DIGISEW: Anisotropic Stitching for Variable Stretch in Textiles

Abhinit Sati
Clemson University
Clemson, SC, USA
asati@clemson.edu

Ioannis Karamouzas
Clemson University
Clemson, SC, USA
ioannis@clemson.edu

Victor Zordan
Clemson University
Clemson, SC, USA
vbz@clemson.edu

ABSTRACT

The textile industry touches many aspects of our daily lives, with clothing, furniture, vehicle interiors and covers, as well as a plethora of medical, sports, and leisure-driven specialized products. This research aims to expand the types of fabric properties that are available for design and manufacturing by introducing methods for modifying material stiffness and tensile characteristics. Specifically, this paper introduces a technique to incorporate anisotropic stitching to control direction and strength of a fabric's stretch through the use of an embroidery machine and computer-driven stitch design and planning. The contributions of this paper include: a method for specifying and controlling direction in stitch planning; a sequential stitch planner that incorporates both density and direction; and a showcase of results that support the value and uniqueness of this new process of manufacturing for textile artifacts.

CCS CONCEPTS

• **Computing methodologies** → **Graphics systems and interfaces**; • **Applied computing** → **Computer-aided manufacturing**.

KEYWORDS

anisotropic stitching, fabrication, stitch planning, smart textile

ACM Reference Format:
Abhinit Sati, Ioannis Karamouzas, and Victor Zordan. 2021. DIGISEW: Anisotropic Stitching for Variable Stretch in Textiles. In *Symposium on Computational Fabrication (SCF '21), October 28–29, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3485114.3485121

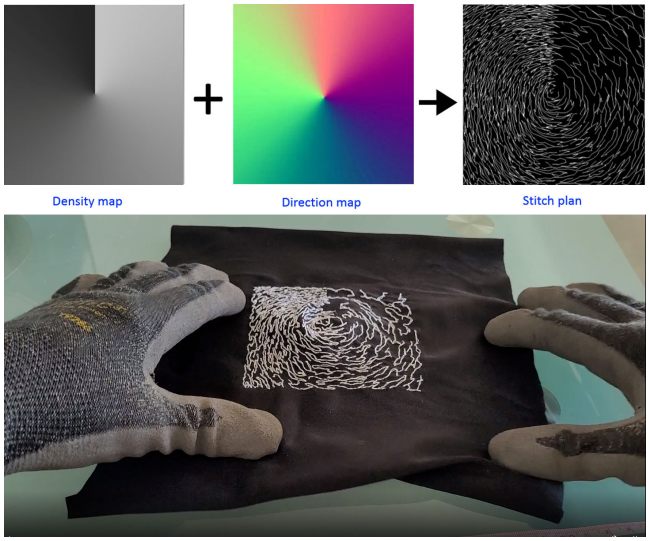

Figure 1: Direction and density combine to create final embroidered cloth with rotational stretch patterns and unique tensile properties across the resulting fabric surface.

1 INTRODUCTION

Considering the profound and substantial capacities of woven textiles and how they impact our lives, there are limited ways in which their core material properties vary, even in the wide array of fabrics available today. For example, while 2- and 4-way stretch materials are commonplace, there is little variation in heterogeneous stretch and even fewer examples that exhibit stretch variation across surfaces in off-the-shelf fabrics. We propose a technology for processing raw woven textile to change tensile properties, yielding fabrics that exhibit arbitrary stretch behavior in arbitrary directions across a fabric's surface. The result is a cloth that has specific local characteristics, including stretch strength and direction that varies and changes in coherent ways (See Figure 1). Our aim in this work is to build next-generation fabric resources that can be incorporated into intelligent textile design at a fundamental level. We approach this unique technology using an embroiderer and exploiting methodologies from computational fabrication. Namely, in this paper, we propose *anisotropic stitching* as a flexible approach to combine path planning and purposeful design with embroidered (physical) stitching. The result is a stitched fabric that adheres to the stretch characteristics present in a plan built to accomplish specific desired anisotropic characteristics.

Computational fabrication is gaining traction as a viable technology for a wide spectrum of products, as computational power has increased rapidly and quality simulation tools have become mainstream. In contrast, textiles have changed very slowly over hundreds of years. With the proposed anisotropic stitching technology and the general approach employed here, we believe we will be able to generate fabrics and designs that have not been made to date - either due to infeasibility in manufacturing - or inability for designers to produce such results by hand. Specifically, this paper introduces: a stitch model that includes anisotropy based on two forms of image-based maps that control density and direction, respectively (see Figure 1); and a process for discretizing and a planning algorithm for stitching - which we compare to previous related results [Moore et al. 2018]. Further, we identify a set of

SCF '21, October 28–29, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-9090-3/21/10…$15.00
https://doi.org/10.1145/3485114.3485121

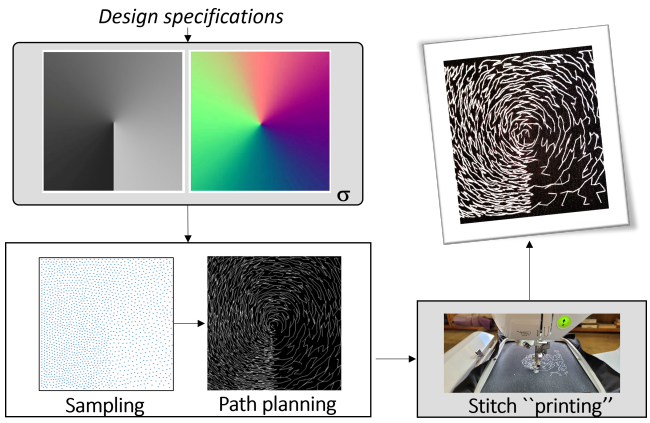

Figure 2: Overview of approach.

representative solutions for a set of conceptual design problems and everyday applications.

1.1 Overview

There are several subcomponents that work in aggregate to realize the desired solution, as shown in Figure 2. At the center of the approach is a stitch planner that builds a sequential list of stitches based on a desired stitch plan. The plan is embedded into an aggregated direction-and-density function, σ, that corresponds to the desired stretch properties across the fabric surface. To realize this plan, we discretize σ into a set of stitch locations using adaptive sampling with Centroidal Voronoi Tesselation [Du et al. 1999]. The planner takes in the function and unordered stitch locations and exports a stitch path based on a modified Dijkstra search[Dijkstra et al. 1959]. This stitch path is then ported to the embroiderer to be fabricated as a list of stitches and their locations.

