arXiv:submit/2056063 [cs.CG] 7 Nov 2017

Shape Representation by Zippable Ribbons

Christian Schüller
ETH Zürich

Roi Poranne
ETH Zürich

Olga Sorkine-Hornung
ETH Zürich

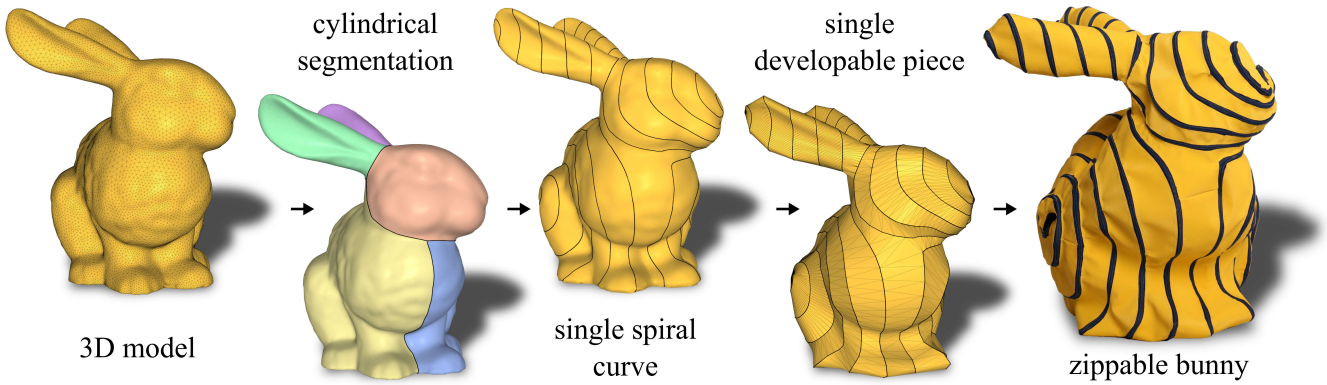

Figure 1: *The pipeline of our approach. Starting from a 3D model the user segments the shape into topological cylinders. Our algorithm then produces a curve on the shape that spirals along the cylinders. It proceeds to cut the shape along the curve and creates a single, developable ribbon-like surface. The ribbon is flattened into 2D and, based on the flattening, plans for laser cutting the ribbon out of fabric are generated. Finally, we stitch a single zipper to the boundary of the ribbon. Zipping it up reproduces the original shape.*

Abstract

Shape fabrication from developable parts is the basis for arts such as papercraft and needlework, as well as modern architecture and CAD in general, and it has inspired much research. We observe that the assembly of complex 3D shapes created by existing methods often requires first fabricating many small flat parts and then carefully following instructions to assemble them together. Despite its significance, this error prone and tedious process is generally neglected in the discussion. We propose an approach for shape representation through a *single* developable part that attaches to itself and requires no assembly instructions. Our inspiration comes from the so-called *zipit* bags [zipit 2017], which are made of a single, long ribbon with a zipper around its boundary. In order to "assemble" the bag, one simply needs to zip up the ribbon. Our method operates in the same fashion, but it can be used to approximate *any* shape. Given a 3D model, our algorithm produces plans for a single 2D shape that can be laser cut in few parts from flat fabric or paper. We can then attach a zipper along the boundary for quick assembly and disassembly, or apply more traditional approaches, such as gluing and stitching. We show physical and virtual results that demonstrate the capabilities of our method and the ease with which shapes can be assembled.

