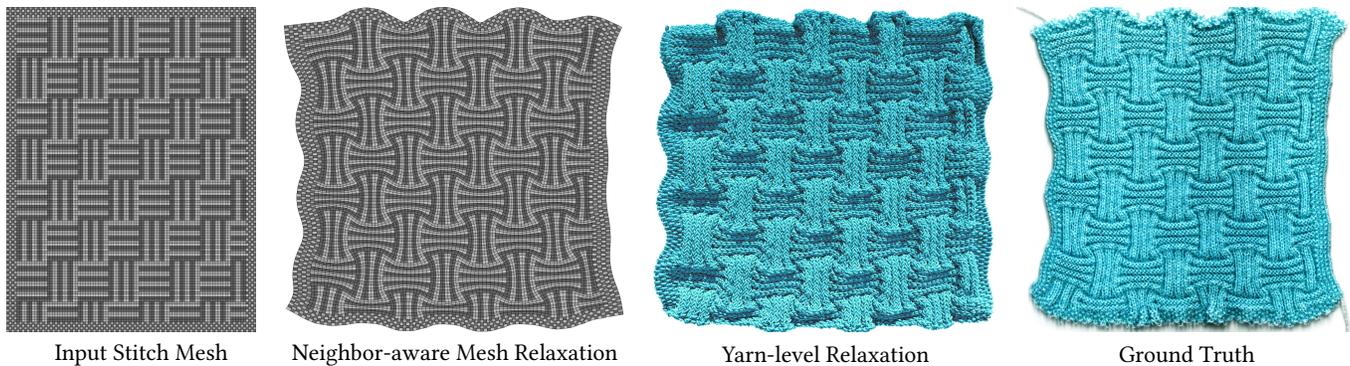


Fig. 1. Given any stitch mesh composed of knit and purl stitches, our neighbor-aware knit model can quickly approximate the rest shape of the knit object using a small set of reference measurements. Our mesh-level relaxation is useful both for quick, interactive previews and for initializing yarn-level simulations.

Lightweight, mesh-level models of knit fabric behavior are useful for both interactive pattern editing and initialization of yarn-level simulations. However, existing mesh-level simulation methods abstract knitting as a homogeneous material, which prevents them from capturing more complicated mixed structures. Furthermore, these methods require different simulation parameters depending on the knit pattern, or arrangement of stitches within the knit. Thus, fitting these parameters to physical examples must be done for each new pattern, even when the same types of stitches are used. To address this, we observe that physical behavior of a stitch is determined not only by its individual structure but also by the stitch types that surround it. In our work, we extend the stitch mesh model to allow for neighbor-aware material properties at the stitch level. Using structural analysis of stitch connections, we derive a finite set of four-way kernels that combine to create general knit-purl patterns for relaxation. From this, we generate a set of reference patterns that can be measured to infer the rest-lengths of the kernels using a linear model. After knitting and measuring these reference patterns, we used the derived kernel rest lengths to run relaxation on our stitch mesh models with mixtures of knits and purls that we then validated against physical examples. Our results show that the 4 neighbors of each stitch is sufficient

to account for much of the neighborhood-dependent deformation, while remaining simple enough to directly fit to measured data with a set of 11 basis swatches. This allows our relaxation method to efficiently estimate the rest shape of mixed knit-purl patterns, which enables fast fabric preview and more accurate yarn-level simulation.

CCS Concepts: • **Computing methodologies** → *Model verification and validation*; Physical simulation; • **Applied computing** → **Computer-aided design**.

Additional Key Words and Phrases: Knitting, Fabric Simulation, Data Fitting, Mechanics-aware, Stitch Meshes

ACM Reference Format:

Yura Hwang, Jenny Han Lin, Jerry Hsu, Benjamin Mastrapolito, James McCann, and Cem Yuksel. 2025. Neighbor-Aware Data-Driven Relaxation of Stitch Mesh Models for Knits. In *SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*, December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3757377.3763890>

1 INTRODUCTION

The world is filled with knitted objects. Not only is knitting one of the major textile fabrication techniques at both an industrial and hobbyist scale, but it has drawn interest from a variety of research fields from soft robotics [Ou et al. 2019], to biomechanics, and to biomedical engineering [Haines et al. 2016; Sydney Gladman et al. 2016; Zeng et al. 2014]. A major reason for this interest is due to the metamaterial properties of knitting; by locally varying the topology and entanglements of a few strands of yarn, it is possible to construct a wide variety of shapes and structures with differing bulk material properties. For example, Figure 2 demonstrates how a simple change

Authors' addresses: Yura Hwang, University of Utah, USA, yura.hwang@utah.edu; Jenny Han Lin, University of Utah, USA, jenny.h.lin@utah.edu; Jerry Hsu, University of Utah, USA, jerry.hsu.research@gmail.com; Benjamin Mastrapolito, University of Utah, USA, ben.mastrapolito@utah.edu; James McCann, Carnegie Mellon University, USA, jmcann@cs.cmu.edu; Cem Yuksel, University of Utah, USA, cem@cemyuksel.com.

Please use nonacm option or ACM Engage class to enable CC licenses  This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

SA Conference Papers '25, December 15–18, 2025, Hong Kong, Hong Kong

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2137-3/2025/12

<https://doi.org/10.1145/3757377.3763890>



Fig. 2. The two halves of this garter (top) and rib (bottom) fabric both consist of alternating stripes of knits and purls, where the only difference is whether the stripes run horizontally or vertically. Despite this, the two halves demonstrate very different rest configurations due to interactions between each loop and its neighbors.

in the arrangement of two stitch types drastically affects the rest dimensions of the fabric.

Due to the complexity of knit structures, computationally modeling knit materials is a challenging task. While yarn-level simulations are able to capture these material details, they remain too expensive for interactive applications. In contrast, mesh-level simulations can achieve significant speedups by working with much lower degrees of freedom. However, existing methods assume a homogeneous material model, which excludes the more complex mixed structures that appear in technical applications. A mesh-level model of knitting that can support heterogeneous knit structures is not only beneficial for interactive applications that need a fast, realistic result, but can also support yarn-level models in offline applications by providing geometry initialization and simulation constraints.

Ideally, we would like to leverage the fact that knitting is composed of low-level building blocks called stitches and measure material parameters for each individual stitch type. However, a study by Singal et al. [2024] observed that the behavior of a knit pattern depends on not just the topology of a stitch in isolation, but also the adjacent stitches it interacts with. Thus we propose a four-way kernel as the fundamental building block that can be composed to describe the overall behavior of a knit object. Our kernel includes a stitch and its neighbors stitches on the sides as well as the top and bottom. By appealing to the underlying symmetry of knitting, we show that 12 discrete kernels are sufficient for describing the space of all knit-purl patterns. Using this formulation, we knit and measure a set of 11 calibration swatches to assign rest lengths to each kernel. We then show that the measurements from this small calibration set can be used to predict the rest behavior of knit fabrics with unseen heterogeneous patterns, assuming they’re made with the same fabrication parameters. We do this by proposing a simple extension to the knitting representation known as *stitch meshes* [Yuksel et al. 2012] that allows for the kernel rest lengths to

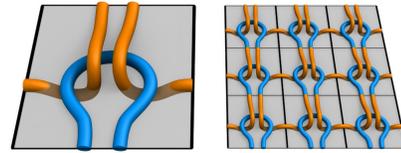


Fig. 3. A stitch contained by a quad (left) is composed to create a stitch mesh (right). Figure recreated from [Yuksel et al. 2012].

be used as material parameters when estimating the rest dimensions of heterogeneous patterns. While we use a simplified model that does not capture out-of-plane, non-linear deformation effects, our results still demonstrate high fidelity in predicting global shape. This suggests that a neighborhood aware material model can be used to efficiently predict the complex behavior of heterogeneous knits while using only a small amount of calibration data.

The main contributions in this paper are as follows:

- A set of twelve kernels that can be used to describe the pattern-level effects of knit objects.
- A linear model for using a discrete set of reference patterns to derive the rest length of these kernels.
- A data driven mesh-level relaxation method that uses these material measurements to predict the high-level features of general knit purl patterns.