2 RELATED WORK

In recent years, there has been a considerable amount of attention on computational techniques employed to alter physical properties of objects for various materials. For example, the work of [Panetta et al. 2015; Schumacher et al. 2015; Zhu et al. 2017] demonstrates how to control material properties, such as Poisson ratio and Young's modulus, for 3D printed objects by manipulating microstructures, allowing the creation of both isotropic and anisotropic structures. [Pérez et al. 2015; Schumacher et al. 2018; Tang et al. 2017] demonstrate analogous efforts focused on lattice and rod structures. Zehnder et al. [2017] alter mechanical properties at the macroscopic level by doping silicone with liquid using a filament-based 3D printer. The works of [Martínez et al. 2016; Martínez et al. 2017] use procedural Voronoi foams to control the rigidity/flexibility of fabricated models by manipulating isotropic and orthotropic properties of the underlying microstructures, respectively. Indeed, the growing field in computational fabrication has focused on a variety of unique problems [Ion et al. 2016; Konaković et al. 2016; Skouras et al. 2014] (among many more). In total, this work inspires our attention of solving computational fabrication to modify material properties as many of the mentioned papers do.

Within the field, in addition, there is a growing body of work that is looking at computational fabrication as it applies to the synthesis and incorporation of textiles and fabric objects. The growing work of McCann and colleagues addresses formalization of computational design as it applies to knitting by proposing a general knitting compiler [Lin et al. 2018; McCann et al. 2016; Narayanan et al. 2018, 2019]. Researchers have also looked at other related problems, including smart embroidery [Stoychev et al. 2017], textiles covering 3D shapes [Mahdavi-Amiri et al. 2015], 3D printing using felted fabric [Peng et al. 2016], user-assisted 3D knitting [Igarashi et al. 2008a,b], and 3D weaving [Wu et al. 2020]. Liu et al. [2017] demonstrate the ability to make custom embroidery mimicking an input photograph. Recently, Liu et al. [2021] have demonstrated techniques for controlling local elasticity of 3D garments for better comfort, by designing stitch meshes composed of soft and firm yarns knitted together to reduce deformations in regions of high stress. In contrast to these works, pinnacle to our effort is an aim on fabrication of a unique finished product, a *woven* "smart-stretch" fabric with directed and varying stretch properties.

Related to our own work is [Zhang et al. 2019] where flat fabric panels are generated from a 3D shape and sewn together for casting. While their end result is a plaster cast, their research motivates our efforts through the use of functional cloth. Likewise, Guseinov et al. [2017] combine 3D printing and stretch fabric to produce 3D models. While their goal is to use a material's stretch to support a transformation from 2D to 3D, we are motivated by their effort in combining fabric and 3D printing. Pérez et al. [2017] accomplish a similar feat by using planar rod networks embedded on pre-stretched fabric that form complex 3D models upon actuation.

Closest to the goal of this paper, Moore et al. [2018] have demonstrated successfully the production of patterns that control stiffness uniformly by solving a path planning problem modelled as a Travelling Salesman Problem. In contrast, we address a new problem, the non-uniform or anisotropic stitch planning problem, and develop novel techniques to control the direction of stitching. Namely, we propose a new *direction* map to this end, as well as a new planning approach to solve the anisotropic stitching problem that relies on a shortest path tree formulation and a novel cost function. Our solution is able to create a wide range of stitch plans, and offers great flexibility to alter the stretch properties of the underlying fabric.

3 ANISOTROPIC STITCH MODEL

We define the idealized fabric stretch within a cloth's surface as a four-dimensional function $\sigma(u, v)$ expressed in the cloth-specific surface coordinate frame, (u, v). In principle, σ parameterizes stitching through a two-dimensional direction (frame) of the major and minor stretch axes $\mathbf{n}$; the major axis desired stretch, s; and the ratio of the major and minor densities, ω. In common fabrics, the value of σ can be thought of (usually) as fixed and global parameters for each σ element suffice to describe the material's behaviors, globally and locally. For example, a simple 2-way stretch allows stretching

that is aligned with the fabric weave and it is (reasonably) inextensible in the weft (orthogonal) direction. This indicates the axes and ratio, and a single description for the stretch s value allows the cloth to be described at all locations (u, v).

For our purposes, the function σ is specified in an idealized manner to be localized within specific regions of the cloth, with no assumptions that the cloth stretch remain constant across the surface. In practice, physical limitations impede the process we employ for constructing smart-stretch fabrics, i.e., through precision limits of stitches via an embroiderer, including the size and strength of the stitch, and the strength of the fabric before and after stitching. Further, the embroidery process expects a sequential stitching process, which cannot yield all stretch patterns, e.g. very high-frequency stretch/no stretch patterns. Nonetheless, the focus of *anisotropic stitching* is to produce a non-uniform σ to support the production of smart stretch fabrics.

In practice, we embed a visual coding of σ into an image pair and use this to visualize and store a discretized spatial field of the four dimensional stretch design. Note, this construction is an idealized expression of the stretch parametrized over the cloth surface, i.e. outside of the physical limitations described. Following [Moore et al. 2018], we employ a grayscale *density map* image to specify the overall stretch value, s. Next, we introduce the use of normal maps [Blinn 1978] to encode directional information for the major and minor axes, $\mathbf{n}$, as well as the ratio of stretch between them, ω. We call this the *direction map*. One interpretation of the direction map is that every (u, v) point on the fabric can be mapped to an idealized orientation in which stitches should be made to align with the desired anisotropic stretch expressed in σ. However, this is just a preference, and the planner may not always be in a position to follow the preferred direction exactly, as further explained in Section 4.

The following converts 24-bit color values to 2D direction vectors $\mathbf{n}$:

$$\mathbf{n} = \begin{pmatrix} \frac{2 \cdot r}{255} - 1 \\ \frac{2 \cdot g}{255} - 1 \end{pmatrix}, \tag{1}$$

where r and g represent the red and green channel values respectively. Thus, this formula maps pixel values [0, 255] to direction values [-1, 1]. While this representation uses two color channels to store a direction which could be represented by a single value, we adopt the convention to follow standards for normal maps. We will refer to these directions as normals in the future. The figure below shows a visualization of how the normals are interpreted for a sample direction map.

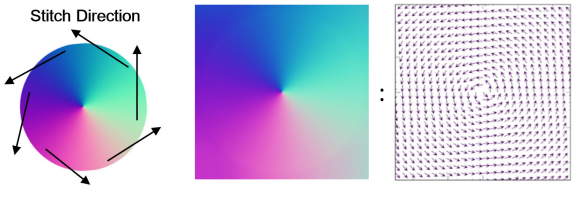

Figure 3: Normal map interpreted as stitch direction map.