1 Introduction

Representing shapes using developable surfaces is a problem with numerous applications, ranging from recreational activities like papercraft fabrication, to large scale industrial design and modern architecture. Our interest in this problem is inspired by a product commercially known as the *zipit* bag [zipit 2017]. This bag is made from a single, long piece of fabric ribbon with a zipper attached around its boundary. When zipped up, the ribbon wraps around to form a simple bag. It takes only a few seconds to assemble and disassemble it, and no instructions are necessary. The simplicity of this concept and the oddly satisfying sensation of seeing it taking form when being zipped up immediately propelled us to ask whether this idea can be generalized to *arbitrary* shapes. To this end, we devise a computational method to create a suitable developable shape that approximates a given target 3D model when zipped up; see e.g. Fig. 1. Our approach generalizes the simple straight ribbons that make up the *zipit* bags by allowing the ribbons to turn, have varying width, and branch (see Fig. 2). Merely finding a 2D shape that approximately reproduces a 3D shape when zipped up is not very challenging. One approach would be to take a 3D mesh and cut it into a single triangle strip [Rossignac 1999; Eppstein and Gopi 2004]. However, in addition to being mesh dependent, the resulting strip would have many sharp turns that make zipping up difficult, if not impossible, and the strip's width would be uneven, which can be visually unappealing. Therefore, we postulate the following desirable "regularity" properties for a zippable 2D ribbon shape:

- The ribbon should curve as little as possible.
- It should have as constant as possible width.

Figure 2: *Our design for a zippable star pillow, made of two differently colored fabrics, flatly attached together.*

The challenge we aim to solve is to compute a single, flat piece that satisfies these two properties and approximates the input shape well when zipped up. We observe that the two properties are nearly trivial to achieve when the target 3D model is a cylinder: it is easy to trace a spiraling curve on a cylindrical surface such that if cut along that curve, the resulting shape is a straight ribbon with fixed width.

For general target surfaces, our approach is based on first decomposing the shape into topological cylinders, and then mapping them onto cylindrical domains with low isometric distortion and in a seamless manner. We then draw "perfect" spirals on the cylinders and map these spirals back onto the input shape. Since the mappings minimize isometric distortion, the mapped curves tend to exhibit low curvature and low variation of the distance between windings. Inspired by [Zhao et al. 2016] we connect the spirals on the different cylinders into a single, long *Fermat spiral* (see inset). The shape is then cut along the computed curve to create a single, possibly bifurcated, but not yet developable ribbon. A simple remeshing process then transforms this ribbon into a developable one, which can be trivially unfolded onto the 2D plane to create the cutting pattern. It is possible that some strands of the flat ribbon overlap in the plane, and the pattern might take up too much space to fully fit onto the laser cutter bed; in both cases we simply divide the ribbon into a few separate pieces that are later sewn back together, before attaching the zipper (see Fig. 17).

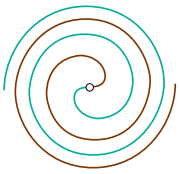

In addition to the assembly process being easy and entertaining, our fabrication process has distinct advantages over papercraft and many other similar methods in this domain. Our assembly is *linear*, i.e., at every instant of the assembly process, the next action is unique and unambiguous, and it requires *no instructions*. In contrast, papercraft and related approaches typically require certain assembly order, especially for the final stages, where all the parts have to come together for the shape to close up; the makers usually must refer to a manual, find the next piece and understand how to attach it to their work. In case our method generates self-overlapping strands of ribbon that must be severed in order to laser-cut them from fabric, sewing the pieces together is simple since both ends are flat, perfectly matching in length and only require flat stitching along a straight line. Also sewing the zipper to the ribbon only requires working on flat sheets of fabric and is therefore straightforward and rather fast, as confirmed by the seamsters and seamstresses we employed. In contrast, attaching multiple parts in papercraft or sewing plush toys from multiple charts usually cannot be done in a flat position, and may necessitate more time and skill for large models.

We demonstrate our technique on a number of various shapes, see e.g. the model in Fig. 3. We show virtual results and physically fabricated objects, which are assembled and disassembled simply by slider a zipper.

2 Related work

Our work relates to the general field of computational fabrication and digital geometry processing. We briefly review the most relevant previous works below.