2 BACKGROUND

Knitting is performed by deforming a single continuous yarn into curves known as *loops*. These loops are secured by pulling a new loop through its center. We refer to the old, existing loop as the *parent* of the new *child* loop that is pulled through it. A horizontal row of loops is connected by its constituent yarn, and called a *course*. Meanwhile, a vertical column of loops consisting of a chain of parents and children is called a *wale*. A *stitch* then refers to a variation on how a child is secured by its parent. The two most fundamental stitches are known as *knit* (K) and *purl* (P). The knit is created by pulling a child loop through its parent from back to front, while the purl stitch is formed from front to back. While in isolation a knit stitch has rotational symmetry with a purl stitch, the larger context of the fabric prevents an individual stitch from transforming between knit and purl. At the fabric level, *pattern* refers to a rectangular arrangement of knits and purls intended for periodic tiling, while a *swatch* is the larger knit fabric constructed by repeating a pattern.

If we shift our perspective from the fabric to that of a single loop, we see that it typically has four neighboring loops: two sides in its course, and a parent and child in its wale. Because of this, the knit fabric can be naturally segmented by a grid to form a quad-dominant mesh known as a *stitch mesh* Figure 3. Note that in order to capture the crossing information that differentiates knits from purls, a single face contains the curved head of one loop and the legs of its children. Thus, the physical connections between stitches cross the edges of each face.

Experiments examining the effect of stitch types and their relation to global elasticity have been done by the physics and textiles

communities. For example, Poincloux et al. [2018] create an elastic model with constraints for a sheet of all knits, where they examined the elastic response under both homogeneous and inhomogeneous deformations. Particularly relevant to this work is the study done by Singal et al. [2024]. They selected four widely used patterns that consisted of knits and purls and studied the elastic response of the fabrics at both the yarn and fabric levels. From this work, they determined that in the low-stress regime, the elastic model of the fabric could be described using a reduced symmetry model, where the different symmetries between a stitch and its neighbor affect the local elastic response. While their model only applies to the four knit patterns examined, this work shows that the concept of a stitch neighborhood can generalize to arbitrary knit-purl patterns.

3 RELATED WORK

Modeling and simulation of knits have progressed alongside the broader field of cloth simulation, which has a rich and well-established history. Much of the graphics research has focused on fast and dynamic cloth behavior, including the seminal work of [Baraff and Witkin 1998]. However, the geometric complexity of knit structures has required unique solutions in order to provide realistic simulation of their material properties. Furthermore, cloth parameter estimation in general has been a challenging problem due to the inherently complex nature of fabrics.

3.1 Data-driven parameter estimation

In order to avoid tedious tuning of simulation parameters in traditional cloth models to match physical examples, data-driven approaches have gained popularity for constructing cloth models and estimating parameters that accurately reflect the behavior of real fabrics. Within the textiles industry, the Kawabata Evaluation System [Kawabata 1980] was used to obtain empirical measurements to simulate draping behavior [Breen et al. 1994; Eberhardt et al. 1996]. These measurements aren't directly compatible with popular graphics models, which encouraged a variety of follow up work. For example, Lloyd et al. [2007] and Clyde et al. [2017] provided methods for translating common elasticity measurements into mass-spring models and Kirchoff-Love thin shells respectively. Others provided alternative calibration schema. For example, Bhat et al. [2003] estimated cloth parameters from a sequence of video frames by analyzing folds, while Miguel et al. [2012] augmented a stretching rig with a computer vision system that was capable of capturing deformation and force in 3D. There has also been interesting work in incorporating simulation results for parameter estimation. Sperl et al. [2022] used a homogenized thin-shell fabric to help estimate yarn-level material parameters from physical fabric measurements, while conversely Zhang et al. [2024] used homogenized yarn-level simulation to estimate elasticity parameters for a shell model. Though these methods can faithfully simulate the samples they measured, any new fabrics require another round of measurements, even if they use the same yarn and fabrication settings. In contrast, our method uses a small calibration set to predict the shape of general mixed knit-purl patterns.

3.2 Knit fabric simulation

The approaches for modeling knit fabric behavior can be roughly divided into yarn-level and mesh-level approaches. The early work of Rémiou et al. [2000] introduced an engine for animating general spline models to handle curve-like objects, laying the foundation for simulating yarn. Kaldor et al. [2008, 2010] pioneered garment simulation at the yarn-level with a contact tracking approach. The addition of persistent contact handling by Cirio et al. [2015, 2017] drastically improved simulation performance and efficiency, and these approaches are able to capture the high-level material properties that emerge from different knit structures. Yarn-level simulation can be further accelerated by introducing periodic constraints to enable efficient simulation of a small patch to be extended to an infinitely large fabric under tension [Leaf et al. 2018]. However, the intricate and complex topology present in non-periodic knit objects means that yarn-level simulations generally remain too expensive for interactive applications. Furthermore, it is unclear how the parameters for a yarn-level model should be fit to physical data.

As a result, modeling cloth as a variable sheet representation is a widely adopted practice. Martin et al. [2010] introduced an approach for simulating elastic materials across multiple dimensions using an integration method called *elaston*. [Schumacher et al. 2018] used a homogenization approach to model cloth as structured sheets of tiled unit cells, effectively bridging the mesoscopic and macroscopic scales. Indeed, a number of works have turned to homogenized yarn models in order to develop the appropriate material parameters specifically for knit fabrics [Sperl et al. 2020; Yuan* et al. 2024]. Mesh-level solutions can also be used in tandem with yarn-level simulation. For example, a hybrid approach can be used where yarn-level simulation is applied where more level-of-detail is needed [Casafranca et al. 2020]. Sperl et al. [2021] employs a hybrid approach that maps the geometric details and physical behavior of yarns onto a displaced mesh-based layer during simulation. However, these mesh-level models treat knitting as a single homogeneous material, as the problem of how to address boundaries between different patterns is unclear. For this reason, we use stitch meshes instead of a traditional thin sheet model.

Instead, we extend the stitch mesh for our mesh-level model. This knit data structure was first introduced by Yuksel et al. [2012], and it encodes the yarn topology into each mesh face. In addition to streamlining the task of modeling the precise yarn-level geometry of knitting for fabric simulation purposes, mesh-based relaxation was used to aid yarn-level relaxation for computing the final rest shape. Furthermore, stitch meshes have also been examined as a data structure for knit object design. Insights on general fabricability constraints [Wu et al. 2018] have been extended to enable automatic compilation of stitch meshes to both hand knitting [Wu et al. 2019] and machine knitting instructions [Liu et al. 2021; Narayanan et al. 2019; Wu et al. 2022]. In fact, Liu et al. [2021] specifically looked at the problem of generating garments with a controllable elastic response by varying the arrangement of different knit patterns while compensating the patterns' variable rest dimensions. However, instead of using the full space of available stitch patterns, they measured and fitted thirteen patterns to a thin elastic sheet model [Wang et al. 2011] and selected three patterns for use in their

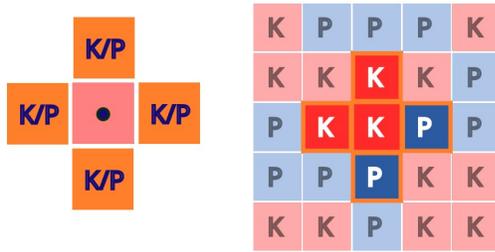


Fig. 4. Our model works on four-way kernels consisting of a stitch and its four neighbors (left). Our model views any knit/purl pattern as being made up of overlapping kernels (e.g., the 5x5 pattern on the right contains 25 kernels – one centered at each stitch).

design pipeline. Simulation and pattern assignment were then done on a separate triangle mesh before being translated to a knittable stitch mesh. Accurate shape prediction of the precise stitch mesh would be useful for both design and animation applications.

4 NEIGHBOR-AWARE STITCH MESH RELAXATION

Our model for neighbor-aware stitch mesh relaxation is built on two key observations. The first is that the smallest necessary neighborhood impacting a given stitch consists of the four directly connected stitches: two in the course direction, and two in the wale. Thus we think of a knit/purl pattern as being made up of overlapping four-way kernels centered on each stitch (Figure 4), where each kernel has a different rest width and height. The second is that within a stitch mesh, the relevant stitch connections move across the edges of a given face. Thus instead of performing relaxation directly on the mesh, we instead use the dual graph, where constraints connect the centers of faces within a kernel. We then geometrically reconstruct the stitch mesh from the relaxed dual graph by averaging the positions of dual vertices.