As we do not make use of the blue color channel to indicate direction, we instead employ the channel value b to control the ratio, ω, which modulates the degree of anisotropy. As such, we compute ω based on the blue channel as

$$\omega = \frac{(b - 127)}{127}. \tag{2}$$

To adhere to the standards of normal maps where b represents the value of the blue channel, we make the assumption that b is in the range from 128 to 255. We describe the exact use of this term in the next section, but one interpretation of this parameter is its control over the amount of randomness in a given direction. Thus $\omega = 0$ indicates perfect directional alignment, while a value of 1 indicates complete isotropy, i.e. uniform directional stretch. Also note, the embroidered stitch is not directed in the manner in which it impedes stretch, and so each direction color includes an opposite direction "complement" which we will exploit in the planner described next.

4 STITCH PLANNING

Given the desired inputs for direction and density, the stitch planner builds the sequential series of stitches that the embroiderer will use to modify the fabric. We handle this in a two step process, first discretizing the continuous input to a set of sample points to stitch and then plan a path to stitch/connect these points. Thereby, we convert the stitch serialization problem to a path search problem and borrow ideas from the motion planning community to find a solution which is specific to our needs. Foremost, to uphold the anisotropic-stitch goal of this paper, the direction of the stitch is paramount. In addition, based on our empirical observations, the stitch should be short, since long stitches tend to be unwieldy and have the tendency to break. Also, when stitching is not strictly anisotropic ($0 < \omega \leq 1$), promoting turning and non-straight stitches is also preferred [Moore et al. 2018]. We seek a planner that can handle all of these considerations in its search.

4.1 Stitch discretization

A simple approach to produce discrete stitch locations is adaptive sampling. Specifically, we divide the stitch surface into subregions and intensity values based on the average desired density for the local area represented in σ. These intensity values are then used to compute the number of points that need to be sampled in the respective subregion. The sampled points are jittered to avoid aliasing artifacts.

While this simple approach works well in many cases, it has a tendency to sample points in clusters that are close to each other on occasion. Consequently, when the stitch planner connects these points by making stitches between them, some stitches end up crossing over each other. When this issue arises, we can turn to Centroidal Voronoi Tesselation (CVT) technique for sampling points [Du et al. 1999]. The CVT produces a weighted centroidal Voronoi diagram, in which each center of mass (centroid) of a Voronoi region is weighted according to the density map defined in σ. We use these centroids as sample points for stitch planning.

While many techniques have been developed over the years to compute a CVT [Hateley et al. 2015; Liu et al. 2009], we employ Lloyd's algorithm with the sampling process performed offline [Du et al. 2006]. In our implementation, we start with sites which are

obtained from a simple dithering approach. And then, we iteratively compute the weighted Voronoi diagram of these sites, integrate Voronoi regions and compute the corresponding new centroids, and finally move the sites towards the newly computed centroids. This process repeats until the distance between sites and centroids becomes sufficiently small. At this point we have a near optimal distribution of sample points. After the process is complete, we jitter all points in one final pass to eliminate bias.

Empirically we found sampling with CVT produces well spaced uniformly random points, which respect the density values encoded in σ, and ensures stitches do not overlap with each other. We employ both CVT and simple adaptive sampling in the results in this paper.

4.2 Path planning

Next, we seek for a stitch plan that is minimal in its structure but covers all sampled points on the fabric whilst keeping stitches directional and reasonably short. To do so, we start with an undirected graph formed by the sampling of the density map, where each vertex denotes a 2D sample point and each edge $(\mathbf{u}, \mathbf{v})$ denotes a potential stitch between the vertices $\mathbf{u}$ and $\mathbf{v}$. From this, we can consider employing the traveling salesman problem (TSP) as formulated by [Moore et al. 2018] where the goal is to find a path (series of stitches) that visits all points exactly once with minimum total cost. However, such a formulation lacks flexibility resulting in jump stitches when a dead end is hit, i.e. when a point has no close available neighbors. Instead, we relax the condition for visiting every sampled point exactly once, and introduce the use of minimum spanning trees, specifically a variant called the shortest path tree (SPT) [Cormen et al. 2009] to produce our stitch serialization. Further, while TSP can offer limited heuristic-based corrections for jump stitches [Moore et al. 2018], SPT supports more flexible clean up as we describe below.

The underlying coverage planning problem that accounts for stitch direction and density can then be formulated as a classic single-source shortest path search [Bellman 1958; Dijkstra et al. 1959]. We randomly pick a sample point as the source that becomes the root of the tree and iteratively construct the SPT using Dijkstra's algorithm. To do so, we propose a novel function to assess the edge cost of a potential stitch that combines two terms. The first term encourages directed stitches to follow σ as close as possible

$$cost_1(\mathbf{u}, \mathbf{v}, \mathbf{n}) = -{\alpha_1}^{-|\mathbf{v}-\mathbf{u}|}{\beta_1}^{|(\mathbf{v}-\mathbf{u})\cdot\mathbf{n}|}, \tag{3}$$

where $(\mathbf{u}, \mathbf{v})$ is the stitch we are evaluating, and $\mathbf{n}$ is the preferred direction at the sampled point $\mathbf{u}$, computed from the direction map. The α_1 term rewards shorter stitches. The exponent on β_1 encourages $(\mathbf{u}, \mathbf{v})$ to align with $\mathbf{n}$ while the absolute value allows both complements from the normal map to be considered equally. α_1 and β_1 control the importance of the two cost terms, with $\alpha_1, \beta_1 \geq 1$. Further, to support the option of uniform (non-directional) stiffening, we adapt a second cost term derived from [Moore et al. 2018]. This cost is computed as

$$cost_2(\mathbf{u}, \mathbf{v}, \mathbf{t}) = -{\alpha_2}^{-|\mathbf{v}-\mathbf{u}|}{\beta_2}^{-|(\mathbf{u}-\mathbf{t})\cdot(\mathbf{v}-\mathbf{u})|}, \tag{4}$$

where $(\mathbf{u}, \mathbf{v})$ is the stitch under consideration for evaluation, $\mathbf{t}$ is the stitch point previous to $\mathbf{u}$, and $\alpha_2, \beta_2 \geq 1$. In this cost term, the two stitches are rewarded for forming right turns, and any straight lines or close will be penalized.

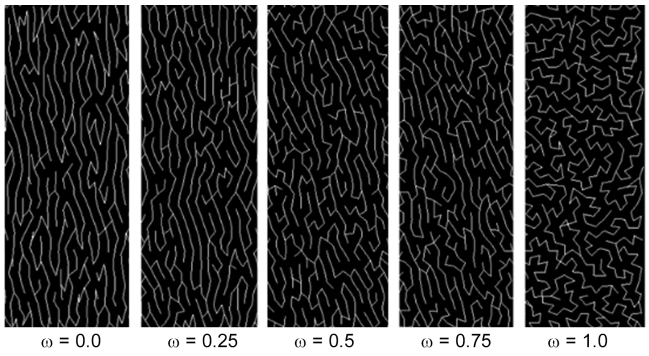

Figure 4: Stitch plans produced with different values of ω that regulate the level of directional alignment. All density and color directions are fixed for all samples, the only variable is the blue channel in σ.