Papercraft and needlework. Some objects, typically of limited size and/or with certain constraints on the shape, can be directly manufactured in one piece using e.g. 3D printing or CNC milling. However, a large number of approaches propose to create shapes composed of several individually fabricated parts. In this space, the most related approaches to ours are papercraft based on cutting and gluing [Mitani and Suzuki 2004; Shatz et al. 2006; Massarwi et al. 2007; Straub and Prautzsch 2011; Takahashi et al. 2011], and manufacturing by sewing [Julius et al. 2005; Mori and Igarashi 2007; Igarashi and Igarashi 2008; Igarashi et al. 2009; Wang 2010; Mahdavi-Amiri et al. 2015]. Both types of methods require the approximation of a given 3D shape by pieces of developable surfaces (in the case of paper) or highly stretch-resistant material (in the case of fabric). This is typically achieved by segmenting or cutting the shape into parts with low Gaussian curvature and parameterizing each part onto the plane. Alternatively, the shape can be a priori modeled or approximated as a piecewise developable surface, which is a current topic of active research [Kolmianič and Guid 2002; Rose et al. 2007; Kilian et al. 2008; Liu et al. 2009; Zeng et al. 2012; Chandra et al. 2015; Tang et al. 2016]. The assembly by gluing or sewing the pieces together requires precision and carefully following the instructions. In contrast, in our case, it is mostly a straightforward and even mindless task.

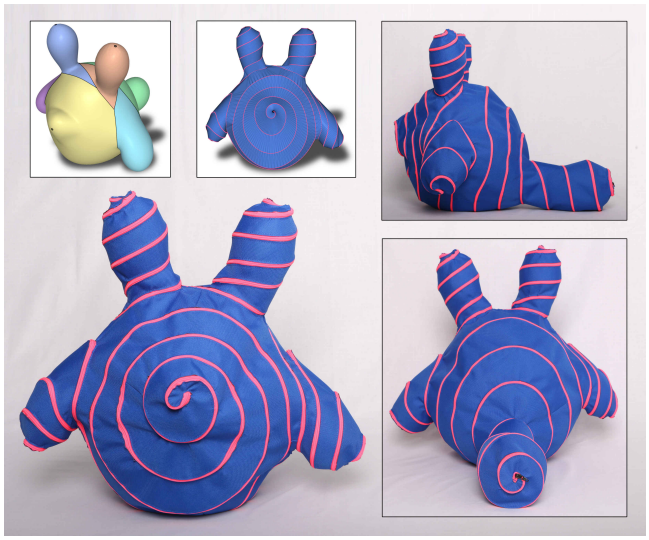

Figure 3: *A zippable model of a character. The zipper starts at the tail and spirals around all extruding parts until it ends at the tip of the nose.*

Soft materials. Designing for fabrication using soft material like fabric or rubber is challenging, since such materials easily deform under stress and gravity. The main goal is to predict the deformation and solve an inverse problem, so that the fabricated model assumes the desired shape when subjected to the stress. Examples include the generation of plush toys [Mori and Igarashi 2007; Igarashi and Igarashi 2008], inflatable structures [Skouras et al. 2012; Skouras et al. 2014] and rubber objects [Bickel et al. 2010; Skouras et al. 2013; Chen et al. 2014], among others. As mentioned, although we use fabric in our results, the behavior is more reminiscent of papercraft, because we use a nearly inextensible cloth and the zipper itself is rather stiff and completely inextensible. Regardless of the type of fabric used, our zipper based approach has one decisive advantage: The whole fabrication process can be done entirely in the flat. In contrast, previous methods result in spaciously curved pieces once they get connected to each other, which makes sewing or gluing much harder as confirmed by professional tailors. The final assembly by zipping up needs no instructions, is fast and fun to do.

Parameterization. Our approach relies on mesh parameterization, which is an extensively studied topic, see e.g. the survey in [Hormann et al. 2007]. The more recent relevant parameterization literature is presented in [Smith and Schaefer 2015; Kovalsky et al. 2016;

Rabinovich et al. 2017]. Our method introduces a *new type* of global parameterization, which extends cylindrical parameterization [Tarini 2012]. Global parameterization is primarily used for quad meshing, where parts of seams in the parameter domain are related to each other by a rotation of integer multiples of $\pi/2$. A recent review can be found in [Bommes et al. 2013]. In our specific case, we require a different form of seamless mapping, based on cylindrical domains [Ray et al. 2006; Knöppel et al. 2015]. Our main inspiration is [Kälberer et al. 2011], where the authors propose an approach for drawing stripes on *tubular* shapes that can be used for generating textures. Their approach is based on a seamless parameterization aligned with a 2-RoSy field obtained from the principal curvature directions. This causes problems near umbilical points, where the principal directions are unstable. In contrast, we employ parameterization based on distortion minimization, which avoids this problem, and generally leads to less distorted mappings [Myles and Zorin 2012].