4.1 Kernel Size Estimation

Before we can relax our knit model we need to determine the rest size of each kernel. To do this, we assume that for a periodic pattern P with a sufficiently small repeat, there is a linear relation between the rest dimensions of the pattern w_P, h_P and the number of times each kernel appears in the pattern $C_P \in \mathbb{N}^{12}$:

$$C_P \times [w_C, h_C] = [w_P, h_P] \quad (1)$$

where $w_C, h_C \in \mathbb{R}^{12}$ are the widths and heights of each kernel. This linear assumption is valid in that the macroscopic behavior can be averaged for each repeating pattern, unless the tile involves significant out-of-plane deformations. We then use this relation to fit a model for the rest lengths of each kernel type.

One might expect there to be 32 distinct kernels (i.e., $|\{k, p\}^5|$). However, the symmetries inherent in knit structures (Figure 5) reduce this to only 12 distinct patterns (Figure 6). From there, we would ideally like to knit and measure twelve patterns that contain one kernel each. However, this is only possible for kernels C1 through C6. Because the kernels overlap with each other, assigning

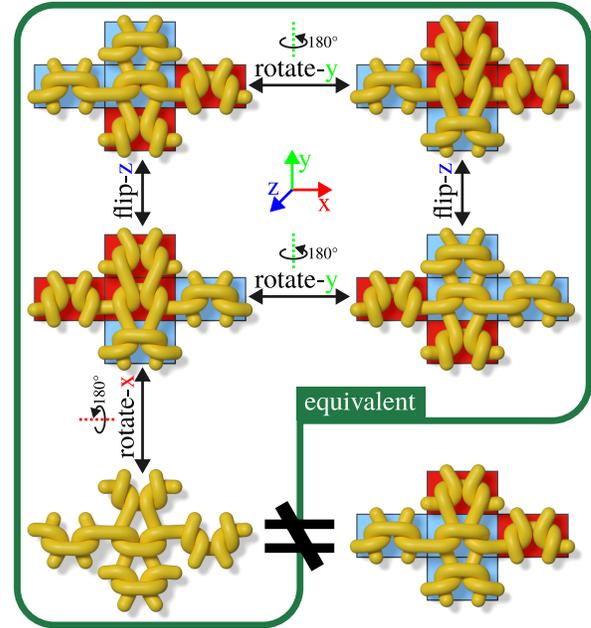


Fig. 5. Rotational and mirror symmetry yields four equivalent kernels of type C7. Note that knitting does not have rotation or reflection symmetry around the horizontal axis, so these transformations cannot be used to further reduce the space of kernels.

a kernel to a pattern not only fixes the identity of five of its stitches, but also restricts the set of valid kernels that can fit in the adjacent stitches. Indeed, upon generating every repeating patterns of size $\{2, 3, 4, 5\} \times \{2, 3, 4, 5\}$ stitches, we observed that the resulting matrix only has rank 10. This is because some kernels always co-occur, which prevents it from spanning the full 12-dimensional space of kernels. While a larger period pattern may produce a unique kernel distribution, larger pattern repeats violate the assumed linear relationship between swatch size and the number of constituent kernels. Therefore, it is not possible to construct a system that spans the full 12-dimensional space when solving for the kernel rest lengths; the best approach is to be conservative and restrict the solution to a 10-dimensional space.

On further examination, we found that patterns could be constructed to isolate kernels C1 : C6, which allow us to span the six dimensions. For the remaining six kernels we add five more patterns to our training set to span the available four dimensions containing, C7 : C12: three patterns using an equal ratio of [C7, C8], [C9, C10], and [C11, C12]; and two patterns using [C12, C9, C8, C10] and [C11, C7, C10, C9] in a 1 : 2 : 2 : 3 ratio. The full set of patterns is shown in Figure 10. We chose these patterns because they span the reachable space of kernels, they each contain relatively few kernels in a relatively even ratio, and they are “balanced” – they have an equal number of knits and purls, so they are not prone to curling. We could, in fact, do without one of these patterns, but we chose to use all of them in order to balance the measurements counts of the various kernels. (We discuss other choices of calibration patterns in Section 7.)

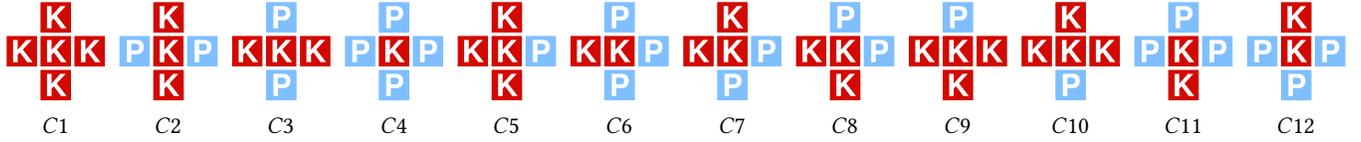


Fig. 6. The set of distinct knit-purl configurations after accounting for symmetry.

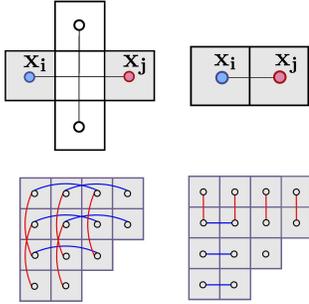


Fig. 7. The spring constraint that connects stitches consists of a kernel spring (left) and a boundary spring (right).

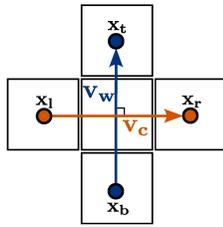


Fig. 8. The shear constraint orthogonally aligns stitches.

Given that we have more unknowns than measurements (and some redundant measurements), we regularize the kernel size estimation problem by choosing kernel sizes which (a) reconstruct the dimensions of our calibration swatches as closely as possible and (b) are as near as possible to the average kernel size in our calibration set, given this constraint. Our code computes this fit by first computing the average kernel size within the calibration swatches, then using the Penrose-Moore pseudo-inverse of the kernel count matrix (C_P^\dagger) to compute the minimum-norm step from this average size that will match the observed measurements:

$$\mu_P \equiv \frac{\text{sum}(w_P)}{\|P\|} \quad \text{average swatch width} \quad (2)$$

$$\mu_C \equiv \frac{\mu_P}{\text{width}} \quad \text{average kernel width} \quad (3)$$

$$w_C \equiv C_P^\dagger(w_P - [\mu_P]) + [\mu_C] \quad \text{fit data} \quad (4)$$

where $\|P\|$ is the size of pattern set P (or the number of patterns used for calibration), width is the number of stitches measured to generate the swatch measurement w_P , and $[s]$ constructs a vector of appropriate size with all elements equal to s . (And similarly for h_C , h_P , height.)

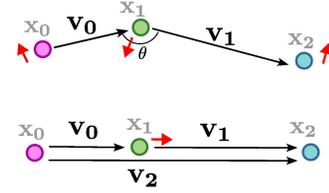


Fig. 9. The bending constraint prevents stitch inversions and in-plane bending to preserve neighboring connections.

4.2 Stitch Model

To perform mesh based relaxation using our measured data, we minimize several constraint based energies defined using the following constraints: *spring*, *shear*, and *bending*. While these constraints alone cannot capture the full complex behavior of knits, it demonstrates how kernel parameters can be combined for a general knit model. We discuss how both the measurement pipeline and model could be augmented to capture these effects in Section 7.

4.2.1 Spring constraint. For each stitch face with a complete kernel (i.e. fully enclosed by valid neighboring stitches) we use linear *kernel springs* to connect its neighboring stitches to each other; one to connect its yarn-wise neighbors in the same course, and another to connect its parent and child in the same wale. While knitting displays non-linear behavior under extreme deformation, the smaller displacements in the static case can be approximated as linear [Singal et al. 2024].

The rest length of each spring is determined by the kernel type. For faces on the boundary where one or more of the neighbor cells is missing, we connect the central face to the existing neighbor using a *boundary spring* (Figure 7). Because we do not have rest length information for partial kernels, we instead copy that information from the closest neighbor face with a full kernel. The rest length for boundary springs is then half rest length of a full kernel. The stiffness for the boundary spring is intentionally set lower to avoid an over-constrained system.