We linearly combine the two cost terms through ω, given the blue signal provided in the normal map, and compute the final cost value as:

$$cost_{total} = (1 - \omega) \cdot cost_1 + \omega \cdot cost_2. \tag{5}$$

When ω = 0.0, the first cost term aims to align stitches with the ideal direction taken from σ. If ω = 1.0 (totally biased towards uniform stitching), we get results similar to ones in the Moore et al. [2018] paper making sharp turns to avoid long stitches and promoting uniform stiffness at the same time. The range in between provides a space of solutions that blends smoothly from the former to the latter – Figure 4 shows plans obtained by interpolating the two cost functions in this manner.

The SPT algorithm computes a stitch plan where both direction is followed and stitch length is kept in check, while adhering to the desired σ. However, it does produce some off-directional stitches, as can be seen in Figure 5b. This is expected, as the cost function is local and dependent on the currently available neighbors' length and direction. To address this issue, we propose a second pass to clean up the undesirable stitches. Consider the example shown in Figure 5a. In the left graph, the arrow at node 'f' represents the preferred direction. The stitch (f, e) goes orthogonal to this direction and will be flagged as an off-directional stitch that needs to be removed. Subsequently, we add a new stitch/edge between the two separated tree components by selecting the one with the lowest cost. Removing and adding edges in trees has been studied at length, and our new sub-problem can be solved by tagging and removing off-directional stitches, and reconnecting the newly separated components. There exists a variety of algorithms for this connected component labelling problem [He et al. 2017]. As we have only two components at any given point, we opt for the simplest solution, a depth first traversal (see Algorithm 2). Results of the cleanup procedure can be seen in 5c.

Our complete stitch planning approach can be summarized in two procedures, shown in Algorithms 1 and 2 below. The result of our solution is a tree rooted at a chosen sample point that is connected to all other samples via local paths that have minimum cost. As our embroiderer requires a continuous path as input, we

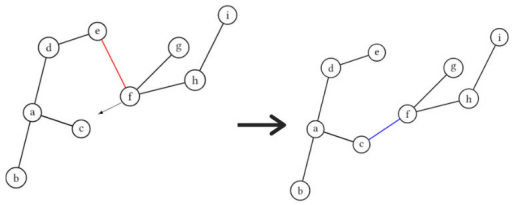

(a)

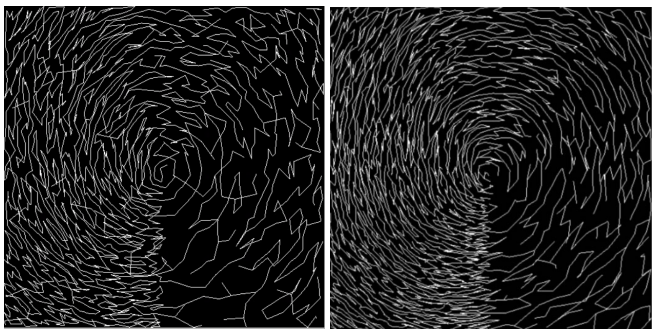

(b) (c)

Figure 5: (a) A typical iteration of the cleanup procedure. (b) The SPT algorithm before a cleanup procedure. (c) Stitch plan following post processing routine.

employ a depth-first traversal of the tree to trace all vertices and return the final, global path.

ALGORITHM 1: Stitch Planner

```
Plan(start, points, normal)
// Run Dijkstra's to generate the SPT
spt ← GenerateSPT(start, points, normal);
// get individual stitches from the SPT
stitches ← GetStitches(spt);
// Fix the off-directional stitches
SecondPass(spt, stitches, normal);
// Convert tree to a path
path ← DFS(spt)
return path
```

ALGORITHM 2: Cleaning up off-directional stitches

```
SecondPass(spt, stitches, normal)
forall e ∈ stitches do
    u ← e.u; v ← e.v;
    spt.remove(e);
    // Label the connected components
    spt.traverse(u, 1);
    spt.traverse(v, 2);
    // Select least cost vertex from the other connected component
    ub ← getBestNode(u, c1, 1);
    vb ← getBestNode(v, c2, 2);
    if c1 < c2 then
        nu ← u; nv ← ub;
    end
    else
        nu ← v; nv ← vb;
    end
    // Compare costs between the old and new edge
    w1 ← cost(nu, nv, normal[nu]);
    w2 ← cost(u, v, normal[u]);
    if w1 < w2 then
        spt.add(edge(nu, nv));
    end
    else
        spt.add(edge(u, v));
    end
end
```

5 IMPLEMENTATION

5.1 Stitch planning

Since our embroiderer can stitch only continuous paths, we traverse the output tree in a depth first fashion to produce a "double stitched" continuous path. The reason we call it double stitched is because the embroiderer has to repeat every stitch it made going down when backtracking or unwinding. This has an unexpected positive consequence of making individual stitches physically stronger, resulting in better performance under greater loads and stress. Given also the flexibility that our formulation provides both in terms of planning and post-processing, we posit that the planner, combined with our novel cost function, employed in this manner is well suited for anisotropic stitching in physical products.

For our Dijkstra implementation, we start from a random sample point as our source vertex, sample neighbors around that vertex based on a user-defined radius (typically set to 5 units), and then pick the vertex with the least cost from the source to continue exploration. A priority queue is maintained to select the next best vertex for exploration. We store the total cost from the source to other vertices and relax vertices when need be. While other approaches such as best first search can be used to generate a tree, our Dijkstra approach leads to the shortest tree, promoting a more continuous looking stitch plan and avoid making too many branches.

We provide further implementation details for the results in the paper in Table 1, including the α and β values of our cost function (Eqs. 3 and 4) used and the discretization method for our different experiments.

5.2 No-go boundaries

The result produced in Figure 6 (top left) induces some undesirable effects in the stitch plan that cross the circular regions. Observe that the two circles at the top and one at the center have stitches crossing over their centers, an artifact not seen in the bottom two circles. The bottom leads to the (desired) formation of soft spots at the centers of these two circles which makes the fabric more compliant at these points. The remaining three circles on the other hand resist forces applied to their centers due to the inextension introduced by the stitches going through their center.

To further enforce constraints we wish the planner to adhere to, we propose the usage of "no-go" boundaries in order to obtain

Table 1: Description of the different alphas and betas and choice of discretization used to produce all results reported.