3 Method

Given a mesh representing the 3D object, our goal is to generate a single, flat, possibly branching shape, which from now on we refer to simply as *ribbon*, that approximates the object when "zipped-up". We can rigorously define zipping-up as an isometric deformation of the flat shape into a 3D shape such that the boundary exactly overlaps with itself. However, we assume that our intention is clear and avoid mathematical rigor at this point. In addition to being "zippable", we wish to enable some creative control by allowing the user to define where the zipper should pass and how it should be oriented or aligned.

Our method is based on the observation that it is trivial to generate a perfect spiral on a cylinder. Assume the cylinder is a topological annulus (i.e., has no caps). We can cut it from its top boundary loop to the bottom one and unfold it to a rectangular shape, where the two boundary curves become the top and bottom edges, and the cut becomes the two side edges. We then place "copies" of the unfolded cylinder side-by-side, and draw a straight line from the bottom corner of the leftmost edge to the top corner of the rightmost edge. Overlaying the copies on top of each other creates several disconnected, parallel line segments on the parameterization of the cylinder, and by folding it back to a cylinder, these segments transform into a perfect connected spiral. Its number of turns depends on the number of copies we made. See Fig. 4 for an illustration.

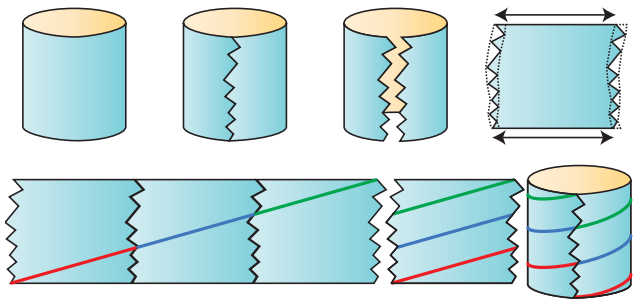

Figure 4: *Drawing a spiral on a cylinder can be done by cutting the cylinder from the top boundary to the bottom one and unfolding it to the plane. We then place copies of the flattened cylinder and draw a straight line that passes from the bottom leftmost corner to the top rightmost one. Overlaying the copies and folding back to a cylinder creates a spiral, where the number of windings is equal to the number of flattened copies.*

Figure 5: *A zippable shape of a kitten. Since it is topologically equivalent to a torus, an additional cut is needed (bottom left inset: marked in red where the tail touches the head). Another zipper could be used to close up this cut, but we opted for using Velcro instead.*

The same approach, termed *cylindrical parameterization*, can be applied to general cylinder-like shapes, which we colloquially continue calling "cylinders". We start in the same manner by cutting the shape from one boundary loop to the other. The cut shape is then mapped to the plane by a distortion minimizing parameterization with *seam constraints*, which force the two sides of the cut to match like puzzle pieces. Minimizing distortion is necessary for the straight line in 2D to be mapped to a smooth and uniform spiral on the surface in 3D. The case of a more complex shape is slightly more involved: we decompose the shape into cylindrical parts and use a global parameterization scheme to smoothly map all the parts to cylinders. We discuss this in more detail in Sec. 3.2.

Once the mapping is found, we turn to designing spirals on the cylinders. The main challenge is to synchronize the spirals, such that one spiral ends where another begins, resulting in a single, long, spiraling curve on the surface. The final task is to cut the surface along the curve and remesh it such that the result is developable. It is then trivially flattened to generate the final 2D shape of the zippable ribbon.

To summarize, the design phase of our method consists of four stages:

1. Decomposition into cylindrical parts.
2. Seamless, global parameterization of the cylinders.
3. Spiral generation.
4. Cutting along the spiral, remeshing and flattening.

See Fig. 1 for a graphical overview.