Given two face positions \mathbf{x}_i and \mathbf{x}_j connected by a spring, the spring energy E_{spring} is defined as

$$E_{spring}(\mathbf{x}_i, \mathbf{x}_j) = \frac{1}{2} k_{spring} (|\mathbf{x}_i - \mathbf{x}_j| - L)^2 \quad (5)$$

where L is the appropriate rest length, and k_{spring} is a spring constant.

4.2.2 Shear constraint. The interlocked yarn loops in a physical knit object constrains the topology and stabilizes the shape of each stitch in a grid-like manner. Thus we introduce a *shear constraint* to preserve a grid-like knit lattice structure by encouraging the course and wale neighbors of a kernel to be orthogonally aligned

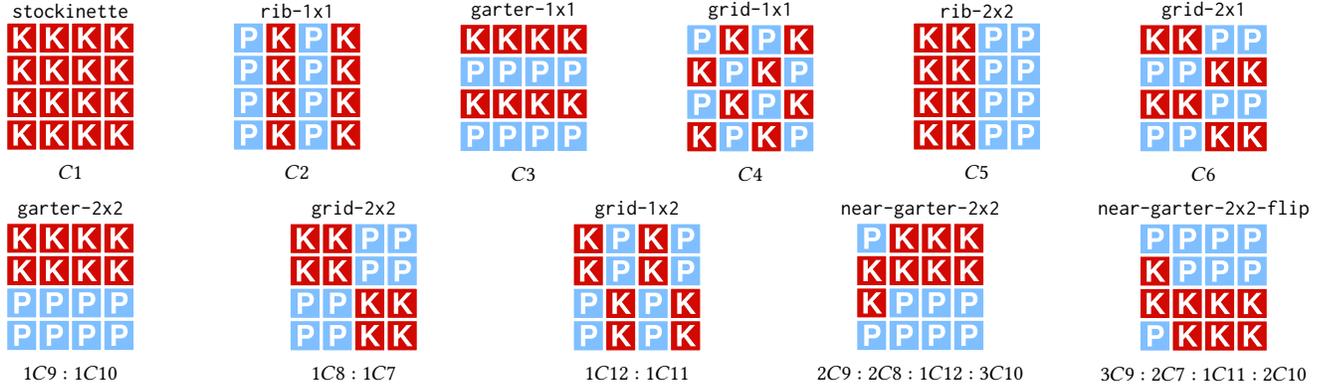


Fig. 10. The patterns used for kernel measurement (tiled over a large area) and their corresponding kernel distributions.

(Figure 8). For the four faces $\mathbf{x}_l, \mathbf{x}_r, \mathbf{x}_t, \mathbf{x}_b$ connected by a given kernel’s boundary/kernel springs, we add an additional shear constraint:

$$E_{shear}(\mathbf{v}_c, \mathbf{v}_w) = \frac{1}{2} k_{shear} (\mathbf{v}_c \cdot \mathbf{v}_w)^2 \quad (6)$$

where $\mathbf{v}_c = \mathbf{x}_r - \mathbf{x}_l$, $\mathbf{v}_w = \mathbf{x}_t - \mathbf{x}_b$, and k_{shear} is an shear stiffness. Although physical knits often exhibit shear that is not orthogonally aligned, the rest angle of each individual kernel is difficult to measure. Thus we adopt a simplified assumption to prevent the course and wale directions from collapsing to zero.

4.2.3 Bending constraint. The lattice structure of the knit grid means a given stitch maintains the same neighbors, regardless of how much they are deformed. Thus, it was crucial to ensure that two neighboring faces do not swap positions.

Furthermore, courses and wales in knit fabrics generally do not bend without external factors. Thus we also include a bending constraint that encourages linearity while also penalizing any inversion caused by in-plane bending of the course and wale directions (Figure 9).

Given three face positions $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ located sequentially along a course/wale, we define bending energy as,

$$E_{bend}(\mathbf{v}_0, \mathbf{v}_1) = \frac{1}{2} k_{bend} \left(1 - \frac{\mathbf{v}_0 \cdot \mathbf{v}_1}{|\mathbf{v}_0| |\mathbf{v}_1|}\right)^2, \quad (7)$$

where $\mathbf{v}_0 = \mathbf{x}_1 - \mathbf{x}_0$, $\mathbf{v}_1 = \mathbf{x}_2 - \mathbf{x}_1$ and k_{bend} is a bending stiffness.

We additionally apply a *sliding regularization* to the bending constraint. At a relaxed stage, the propagated tension along the swatch will distribute the stitches relatively uniformly along the knit object. We model this with a regularization that encourages the center face \mathbf{x}_1 to be equidistant between its neighbors \mathbf{x}_0 and \mathbf{x}_2 :

$$E_{slide}(\mathbf{v}_0, \mathbf{v}_1, \mathbf{v}_2) = \frac{1}{2} k_{slide} \left(\frac{|\mathbf{v}_1|^2 - |\mathbf{v}_0|^2}{|\mathbf{v}_2|^2}\right)^2, \quad (8)$$

where $\mathbf{v}_2 = \mathbf{x}_2 - \mathbf{x}_0$, and k_{slide} is a sliding stiffness.

5 KNIT KERNEL DATA COLLECTION

Having established our kernels and neighbor-aware model, we must fabricate and measure our reference patterns (Figure 10) in order to infer the necessary kernel rest dimensions. This is actually non-trivial given the flexible nature of fabrics, the impacts of boundary effects, and how small stitch feature can get.

5.1 Material Fabrication and Measurement

Each pattern was tiled to fill a 100 stitch wide by 120 stitch tall swatch, and a border consisting of 10 columns of seed stitch was added at the edges. To aid with measurement, we used multiple yarns to provide clearer landmarks. Two contrasting colors¹ of Tamm Petit yarn (2/30 acrylic yarn) were used for the body of the swatches, with the colors alternating every two rows. These contrasting colors were also used to knit a label at the top of each swatch showing the pattern used in the swatch, to aid in tracking swatches. Finally, we tasked additional yarn carriers to vertically inlay thin rayon yarns (Winning brand) as markers next to the 30th and 70th column in the pattern, to help locate these positions in the finished swatch. Carrier numbers were assigned such that the darker body yarn would pass behind the trailing marker yarns and the lighter body yarn would pass in front of them, locking the marker yarns in place and creating a clear “dashed line” to aid in measurement.

Swatches were knit on a Shima Seiki SWG91N2 15-gauge v-bed knitting machine at stitch size 35/20. Swatches were knit joined vertically into groups of 3-4, cut into individual swatches, and the label sections were mostly severed to minimize any shape distortion they might cause.

All swatches sat for at least 12 hours after coming off the machine, then were tugged sharply in the course direction twice, and in the wale direction once, then let rest for a few more minutes before scanning at 600dpi on an office multifunction scanner/copier. Scans were manually annotated by first roughly finding the center of the swatch, then counting up and down 24 rows to find a 40 (wide) x 48 (tall) patch of stitches, and placing five horizontal and five vertical lines at roughly equally-spaced locations on this patch (Figure 11).

5.2 Hysteresis Effects

Knit fabric has notable hysteresis – the size and shape of a swatch of knitting depends on how it was last handled and distorted. While we attempted to control for this in our measurement procedure, it is to be expected that some variation would be present. To characterize the extent of this hysteresis, we performed additional measurements

¹We used 4275 [Turquoise] for the darker color and 4280 [Pastel Green] or 4260 [Pastel Blue] for the lighter yarn – we ran out of 4280 during sample production.

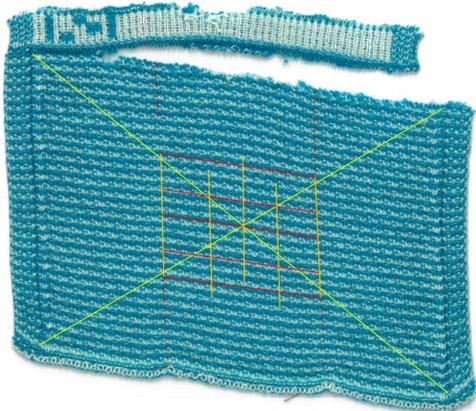


Fig. 11. Example of a marked-up swatch after measurement.