Discretization and Cost Parameters							
Density	Normal	ω	$\alpha 1$	$\beta 1$	$\alpha 2$	$\beta 2$	Sample
Keyhole	Curvy	0.0	1.1	1.5	1.1	5.1	Adapt
Sunset	Cone	0.0	1.1	1.5	1.1	5.5	Adapt
Uniform	Polka	0.0	1.1	1.5	1.1	5.2	Adapt
Uniform	Vertical	vary	1.05	2.05	2.55	4.25	CVT

consistent (desired) results. The no-go boundary defines a set of segments on the fabric that no stitch may cross over. Figure 6 (bottom left) shows an example of no-go boundaries visually over a normal map, highlighted in white. We employ a straightforward approach to solving the problem by disallowing all stitches that violate the boundaries chosen by the user. With this approach, we can guarantee that there wouldn't be any stitches going over the centers of the circles, producing soft spots in each of the circle centers.

5.3 Printing

We employ a Brother 800 SE Embroiderer that is intended for consumer use. In practice, this was sufficient for out prototyping but a professional grade embroiderer will add additional benefit, including scaling the final products. Our base fabric is a 4-way elastane material similar to a medium weight spandex that might be used for fitted swimwear. To make this stretch fabric printable, we treat the base fabric with Solvy, a water soluble stabilizer, that makes stretch fabric temporarily inextensible during embroidering process. This stabilizer is removed before use (or measurement) by rinsing the stitched product. We employ standard non-stretch nylon embroidery thread for the stitching and bobbin in the embroidery process. We note that the choice of base fabric and thread will impact the effect of the final product as the stitches modify this fabric's properties. A light-weight or heavy-weight stretch base may be more appropriate based on the anticipated application. Likewise employing stretch thread proposes an interesting direction for future work.

6 ANALYSIS AND APPLICATIONS

6.1 Experimental validation

We performed experimental validation by measuring the strain under load for printed artifacts to provide evidence of the efficacy of the technique. Because our goal is variable stretch, including changing direction, it is difficult to measure in a complex design. Thus, we devise a simple mechanical testing set-up (shown in Figure 7, left) to assess aggregate effects, and opt to measure the displacement of uniform samples of material. In practice, we convert displacement to strain to normalize for any length variations between samples - although all printed examples were made to be of the same original size, a 10x10cm area. As measures showed efficacy in uniform density stitch plans at changing (uniform) stiffness previously [Moore et al. 2018], we focus our analysis on the impact of anisotropic stitching plans and its blending with uniform stitching in this paper. Specifically, in the plot in Figure 7, we show the response to the loading for the base fabric, in the warp and weft directions, and we compare this to two variants with fixed density and uniform

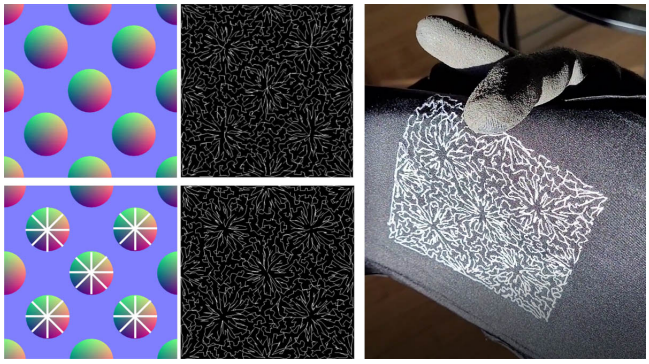

Figure 6: Anisotropic stitch example with polka dots, without (top-left) and with (bottom-left) no-go regions . The subsequent printed stitch output, shown at the right, has "relief" regions in the centers of each polkadot with reduced stiffness producing a "textured" fabric with unique stretch characteristics not unlike an "egg-crate" style foam pad.

anisotropic maps under two blend conditions. The normal aligns the stitching with the warp direction. Note we performed this experiment in both alignments (warp- and weft- aligned) and the result is similar for both. We opt to show the result here (for warp aligned) as the slight differences in warp and weft are at odds with the planned anisotropy, and therefore the base bias does not explain the anisotropy in the result, showing it is derived from the stitch procedure. We opt for $\omega = 0.25$ which is mostly anisotropic, and $\omega = 0.75$, mostly isotropic as we felt these were representative of both anisotropic and isotropic blend conditions.

From the figure, the plot shows the strain increases fairly smoothly as the load increases in all measured directions. We also observe that the base fabric, in warp and weft, have the lowest stiffness (revealed by the largest strain) over the tested range of loads - not surprisingly. In comparison, the two uniform density, uniform direction samples reveal a drop in stretch (corresponding to lower strain). While the 0.25-blend example shows the largest drop in strain (equivalent to greater resistance to stretch) in the warp direction, it shows a low change in the weft. This indicates we are indeed stiffening the desired directions in the anticipated (anisotropic) manner. In contrast, as the 0.75-blend example treats the directions more evenly, the measurements reveal a change in both warp and weft with a preference to increase stiffening in weft over the previous sample as the small amount of anisotropy would warrant.

6.2 Stitch planning analysis

Next, we analyze the quality of our planning algorithm based on computational errors from the idealize σ design and the stitched output plan. To provide a ground for measurement, we show how our proposed (Dijkstra-based) stitch planning approach compares to the TSP formulation as previously proposed by Moore et al. [2018]. Note, as the previous work focuses only on uniform stitching, we had to modify their solution to account for directionality using our anisotropic cost function from Equation 5. We design a simple test to gauge the efficacy of the two solutions for non-uniform stitching. Namely, we generate a stitch plan with each algorithm,

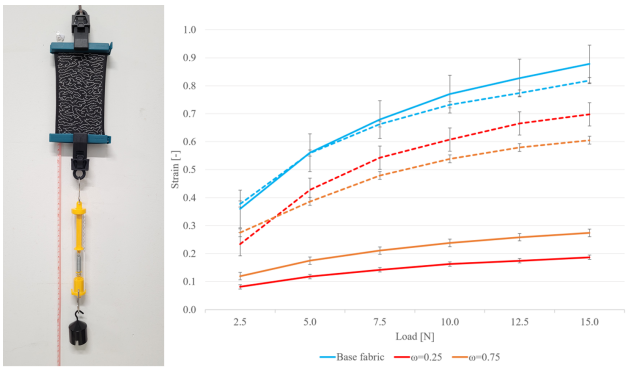

Figure 7: (Left) Experimental setup to validate the efficacy of our approach to change the properties of a physical fabric. (Right) Corresponding results that compare the base fabric (blue) with two stitched examples. Warp direction is solid, weft is dashed. Both samples have anisotropy directed to increase stiffness in the warp direction more than the weft as is clearly demonstrated. Between the two stitched samples, one is highly anisotropic in the warp direction (red, $\omega = 0.25$) and one which is more isotropic but with bias to the warp direction (gold, $\omega = 0.75$). Each stitch print was printed two times and measures were taken three times each to produce the error bars shown.

and use the directions of the output stitches to produce errors color maps (see Figure 8). The points in these maps match the sampled points from our discretization (Section 4.1). To compute color values, we compare the similarity of the stitch direction vector at a given point with the preferred direction in the normal map from σ corresponding to that point, and use that result (ranging from 0.0 to 1.0) as a correctness value. Above a small threshold the stitch is indicated as misaligned and colored red (see figure) otherwise the stitch is drawn in black. Note, for points with multiple stitches, we average all stitches at each point to get a final correctness value for that point. We refer to Figure 8 for examples of these error maps.