3.1 Decomposition into cylindrical parts

We decompose the input shape S into topological cylinders $\mathcal{C}_i$, i.e., 2-manifolds with two boundary loops, more commonly termed *annuli*. The decomposition plays a substantial role in the final appearance of the spiral, since the resulting curve is aligned to the boundaries of the cylinders. It enables a flexible interface for artistic exploration and is achieved by interactively tracing the boundaries of the segmentation using our software. Alternatively, more automatic ways of cylindrical decomposition such as [Zhou et al. 2015; Livesu et al. 2017] could be applied, but we have not explored this option. Since surfaces without boundaries cannot be decomposed into topological cylinders, we also allow the user to cut the shape open and place new boundaries, e.g., small circular holes or curves on the shape

that act as a single cylinder boundary (see Fig. 6). We distinguish between *transition* boundaries separating two adjacent cylinders and *open* boundaries, which are actual boundaries of the shape. In the spiral design stage, the curve will pass through transition boundaries, *spiraling in* toward an open boundary and then spiraling out toward another open boundary. This behavior is reminiscent of the Fermat spiral, which we discuss in Sec. 3.3.

3.2 Seamless parameterization

Once the cylindrical decomposition of S into parts $\mathcal{C}_i$ is available, we proceed to compute the parameterization. This step determines how the equally spaced, straight lines in the 2D parameter domain transform into spirals on S. To obtain a spiral that maintains the even spacing on the 3D shape, the parameterization must minimize isometric distortion. Additionally, we require the parameterization to be bijective in order for the mapping from the 2D lines to the 3D curve to be well defined. Fig. 7 compares curves generated with the initial (i.e. suboptimal) and optimized parameterization.

We assume that S has been cut along the boundaries of the $\mathcal{C}_i$'s, and each cylinder is cut from one boundary to the other, analogous to the example of one cylinder (see Fig. 4 top row). The edges and vertices along the cuts are duplicated, and we keep correspondences between the copies. We generate a seamless bijective parameterization of each $\mathcal{C}_i$, with seamless transitions between adjacent $\mathcal{C}_i$'s. We first explain the case where there is only one cylinder, and the general case of multiple cylinders immediately follows.

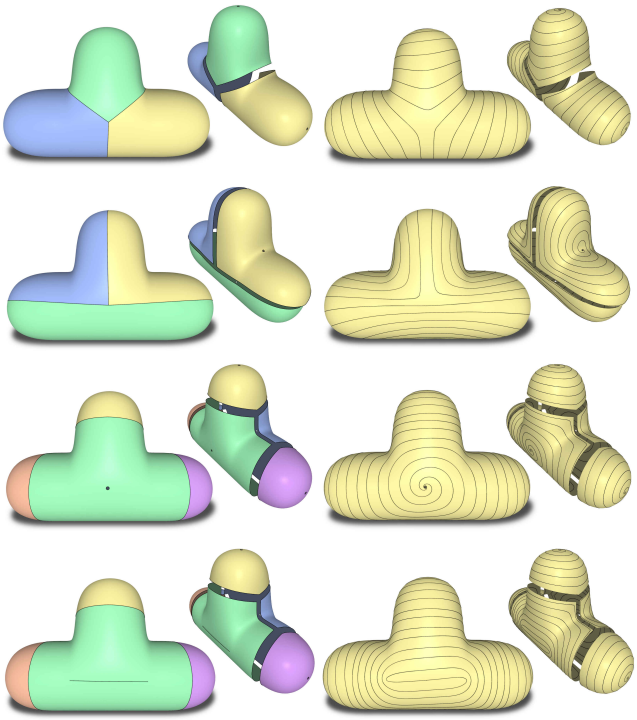

Figure 6: *Different cylindrical segmentations of a T shape. Each cylinder has one transition boundary and one open boundary. Note the small holes in the middle of the colored parts. These are the open boundaries for the corresponding cylinders. In the last row we show an example of a straight curve cut, serving as the open boundary of the green cylinder. Compare the resulting spiral to the segmentation in the third row to see the effect of the curve cut.*