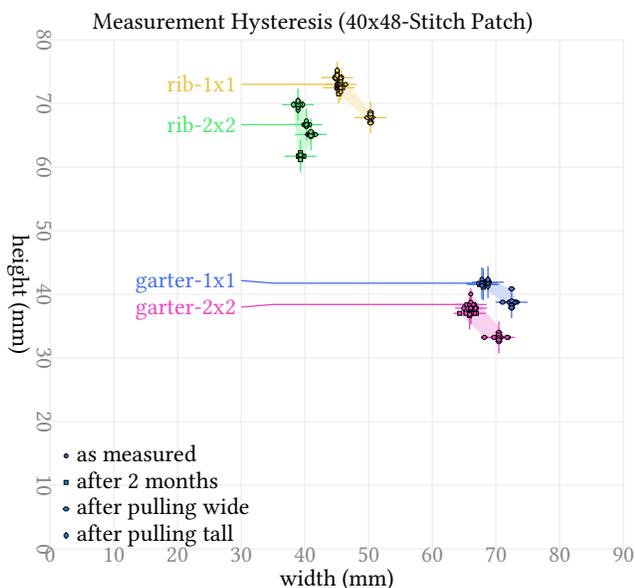


Fig. 12. To characterize the handling-dependence of knit pattern sizes, we re-measured four of our training swatches under different handling conditions. Handling seems to change the size by around 10%, suggesting that any history-independent estimation procedure would do no better.

on four of our training patterns, and present these measurements in Figure 12.

Each swatch was re-measured three more times: after 2 months of rest, immediately after a sharp tug in the course (width) direction, and immediately after a sharp tug in the wale (height) direction. Both sharp-pull measurements were also conducted after the two-month wait. From this data we conclude that handling hysteresis can account for 10% variation among average width and height measurements, with some pattern dependence; for example, *garter-2x2* has the largest variation – 15% in height – and *rib-1x1* has the smallest – 5% in width. We additionally observe that there isn’t a

simple width-height trade-off present – all swatches were smaller after being left for two months, with *rib-2x2* being particularly striking in this regard. Also, even though pulling sharply in the wale direction was the standard for initial measurements, the “as measured” and “after pulling tall” measurements do not perfectly agree, suggesting either the impact of aging the swatches or (more likely) variation in the application of a “sharp pull” in our measurement protocol.

6 RESULTS

We evaluate our relaxation method against two datasets: 33 small tiling patterns consisting of the 11 test patterns and 22 unknowns, as well as 6 heterogeneous patterns that include irregular patches with different stitch patterns as well as patterns with larger repeats. Each stitch mesh is initialized using the average kernel size and then relaxed by using *vertex block descent* (VBD) [Chen et al. 2024], implemented in C++. VBD essentially uses a Newton’s method to solve the optimization problem

$$\mathbf{x}^{t+1} = \arg \min_{\mathbf{x}} G(\mathbf{x}). \quad (9)$$

for face positions \mathbf{x} , where variational energy $G(\mathbf{x})$ involves two terms; the *inertia potential* and the *total potential energy*. As our simulation is a relaxation, we can exclude the inertial potential from $G(\mathbf{x})$, which is equivalent to minimizing the total potential energy $E(\mathbf{x})$ at \mathbf{x} .

For all simulated results, we use the $k_{spring} = 2.0$ for kernel springs, $k_{spring} = 1.0$ for boundary springs, and $k_{shear} = 1.0$ $k_{bend} = 30.0$, $k_{slide} = 2.0$. Preliminary experiments suggest that there is a relatively wide range of stiffness parameters that yield similar results. This behavior is expected, given that the bending, shear, and slide energies primarily serve as regularization to prevent arrangements of degenerate dual vertices. Note that the bending stiffness k_{bend} is comparatively large because the bending constraint defined in Section 4.2 uses normalized vectors, whereas the shear constraint uses unnormalized vectors. When we modified the angle constraint to also use normalized vectors, reducing k_{bend} by an order of magnitude produced similar results.

For the tiling patterns, we estimate the swatch measurements described in Section 5 by both calculating the linear combination of kernels present in the 40×48 swatch, and by relaxing a 100×120 stitch mesh and taking centered width and height measurements. From the results in the training set (Figure 13), we can see that the kernel fitting process worked well, since the linear reconstruction agrees closely with the measurements. As expected the patterns used to derive linearly independent kernels $C1 : C6$ are almost perfectly reproduced by the linear reconstruction, while the patterns with under constrained kernels $C7 : C12$ have a reconstruction error of at most 3.2%. These deviations are a result of us using two versions of the near-*garter-2x2* pattern (one of which was flipped vertically), which, in turn, over-constrains these kernels. The results from the full relaxation also – unsurprisingly – agree with the linear estimate – there is no heterogeneous global deformation for the simulator to resolve.

We report similar results for the set of validation patterns. The simulation once again aligns closely with the basic linear size

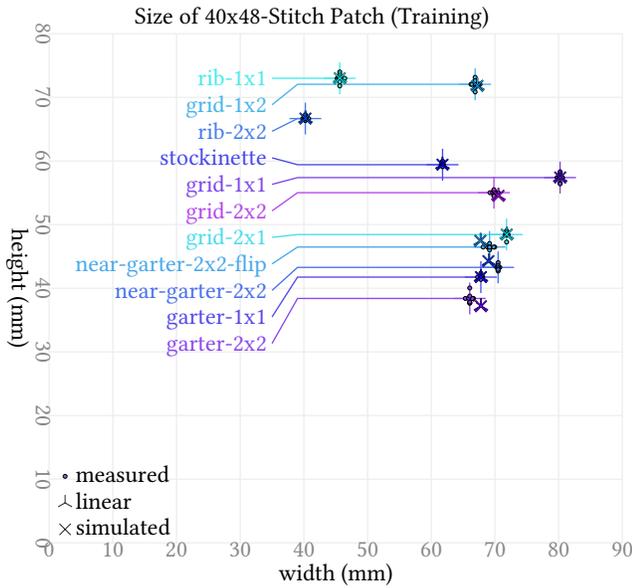


Fig. 13. Measured data for the patterns we used in fitting kernel sizes. The small points are the individual width and height measurements, the inverted Y is a linear reconstruction of the swatch size using a weighted average of the contained kernel sizes, and the X is the reconstruction with our simulator.

estimate, although the predictions now deviate further from the measured size. Despite this, this deviation is below 10% of the overall size for 22 of the swatches, which is comparable to potential hysteresis effects. The ten worst behaving examples are plotted in Figure 14, while the estimation results for the full dataset are available in the supplemental material.

The largest deviation is present in rough-rib-2x2, which is 21.62% too wide, and 1.97% too short. If we examine the ground truth pattern (Figure 15), we see that there is a ridged effect due to out-of-plane deformation. This is because knit stitches tend to curl, which is a property that is not directly captured by our model. The high error swatches involve kernels particularly impacted by this curling effect.

For the heterogeneous patterns, we perform mesh-level relaxation on the exact pattern used for the physical example. In addition, we then use the relaxed result to initialize the geometry for a yarn-level relaxation (Figure 1, Figure 16). The yarn-level relaxation was performed using [Hsu et al. 2025]. This generated an initial curve with a set radius from the stitch mesh’s yarn topology information. We then simulate this curve using Cosserat rods to resolve collisions before scaling the vertices down [Leaf et al. 2018]. As we keep the curve radius constant, the scaling has the effect of slowly tightening the yarn pattern while respecting collisions. In addition, the yarn relaxation included two planes pressing on the knit to imitate the compression of the scanner in the physical measurement. We repeat this procedure until the gaps between knits close. Note that the yarn relaxation was a separate step that does not include any forces

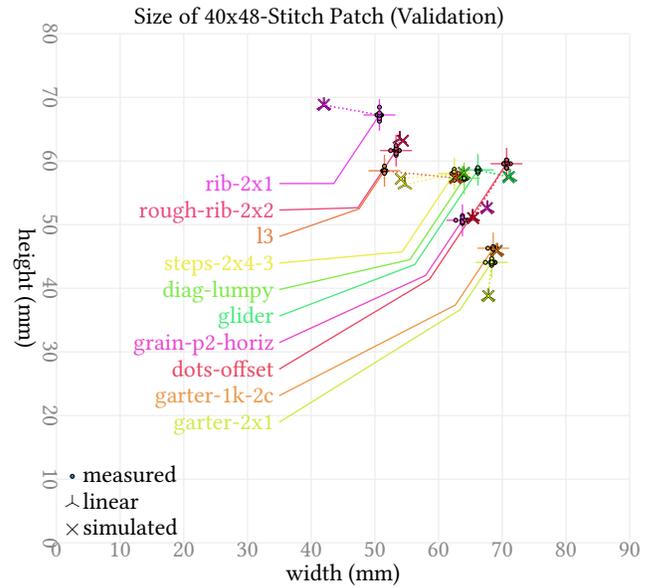


Fig. 14. Measured data for additional patterns not used in fitting kernel sizes. Deviation is generally within expected error due to hysteresis, with a few outliers reaching > 10% error. This plot includes only selected patterns for readability; a plot including every pattern we measured is available in the supplemental material.