Quantitatively, Table 2 shows the Root Mean Squared Error (RMSE) values between the reversed map and the original normal map under a variety of conditions. Our approach consistently outperforms TSP highlighting the benefit of our formulation. Additional evaluations and details appear in the following thesis [Sati 2021].

6.3 Everyday things

To motivate this work as part of a greater exploration of such technology in the design of textile artifacts, we highlight a number of extended conceptual designs for everyday products. Notably, a wide spectrum of domains use fabrics in specialized manners. For example in medicine, custom harnesses immobilize hip dysplacia in infants, while compression sleeves and braces restrict joint movement for injury care and promote blood circulation. Likewise, there are entire apparel industries focused on body shape "sculpting" and support through textile and textile-rich products. To showcase the

Table 2: Comparison between TSP and our approach. The reported numbers denote the RMSE between the given normal map in σ and the correction map computed from the generated plan. The errors are normalized to lie in the range of 0-100.

Correction map errors			
Density map	Normal map	RMSE TSP	RMSE Ours
Keyhole	Curvy	6.6	5.2
Target	Cone	9.4	5.1
Sunset	Cone	9.9	4.9
Polka	Polka	10.2	6.9
Uniform	Polka	8.4	4.4

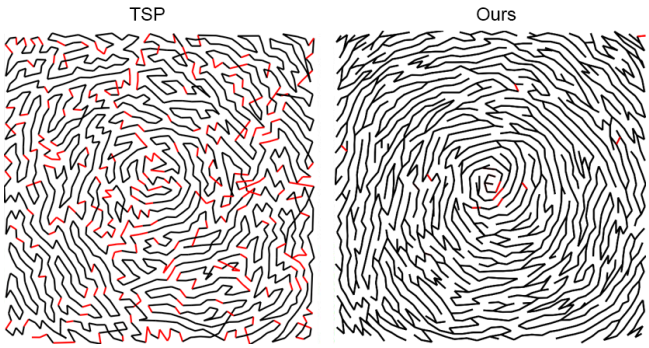

Figure 8: Left: Stitch plan rendered with TSP. Right: Stitch plan rendered with our proposed approach. Black represent the stitches aligned with the plan while red shows deviation from the plan, in both images.

benefit of this technology in the wide array of textile designs, we introduce a small suite of conceptual design examples.

We envision that cost-effective custom pressure insoles are a plausible fast-to-market example, Figure 9 (left-top), conveys a sample manufacturing approach. The pressure profile of an individual is commonplace (e.g. in many pharmacies today). We can employ this to produce a stitch plan that conforms and relieves pressure using the proposed technology. As shown in the figure, we anticipate embedding this custom fabric stitch "print" between two foam layers that provide cushion and lateral support. We do not employ a direction map for this example.

Another motivating example is the unique sports shoe, shown in Figure 9 (left-middle). This shoe is designed to be made from a single fabric upper which is both inexpensive and lightweight (akin to an elastane sock plus rubber sole). While similar shoes exist to date, the added value is shown as directional functional stitching to replace laces and the need for additional elastic webbing. The normal map directs stitching in the fabric space to support and lock in the foot when the shoe is worn while allowing flexibility for the foot to bend and remain comfortable.

High-end seating and specialized needs for chair-bound individuals motivate our next example. Figure 9 (left-bottom) shows a similar pressure profile as describe in the insole example, this

Figure 9: Potential applications of the proposed technology. (Top-left) Concept example for customized insole from foot pressure reading. (Middle-left) Shoe model where stitches create a comfortable stylish and low cost upper that is also supportive like a lace or buckle. (Bottom-left) Example of director's chair that supports pressure profile shown. (Right) Real-world prototype of variable stretch for support bra.

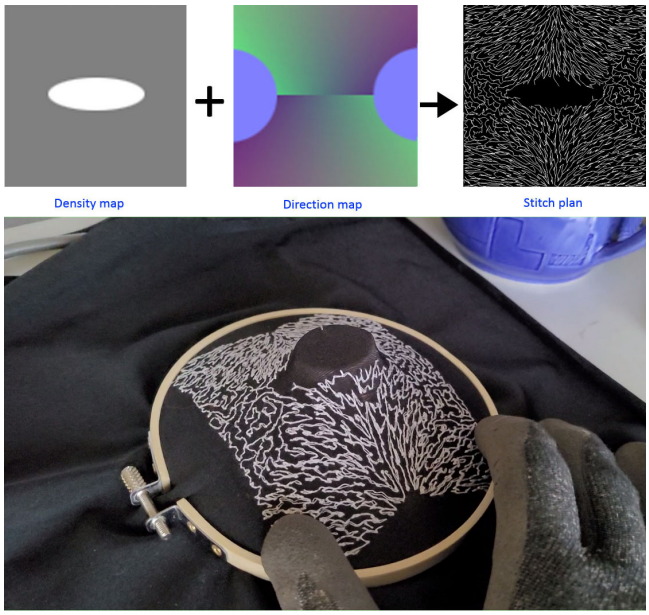

Figure 10: This example shows an anisotropic stitched fabric under a loaded condition. This physical print (bottom) shows that the 3D shape of this tiny "tent" is affected by the stitching.

time for a sling-style fabric seat. As a demonstration, we use the cone normal map which directs the weight to the center of the chair to support an upright posture. A similar example could be smart-stretch light-weight hammocks which conform to the shape of the body, such as providing a dedicated impression for the head and a more distributed pressure across the back. We note, such like custom fabric seats for wheelchair-bound individuals are not yet considered to date and could provide critical relief from bed sores and the like. We intend to explore this direction in the future.

Finally, we have a prototype bra worn by a mannequin in the right-hand photograph of Figure 9. This physical example exploits an anisotropic stitch pattern with a varying density plan that creates greater support and shaping. This exemplary application would benefit from expert designers that may also opt to add visual appeal derived from the stitching as well.