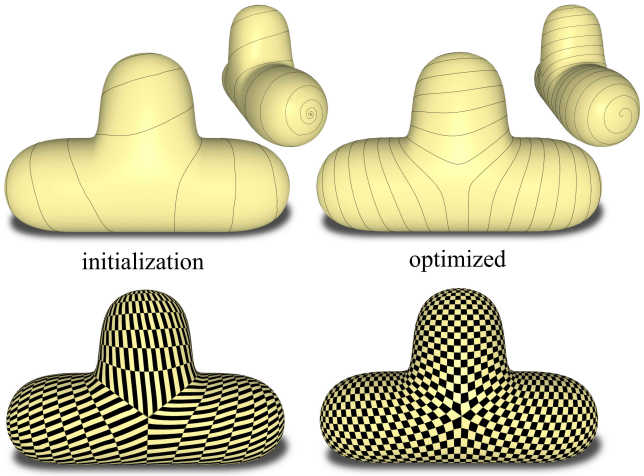

Figure 7: *Comparison between the curve obtained before and after minimizing isometric distortion of the parameterization. Note that the non-optimized spirals have a much greater variation in the spacing between the windings. We show a uniform grid texture to illustrate the difference in distortion.*

Minimizing isometric distortion. An isometric distortion measure quantifies the difference between a given flattening of a shape and a perfect isometry; most formulations define it as a sum of the distortions of individual triangles. In this paper we use the recently proposed *symmetric Dirichlet* distortion measure; see [Smith and Schaefer 2015; Kovalsky et al. 2016; Rabinovich et al. 2017] for details.

We denote the coordinates of a vertex in the parameter domain by $\mathbf{x} = (x, y)$ and stack all the coordinates in a vector $\mathbf{X}$. The distortion of a triangle t is a function of the positions of its vertices in the plane. We denote the symmetric Dirichlet measure of triangle t by $D_t(\mathbf{X})$. Then the optimization problem to solve is

$$\operatorname*{argmin}_{\mathbf{X}} \sum_{\text{triangle } t} A_t D_t(\mathbf{X}), \tag{1}$$

where A_t is the area of t in the original mesh. In our work, we use a modified Newton's method [Shtengel et al. 2017] with a feasible starting point to solve this problem. Since we add several constraints in the following, we defer a detailed discussion to the end of this section.

Seamless cylindrical parameterization. To map a single topological cylinder $\mathcal{C}_i$ to the plane with minimal distortion in a seamless manner, it is cut to form a disk topology and then parameterized while adhering to seam constraints. The role of the seam constraints is to ensure that the parameterization is invariant to the cut [Myles and Zorin 2012]. In the cylinder case, the seam constraints call for each edge on one side of the seam to be a translation of its twin edge on the other side of the seam. More precisely, assume the cut contains n consecutive vertices and let $\mathbf{x}_j^{\mathrm{L}}$ and $\mathbf{x}_j^{\mathrm{R}}$, $j = 1, \ldots, n$, be the two copies of each vertex in the parameterization (superscripts L, R stand for Left and Right). Then the cylindrical seamlessness constraints are

$$[\mathrm{Cyl}(\mathcal{C}_i)] \quad \mathbf{x}_j^{\mathrm{L}} - \mathbf{x}_{j-1}^{\mathrm{L}} = \mathbf{x}_j^{\mathrm{R}} - \mathbf{x}_{j-1}^{\mathrm{R}}, \quad j = 2, \ldots, n. \tag{2}$$

We use the differential form of the seam constraints in order to avoid introducing auxiliary variables. The equivalent positional form is $\mathbf{x}_j^{\mathrm{L}} = \mathbf{x}_j^{\mathrm{R}} + \mathbf{t}$, $j = 1, \ldots, n$, where $\mathbf{t}$ is an unknown offset (the same for all vertices). For conciseness, we refer to the set of constraints in (2) as $\mathrm{Cyl}(\mathcal{C}_i)$ for a given $\mathcal{C}_i$, or Cyl in general.