Fig. 15. The ground truth pattern of rough-rib-2x2.

from the mesh relaxation nor are its parameters calculated from measured data.

From a qualitative perspective, we see that for the majority of results, the relaxation reproduces high-level features like the size and shape of the circles present in diagonal and rib-garter-circles, and the reflection asymmetry tiles, shown in the teaser figure. In addition, we see that when our method fails on ribs, this is again due to the presence of out-of-plane curvature, which our model does not capture. Adding in a yarn-relaxation stage enhances visual fidelity, though we do not exactly match the ground truth. This is unsurprising given that our data collection does not directly inform the yarn simulation. A more integrated approach may further improve accuracy of the results.

From a quantitative perspective, these complex patterns do not have consistent or representative reference points to report. We instead quantify error through affine mask matching. For each complex example, we extract binary masks from our original scan and the relaxed mesh. We then optimize for an affine transform on our

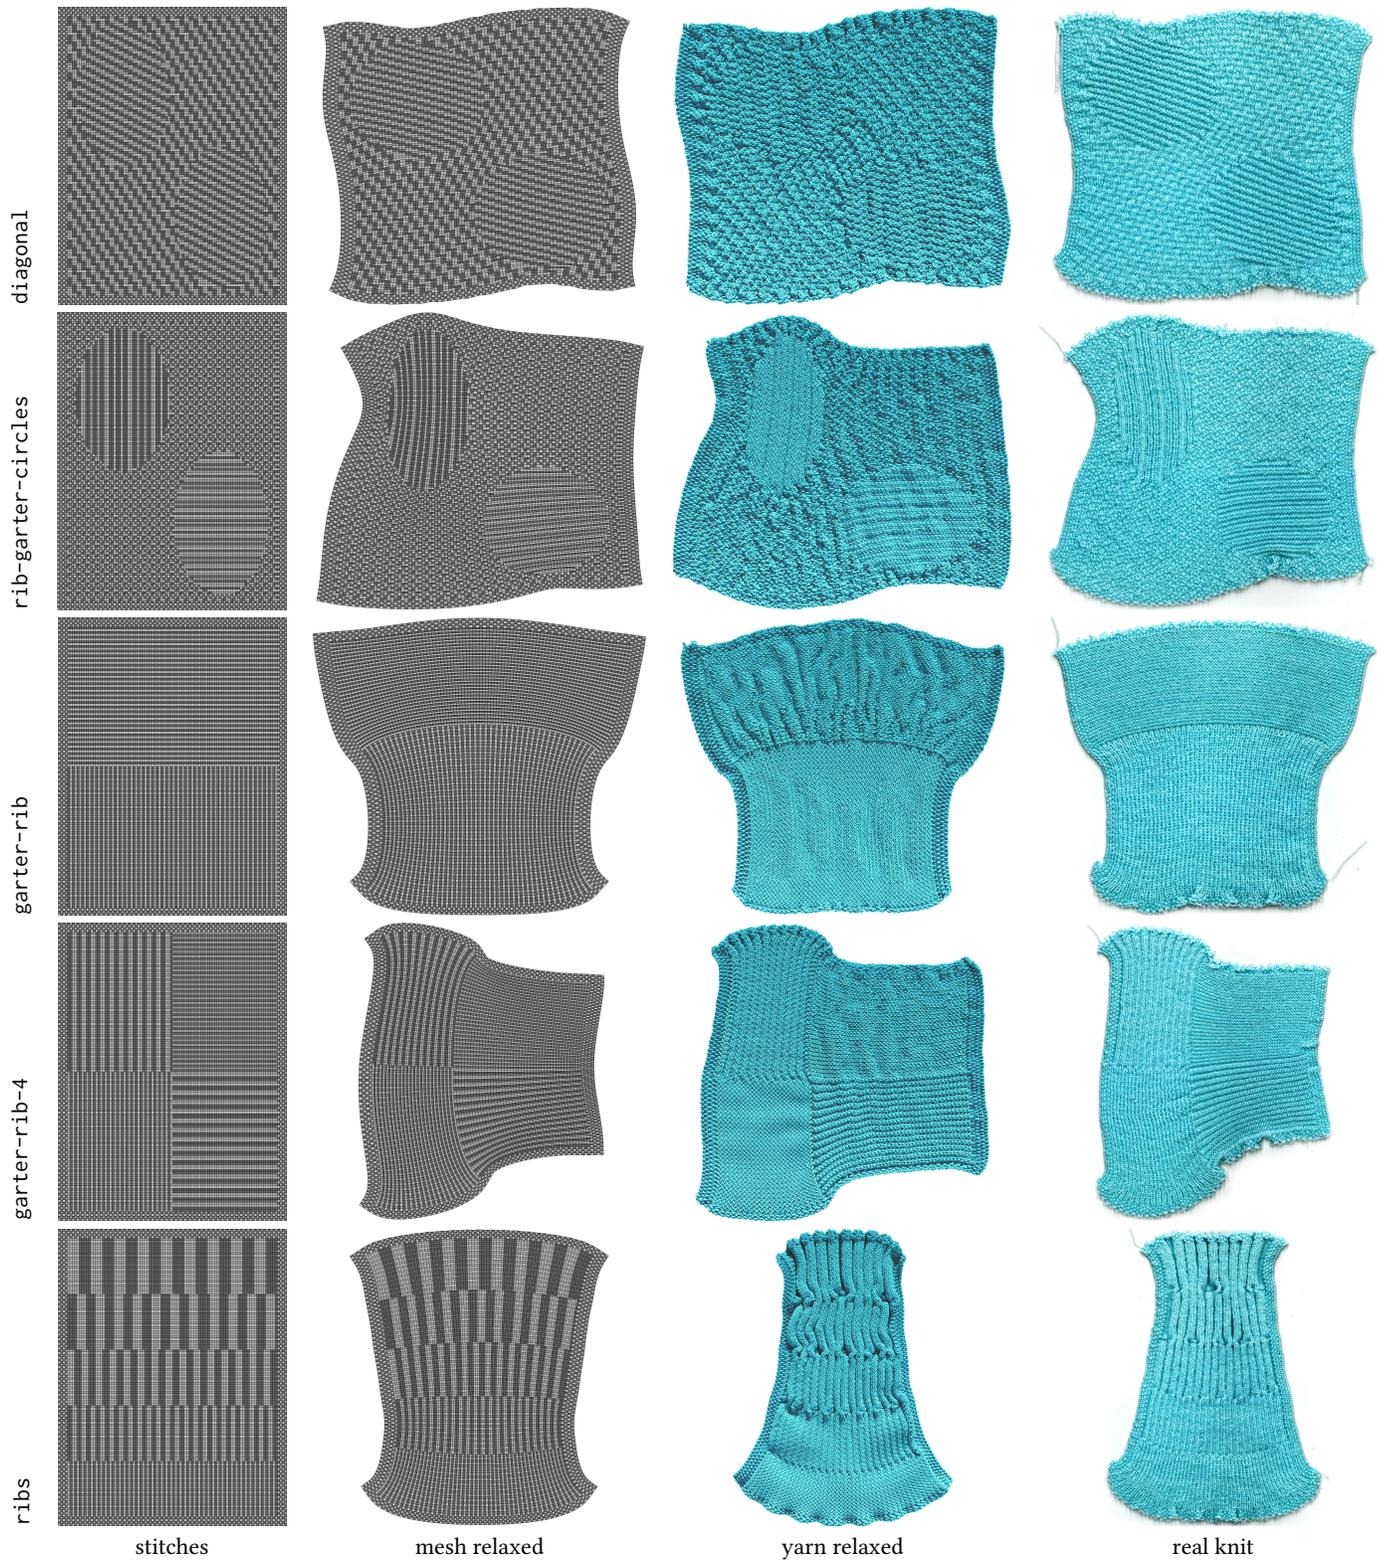


Fig. 16. Comparison to fabricated heterogeneous patterns. (Far Left) Stitch meshes before relaxation. (Center Left) Stitch meshes after relaxation. (Center Right) Yarn relaxation. (Far Right) Photos of real knits.