7 CONCLUSION

We introduce anisotropic stitching through computational stitch-level control and planning to manipulate the stretch properties of textiles. We propose an anisotropic design model derived from the density and direction of stitch augmentation of fabrics expressed through grayscale and normal image representations (in aggregate σ). This formulation makes way to direct design specification that is solved via a discretization and stitch plan algorithm based on a shortest path tree formulation. We provide a set of experiments that highlight stitch planning on manually designed stitch models, and experimentally validate the efficacy of our approach to change the stretch in physical fabric through an embroidery machine. We also perform comparisons with a state-of-the-art stitch planning approach and our proposed approach. Finally, we showcase examples that are justified and motivated from real-world applications.

While simple 2- and 4-way stretch is commonplace, our results showcase materials with arbitrary stretch in arbitrary directions across the fabric surface, the likes of which could support the next generation of "smart-stretch" fabrics. Of course, we are limited by what is possible within the physical properties – this limit reveals itself in the control for extreme anisotropic and density levels differences. Likewise, the constraints of the stitch plan as a sequential path of stitches also limits to how we may generate stitches and how well we can control the physical fabric. Even so, our experiments reveal that we are able to take a significant step forward using the proposed approach with both qualitative and quantitative evidence supporting this claim.

In contrast to the state of the art, we note the proposed use of SPT and trees in general gives us the important advantage of halting at a point and continuing exploration from another node, e.g., in case of a dead end. While a tree is not completely free from jump stitch issues, it does provide a greater degree of freedom in planning, since a vertex is allowed to have an arbitrary degree compared to a path where the connection is fixed to be at most two. The limitation of the tree structure in contrast is its traversal. However, by allowing multiple stitches at each point, the traversal of the tree becomes trivial (e.g. a depth-first walk will traverse each edge twice) and the resulting embroidery benefits from a stronger stitch path as well.

In terms of limitations, we have observed that our stitch plan solution can introduce artifacts such as long rows of stitches in extremely low uniform density regions. These rows focus extreme

loading when stretching the real fabric that introduces a vulnerability, such as stitch breakage. Also, in a way, such loading has a complete opposite effect than the intended because the region is largely inextensible in the direction of the stitch (before failure) which is not expected from a low density stitch area. Further, the coupling of off- and on-axis stitches controlled via ω rebalances the count of stitches in a given direction (e.g. when density is fixed) and this means we do not completely control the outcome stiffness. Figure 7 shows this, with the 0.75-blend example being lower than the 0.25-blend example.

As for future directions, the limitations need to be addressed and new areas are poised for exploration. For example, as our current implementation assumes a linear relationship between stitch amount and stiffness, it is likely that this needs deeper analysis and should be revisited, perhaps by employing a data-driven approach [Miguel et al. 2012; Wang et al. 2011]. To this end, we would like to study the density and direction relation in real stitched fabric samples obtained by our proposed approach and build feedback to obtain careful calibration for exact control. Another exciting direction for future work is in the production of shaping textiles, as suggested in the print from Figure 9 as well as the bra prototype. Shape control for textiles has numerous applications from body tensioning garments to furniture and textile tension structures (e.g., tents). Future work may also lead to the investigation of specific application domains such as sports or medical compression sleeves/socks/apparel indicating the need for improved pressure and collision handling/contacts [Harmon et al. 2008; Jiang et al. 2017; Li et al. 2018].

A main bottleneck we are facing is in the specification of design criteria that can further highlight the power and potential of the proposed technology. We believe that our overall pipeline could benefit from an inverse design optimization approach (such as [Skouras et al. 2014]) coupled with a physical simulator, in which the user specifies a problem to be solved and a σ map is automatically computed that can be given as input to our stitch planner. Likewise, more user-focused design input tool, e.g. sketch-based, could be a viable aid in this technology. In any case, as the technology promoted in this paper is fairly unexplored, we spent some effort to showcase a set of plausible real-world applications that exploit the anisotropic (and uniform) stitching to make a set of motivational examples.

ACKNOWLEDGMENTS

The authors would like to thank Connie Wen for her help in the analysis and data collection, Robin Ackermann for her help and input on embroidering, and Ella Moore and Michael Porter for their continuing support and contributions. This research was funded in part by the Clemson University Research Foundation.

REFERENCES

Richard Bellman. 1958. On a routing problem. *Quarterly of applied mathematics* 16, 1 (1958), 87–90.

James F Blinn. 1978. Simulation of wrinkled surfaces. *ACM SIGGRAPH computer graphics* 12, 3 (1978), 286–292.

Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2009. *Introduction to algorithms.* MIT press, Cambridge, MA, USA.

Edsger W Dijkstra et al. 1959. A note on two problems in connexion with graphs. *Numerische mathematik* 1, 1 (1959), 269–271.

Qiang Du, Maria Emelianenko, and Lili Ju. 2006. Convergence of the Lloyd algorithm for computing centroidal Voronoi tessellations. *SIAM journal on numerical analysis* 44, 1 (2006), 102–119.

Qiang Du, Vance Faber, and Max Gunzburger. 1999. Centroidal Voronoi tessellations: Applications and algorithms. *SIAM review* 41, 4 (1999), 637–676.

Ruslan Guseinov, Eder Miguel, and Bernd Bickel. 2017. CurveUps: Shaping Objects from Flat Plates with Tension-actuated Curvature. *ACM Transactions on Graphics* 36, 4 (2017), 64:1–64:12.

David Harmon, Etienne Vouga, Rasmus Tamstorf, and Eitan Grinspun. 2008. Robust Treatment of Simultaneous Collisions. *SIGGRAPH (ACM Transactions on Graphics)* 27, 3 (2008), 1–4.

James C Hateley, Huayi Wei, and Long Chen. 2015. Fast methods for computing centroidal Voronoi tessellations. *Journal of Scientific Computing* 63, 1 (2015), 185–212.

Lifeng He, Xiwei Ren, Qihang Gao, Xiao Zhao, Bin Yao, and Yuyan Chao. 2017. The connected-component labeling problem: A review of state-of-the-art algorithms. *Pattern Recognition* 70 (2017), 25–43.

Yuki Igarashi, Takeo Igarashi, and Hiromasa Suzuki. 2008a. Knitting a 3D model. In *Computer Graphics Forum*, Vol. 27. Wiley, Online Library, 1737–1743.

Yuki Igarashi, Takeo Igarashi, and Hiromasa Suzuki. 2008b. Knitty: 3D Modeling of Knitted Animals with a Production Assistant Interface.. In *EG Short Papers*. Eurographics, Citeseer, 17–20.

Alexandra Ion, Johannes Frohnhofen, Ludwig Wall, Robert Kovacs, Mirela Alistar, Jack Lindsay, Pedro Lopes, Hsiang-Ting Chen, and Patrick Baudisch. 2016. Metamaterial Mechanisms. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, new York, NY, USA, 529–539.