Table 1. The shape error between masks of scanned swatches and relaxed meshes.

name	x Size (strain)	y Size (strain)	Shear (Degree)	Shape Error (%)
diagonal	0.97	0.97	-0.60	3.58
garter-rib-4	1.02	0.98	2.89	3.76
garter-rib	1.00	0.97	1.95	2.81
rib-garter-circles	1.02	0.99	0.25	4.67
ribs	1.40	0.99	1.28	20.62
tiles	1.08	0.99	1.47	2.44

mask which minimizes the intersection error. We then remove rotation and translation from this affine transform to extract axial strains and shearing. The result is presented in Table 1. Through this comparison, we see that ribs indeed performs the worst, with an intersection error of 20.62%. The remaining patterns all have an intersection error of less than 5%, with minimal strain and shear. The actual mask alignments can be found in the supplemental.

7 DISCUSSION AND FUTURE WORK

In this work, we used neighborhood information to estimate the rest shapes of the patterns that comprise of knits and purls. Despite using a minimal neighborhood and simple set of energies for relaxation, we still managed to obtain compelling results, which suggests that a stitch-level model of knitting serves as a useful middle ground between the precise geometry control afforded by yarn-level models and the abstraction available to mesh-level representations. We believe this means there is a wealth of future directions to explore for refining the accuracy and improving complexity of patterns that can be captured using stitch level models.

For example, this work looks at patterns constructed of the most fundamental stitch types: knit and purl. However, there are a wealth of additional stitch types used in knitting to create even more complicated shapes. For example, increase and decrease stitches may be used to construct curved surfaces by essentially introducing non-quad faces to the stitch mesh. There exist a variety of ways to construct these shaping stitches, each of which have different stiffnesses and rest dimensions. What’s more, there are stitch types such as cables and misses that result in more complex stitch connections than just the immediate four neighborhood. Capturing these complex stitch types would require both extensions to the knit model as well as careful derivation of reasonable calibration sets. Another interesting follow-up direction would be to explore whether stitch-level calibration could be augmented by yarn material measurements. Currently, our system requires fabrics knit with different yarns or machine settings to be re-measured to acquire the corresponding kernel measurements. Given that high-level knit material properties are dominated by the stitch topology [Singal et al. 2024], it would be interesting to see whether a simpler material measurement scheme could be used to adapt the measurements from a single kernel calibration set.

Another interesting future direction is measuring and modeling more complex aspects of knit materials. Many physical knit patterns exhibit out-of-plane, nonlinear behavior, where even deformations

rough-garter-2x2



1C1 : 1C8 : 1C7 : 1C3

Fig. 17. We considered using rough-garter-2x2 instead of stockinette in our training set.

that appear globally in-plane—such as the course-wise contraction observed in 2×2 rib structures—often arise from localized out-of-plane deformations. This makes it challenging to directly measure the role of out-of-plane effects in a given pattern. While simulating such deformations is relatively straightforward, producing accurate and predictive simulations requires material characterization beyond static, rest configuration measurements. Consequently, prior studies that investigate out-of-plane knit behavior typically rely on simulation data or focus on a limited set of specific patterns. Developing a measurement pipeline capable of efficiently capturing the out-of-plane, nonlinear deformation behavior of knitted materials is an interesting open problem.

The measurement of new, complex material parameters would also require a nuanced model for simulating knit-purl patterns. While linear spring forces are a reasonable approximation of knitting in the low stretch regime, prior studies show that the elastic response becomes non-linear as the fabric is stretched [Singal et al. 2024]. This is in part due to the slight curvature present in stitches that compounds to create the out-of-plane deformation. Properly modeling this curvature would not only improve the accuracy of static relaxation results, but could also enable dynamic mesh-level simulations.

Another adjustment to the material model would be to allow different kernels to skew according to measured rest angles. Again, the primary challenge lies in accurately measuring these rest angles, due to the correlation and overlap among kernels. Identifying methods to address this limitation would be an interesting future research.

Finally, our choice of training set was motivated by minimality and linear independence, but there are several such sets that work similarly. Particularly, in developing this work we tried training sets that substituted rough-garter-2x2 (Figure 17) for stockinette; since this makes measurement easier (stockinette has a strong tendency to curl). But this alternative set seems to result in an underestimate of the size of the C1 (all-K) kernel – since they always appear in places where they are deformed out-of-plane by surrounding stitches. In the future, it would be interesting to use a 2.5D capture method to capture this deformation.

8 CONCLUSION

We have presented a neighbor-aware stitch mesh model, which accounts for the stitch types that spatially encompass individual knit structures. Using a finite set of cross-shaped kernels, where the rest configurations are extrapolated from a set of knit patterns, we perform a mesh relaxation on our model. Our results demonstrate

that our neighbor-aware stitch mesh model combined with the data-driven relaxation could estimate the rest shape of the patterns that consisted of mixtures of knits and purls with high accuracy. We believe our stitch-level neighborhood-aware abstraction can be used to capture more dynamic and complex scenarios in the future, and expect that our work will contribute to achieving not only visual but also physical fidelity for virtual knit garments.

ACKNOWLEDGMENTS

This project was supported in part by NSF grant #1956085.