Chenfanfu Jiang, Theodore Gast, and Joseph Teran. 2017. Anisotropic elastoplasticity for cloth, knit and hair frictional contact. *ACM Transactions on Graphics* 36, 4 (2017), 1–14.

Mina Konaković, Keenan Crane, Bailin Deng, Sofien Bouaziz, Daniel Piker, and Mark Pauly. 2016. Beyond Developable: Computational Design and Fabrication with Auxetic Materials. *ACM Transactions on Graphics* 35, 4, Article 89 (2016), 11 pages.

Jie Li, Gilles Daviet, Rahul Narain, Florence Bertails-Descoubes, Matthew Overby, George E Brown, and Laurence Boissieux. 2018. An implicit frictional contact solver for adaptive cloth simulation. *ACM Transactions on Graphics* 37, 4 (2018), 1–15.

Jenny Lin, Vidya Narayanan, and James McCann. 2018. Efficient transfer planning for flat knitting. In *Proceedings of the 2nd ACM Symposium on Computational Fabrication*. ACM, Cambridge,MA,USA, 1–7.

Chenxi Liu, Jessica Hodgins, and James McCann. 2017. Whole-cloth quilting patterns from photographs. In *Proceedings of the Symposium on Non-Photorealistic Animation and Rendering*. ACM, Los Angeles, CA, USA, 1–8.

Yang Liu, Wenping Wang, Bruno Lévy, Feng Sun, Dong-Ming Yan, Lin Lu, and Chenglei Yang. 2009. On centroidal Voronoi tessellation—energy smoothness and fast computation. *ACM Transactions on Graphics* 28, 4 (2009), 1–17.

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 Trans. Graph.* 40, 4, Article 62 (July 2021), 16 pages.

Ali Mahdavi-Amiri, Philip Whittingham, and Faramarz Samavati. 2015. Cover-it: An Interactive System for Covering 3D Prints. In *Proceedings of the 41st Graphics Interface Conference (GI '15)*. Canadian Information Processing Society, Toronto, Ont., Canada, Canada, 73–80.

Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. 2016. Procedural Voronoi Foams for Additive Manufacturing. *ACM Transactions on Graphics* 35, 4 (2016), 44:1–44:12.

Jonàs Martínez, Haichuan Song, Jérémie Dumas, and Sylvain Lefebvre. 2017. Orthotropic k-nearest foams for additive manufacturing. *ACM Transactions on Graphics* 36, 4 (2017), 1–12.

James McCann, Lea Albaugh, Vidya Narayanan, April Grow, Wojciech Matusik, Jennifer Mankoff, and Jessica Hodgins. 2016. A Compiler for 3D Machine Knitting. *ACM Transactions on Graphics* 35, 4 (2016), 49:1–49:11.

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.

Ella Moore, Michael Porter, Ioannis Karamouzas, and Victor Zordan. 2018. Precision control of tensile properties in fabric for computational fabrication. In *Proceedings of the 2nd ACM Symposium on Computational Fabrication*. ACM, Cambridge, MA, USA, 1–7.

Vidya Narayanan, Lea Albaugh, Jessica Hodgins, Stelian Coros, and James Mccann. 2018. Automatic Machine Knitting of 3D Meshes. *ACM Trans. Graph.* 37, 3, Article 35 (Aug. 2018), 15 pages.

Vidya Narayanan, Kui Wu, Cem Yuksel, and James McCann. 2019. Visual knitting machine programming. *ACM Transactions on Graphics* 38, 4 (2019), 1–13.

Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic Textures for Additive Fabrication. *ACM Transactions on Graphics* 34, 4, Article 135 (2015), 12 pages.

Huaishu Peng, Scott Hudson, Jennifer Mankoff, and James McCann. 2016. Soft printing with fabric. *XRDS: Crossroads, The ACM Magazine for Students* 22, 3 (2016), 50–53.

Jesús Pérez, Miguel A Otaduy, and Bernhard Thomaszewski. 2017. Computational design and automated fabrication of kirchhoff-plateau surfaces. *ACM Transactions on Graphics* 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* 34, 4 (2015), 138:1–138:12.

Abhinit Sati. 2021. *A System for Programming Anisotropic Physical Behaviour in Cloth Fabric.* MS Thesis 3621. Clemson University, Clemson, SC.

Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to Control Elasticity in 3D Printing. *ACM Transactions on Graphics* 34, 4, Article 136 (2015), 13 pages.

Christian Schumacher, Steve Marschner, Markus Gross, and Bernhard Thomaszewski. 2018. Mechanical characterization of structured sheet materials. *ACM Transactions on Graphics* 37, 4 (2018), 1–15.

Mélina Skouras, Bernhard Thomaszewski, Peter Kaufmann, Akash Garg, Bernd Bickel, Eitan Grinspun, and Markus Gross. 2014. Designing Inflatable Structures. *ACM Transactions on Graphics* 33 (2014), 63:1–63:10.

Georgi Stoychev, Mir Jalil Razavi, Xianqiao Wang, and Leonid Ionov. 2017. 4D Origami by Smart Embroidery. *Macromolecular rapid communications* 38, 18 (2017), 6 pages.

Yunlong Tang, Guoying Dong, Qinxue Zhou, and Yaoyao Fiona Zhao. 2017. Lattice structure design and optimization with additive manufacturing constraints. *IEEE Transactions on Automation Science and Engineering* 15, 4 (2017), 1546–1562.

Huamin Wang, Ravi Ramamoorthi, and James F. O'Brien. 2011. Data-Driven Elastic Models for Cloth: Modeling and Measurement. *ACM Transactions on Graphics* 30, 4 (2011), 71:1–11.

Rundong Wu, Claire Harvey, Joy Xiaoji Zhang, Sean Kroszner, Brooks Hagan, and Steve Marschner. 2020. Automatic structure synthesis for 3D woven relief. *ACM Transactions on Graphics* 39, 4 (2020), 102–1.

Jonas Zehnder, Espen Knoop, Moritz Bächer, and Bernhard Thomaszewski. 2017. Metasilicone: design and fabrication of composite silicone with desired mechanical properties. *ACM Transactions on Graphics* 36, 6 (2017), 240.

Xiaoting Zhang, Guoxin Fang, Melina Skouras, Gwenda Gieseler, Charlie C. L. Wang, and Emily Whiting. 2019. Computational Design of Fabric Formwork. *ACM Trans. Graph.* 38, 4, Article 109 (July 2019), 13 pages.

Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. 2017. Two-scale topology optimization with microstructures. *ACM Transactions on Graphics* 36, 4 (2017), 1.