REFERENCES

- David Baraff and Andrew Witkin. 1998. Large steps in cloth simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. Association for Computing Machinery, New York, NY, USA, 43–54. <https://doi.org/10.1145/280814.280821>
- 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.
- David E. Breen, Donald H. House, and Michael J. Wozny. 1994. Predicting the drape of woven cloth using interacting particles. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. Association for Computing Machinery, New York, NY, USA, 365–372. <https://doi.org/10.1145/192161.192259>
- Juan J. Casafranca, Gabriel Cirio, Alejandro Rodriguez, Eder Miguel, and Miguel A. Otaduy. 2020. Mixing Yarns and Triangles in Cloth Simulation. *Computer Graphics Forum* (2020). <https://doi.org/10.1111/cgf.13915>
- Anka He Chen, Ziheng Liu, Yin Yang, and Cem Yuksel. 2024. Vertex Block Descent. *ACM Trans. Graph.* 43, 4, Article 116 (July 2024), 16 pages. <https://doi.org/10.1145/3658179>
- Gabriel Cirio, Jorge Lopez-Moreno, and Miguel A. Otaduy. 2015. Efficient simulation of knitted cloth using persistent contacts. In *Proceedings of the 14th ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (Los Angeles, California) (SCA '15). Association for Computing Machinery, New York, NY, USA, 55–61. <https://doi.org/10.1145/2786784.2786801>
- Gabriel Cirio, Jorge Lopez-Moreno, and Miguel A. Otaduy. 2017. Yarn-Level Cloth Simulation with Sliding Persistent Contacts. *IEEE Transactions on Visualization and Computer Graphics* 23, 2 (Feb. 2017), 1152–1162. <https://doi.org/10.1109/TVCG.2016.2592908>
- David Clyde, Joseph Teran, and Rasmus Tamstorf. 2017. Modeling and data-driven parameter estimation for woven fabrics. In *Proceedings of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (Los Angeles, California) (SCA '17). Association for Computing Machinery, New York, NY, USA, Article 17, 11 pages. <https://doi.org/10.1145/3099564.3099577>
- B. Eberhardt, A. Weber, and W. Strasser. 1996. A fast, flexible, particle-system model for cloth draping. *IEEE Computer Graphics and Applications* 16, 5 (1996), 52–59. <https://doi.org/10.1109/38.536275>
- Carter S. Haines, Na Li, Geoffrey M. Spinks, Ali E. Aliev, Jiangtao Di, and Ray H. Baughman. 2016. New twist on artificial muscles. *Proceedings of the National Academy of Sciences* 113, 42 (2016), 11709–11716. <https://doi.org/10.1073/pnas.1605273113> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.1605273113>
- Jerry Hsu, Tongtong Wang, Kui Wu, and Cem Yuksel. 2025. Stable Cosserat Rods. In *SIGGRAPH 2025 Conference Papers* (Vancouver, Canada) (SIGGRAPH '25). ACM Press, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3721238.3730618>
- Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2008. Simulating knitted cloth at the yarn level. In *ACM SIGGRAPH 2008 Papers* (Los Angeles, California) (SIGGRAPH '08). Association for Computing Machinery, New York, NY, USA, Article 65, 9 pages. <https://doi.org/10.1145/1399504.1360664>
- Jonathan M. Kaldor, Doug L. James, and Steve Marschner. 2010. Efficient yarn-based cloth with adaptive contact linearization. *ACM Trans. Graph.* 29, 4, Article 105 (July 2010), 10 pages. <https://doi.org/10.1145/1778765.1778842>
- Sueo Kawabata. 1980. *The Standardisation and Analysis of Hand Evaluation* (2nd ed.). Textile Machinery Society of Japan, Osaka.
- Jonathan Leaf, Rundong Wu, Eston Schweickart, Doug L. James, and Steve Marschner. 2018. Interactive Design of Yarn-Level Cloth Patterns. *ACM Trans. Graph. (Proceedings of SIGGRAPH Asia 2018)* 37, 6 (11 2018).
- Zishun Liu, Xingjian Han, Yuchen Zhang, Xiangjia Chen, Yu-Kun Lai, Eugeni L. Doubrovski, Emily Whiting, and Charlie C. L. Wang. 2021. Knitting 4D garments with elasticity controlled for body motion. *ACM Transactions on Graphics* 40, 4, Article 62 (July 2021), 16 pages. <https://doi.org/10.1145/3450626.3459868>
- Bryn Lloyd, Gabor Szekely, and Matthias Harders. 2007. Identification of Spring Parameters for Deformable Object Simulation. *IEEE Transactions on Visualization and Computer Graphics* 13, 5 (2007), 1081–1094. <https://doi.org/10.1109/TVCG.2007.1055>
- Sebastian Martin, Peter Kaufmann, Mario Botsch, Eitan Grinspun, and Markus Gross. 2010. Unified simulation of elastic rods, shells, and solids. *ACM Trans. Graph.* 29, 4, Article 39 (July 2010), 10 pages. <https://doi.org/10.1145/1778765.1778776>
- 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 (May 2012), 519–528.
- Vidya Narayanan, Kui Wu, Cem Yuksel, and James McCann. 2019. Visual knitting machine programming. 38, 4, Article 63 (July 2019), 13 pages. <https://doi.org/10.1145/3306346.3322995>
- Jifei Ou, Daniel Oran, Don Derek Haddad, Joseph Paradiso, and Hiroshi Ishii. 2019. SensorKnit: Architecting textile sensors with machine knitting. *3D Printing and Additive Manufacturing* 6, 1 (2019), 1–11.
- Samuel Poincloux, Mokhtar Adda-Bedia, and Frédéric Lechenault. 2018. Geometry and Elasticity of a Knitted Fabric Using a stockinette pattern. Poincloux et al. [Poincloux et al. 2018] create an elastic model with constraints that conserved the total length and bending energy of yarns. They went through elastic response experiments in both homogeneous and inhomogeneous deformations. *Phys. Rev. X* 8 (Jun 2018), 021075. Issue 2. <https://doi.org/10.1103/PhysRevX.8.021075>
- Yannick Rémon, Jean-Michel Nourrit, and Didier Gillard. 2000. A dynamic animation engine for generic spline objects. *The Journal of Visualization and Computer Animation* 11, 1 (2000), 17–26. [https://doi.org/10.1002/\(SICI\)1099-1778\(200002\)11:1<17::AID-VIS213>3.0.CO;2-9](https://doi.org/10.1002/(SICI)1099-1778(200002)11:1<17::AID-VIS213>3.0.CO;2-9)
- Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. 2018. Mechanical characterization of structured sheet materials. *ACM Trans. Graph.* 37, 4, Article 148 (July 2018), 15 pages. <https://doi.org/10.1145/3197517.3201278>
- Krishma Singal, Michael S. Dimitriyev, Sarah E. Gonzalez, A. Patrick Cachine, Sam Quinn, and Elisabetta A. Matsumoto. 2024. Programming mechanics in knitted materials, stitch by stitch. *Nature Communications*, Article 2622 (2024).
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2020. Homogenized yarn-level cloth. *ACM Trans. Graph.* 39, 4, Article 48 (Aug. 2020), 16 pages. <https://doi.org/10.1145/3386569.3392412>
- Georg Sperl, Rahul Narain, and Chris Wojtan. 2021. Mechanics-Aware Deformation of Yarn Pattern Geometry. *ACM Transactions on Graphics (TOG)* 40, 4 (2021).
- Georg Sperl, Rosa M. Sánchez-Banderas, Manwen Li, Chris Wojtan, and Miguel A. Otaduy. 2022. Estimation of yarn-level simulation models for production fabrics. 41, 4, Article 65 (July 2022), 15 pages. <https://doi.org/10.1145/3528223.3530167>
- A. Sydney Gladman, Elisabetta A. Matsumoto, Ralph G. Nuzzo, L. Mahadevan, and Jennifer A. Lewis. 2016. Biomimetic 4D printing. *Nature Materials* 15, 4 (01 Apr 2016), 413–418. <https://doi.org/10.1038/nmat4544>
- Huamin Wang, James F. O'Brien, and Ravi Ramamoorthi. 2011. Data-driven elastic models for cloth: modeling and measurement. *ACM Trans. Graph.* 30, 4, Article 71 (July 2011), 12 pages. <https://doi.org/10.1145/2010324.1964966>
- Kui Wu, Xifeng Gao, Zachary Ferguson, Daniele Panozzo, and Cem Yuksel. 2018. Stitch meshing. 37, 4, Article 130 (July 2018), 14 pages. <https://doi.org/10.1145/3197517.3201360>
- Kui Wu, Hannah Swan, and Cem Yuksel. 2019. Knittable Stitch Meshes. 38, 1, Article 10 (Jan. 2019), 13 pages. <https://doi.org/10.1145/3292481>
- Kui Wu, Marco Tarini, Cem Yuksel, James McCann, and Xifeng Gao. 2022. Wearable 3D Machine Knitting: Automatic Generation of Shaped Knit Sheets to Cover Real-World Objects. *IEEE Transactions on Visualization and Computer Graphics* 28, 9 (2022), 3180–3192. <https://doi.org/10.1109/TVCG.2021.3056101>
- Chun Yuan*, Haoyang Shi*, Lei Lan, Yuxing Qiu, Cem Yuksel, Huamin Wang, Kenfanfu Jiang, Kui Wu, and Yin Yang. 2024. Volumetric Homogenization for Knitwear Simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2024)* 43, 6, Article 207 (12 2024), 19 pages. <https://doi.org/10.1145/3687911> (*Joint First Authors).
- 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, Article 37 (July 2012), 12 pages. <https://doi.org/10.1145/2185520.2185533>
- Wei Zeng, Lin Shu, Qiao Li, Song Chen, Fei Wang, and Xiao-Ming Tao. 2014. Fiber-Based Wearable Electronics: A Review of Materials, Fabrication, Devices, and Applications. *Advanced Materials* 26, 31 (2014), 5310–5336. <https://doi.org/10.1002/adma.201400633> arXiv:<https://advanced.onlinelibrary.wiley.com/doi/pdf/10.1002/adma.201400633>
- Joy Xiaojing Zhang, Gene Wei-Chin Lin, Lukas Bode, Hsiao-Yu Chen, Tuur Stuyck, and Egor Larionov. 2024. Estimating Cloth Elasticity Parameters From Homogenized Yarn-Level Models. In *Proceedings of the 17th ACM SIGGRAPH Conference on Motion, Interaction, and Games* (Arlington, VA, USA) (MIG '24). Association for Computing Machinery, New York, NY, USA, Article 7, 12 pages. <https://doi.org/10.1145/3677388.3696340>