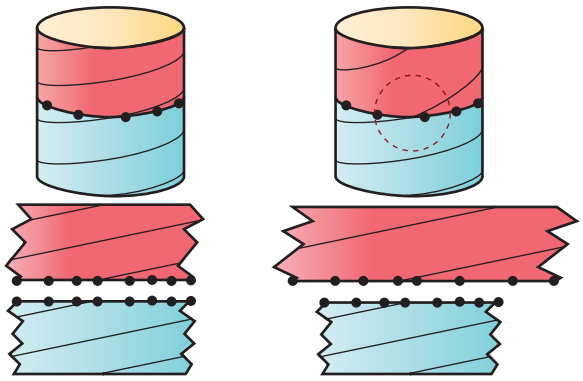

Figure 8: *Inter-cylinder constraints ensure that the transitions between cylinders are smooth (left). Note the kink that appears in the curve when these constraints are missing (right).*

Cylinder boundary constraints. In addition to the cylinder seam constraints Cyl, we also require the boundary loops of the cylinders to be mapped to straight lines. This serves two purposes: First, together with the Cyl, it guarantees bijectivity, and second, it allows for a better surface coverage by the spiral. Indeed, when the boundaries are not kept straight and allowed to "spill out" in the 2D domain, the spilled region is not covered by the spiral (see illustration in the inset and Fig. 22). Due to Cyl, the straight lines of the boundaries must be parallel, hence, without loss of generality, we can make them parallel to the horizontal x-axis. Let y_k^{Top}, $k = 1,\dots,m^{\text{Top}}$ and y_l^{Bot}, $l = 1,\dots,m^{\text{Bottom}}$ be the y-coordinates of the vertices of the top and bottom boundaries in the parameter domain. We again use the differential form for the straight line constraints, given by

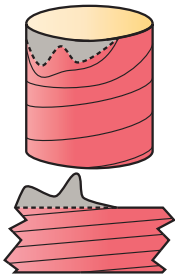

$$[\,\text{Str}(\mathcal{C}_i)\,] \quad \begin{aligned} y_k^{\text{Top}} - y_{k-1}^{\text{Top}} &= 0, \quad k = 2,\dots,m^{\text{Top}} \\ y_l^{\text{Bot}} - y_{l-1}^{\text{Bot}} &= 0, \quad l = 2,\dots,m^{\text{Bottom}}. \end{aligned} \tag{3}$$

We denote the constraints of each $\mathcal{C}_i$ in (3) by $\text{Str}(\mathcal{C}_i)$, and the entire set of these constraints as Str. The constraints Cyl and Str together already result in a nice spiral on each $\mathcal{C}_i$ separately. However, without dedicated treatment, there could be a visible kink in the spiral when transitioning between $\mathcal{C}_i$'s (see illustration in Fig. 8). Indeed, if copies of the same edge on a transition boundary are parameterized to two edges with different size and direction, then the parameterization will not appear smooth across that edge. We handle this issue in the following.

Inter-cylinder seamlessness. In order for the transition between $\mathcal{C}_i$'s to appear smooth, we apply seam constraints on cuts between each pair of neighboring $\mathcal{C}_i$'s. These constraints enforce a rigid transformation between the two sides of each seam (see illustration in Fig. 8). Since we have the freedom to define the exact transformation, we choose a rotation by π. Thus, for every two neighboring $\mathcal{C}_P$ and $\mathcal{C}_Q$, we let $\mathbf{x}_r^P$ and $\mathbf{x}_r^Q$, $r = 1,\dots,s$, be the coordinates of the two copies of each vertex along the seam between $\mathcal{C}_P$ and $\mathcal{C}_Q$. Then the inter-cylinder seam constraints can be written in differential form as

$$\big[\,\text{Int}(\mathcal{C}_P,\mathcal{C}_Q)\,\big] \quad \mathbf{x}_r^P - \mathbf{x}_{r-1}^Q = \mathbf{x}_{r-1}^P - \mathbf{x}_r^Q, \quad r = 2,\dots,s. \tag{4}$$

We denote the constraints in (4) for each pair $\mathcal{C}_P,\mathcal{C}_Q$ by $\text{Int}(\mathcal{C}_P,\mathcal{C}_Q)$.

Global cylindrical parameterization. With all the types of constraints defined, we can formulate the optimization problem for parameterizing the surface:

$$\begin{aligned} \underset{\mathbf{X}}{\text{argmin}} \quad & \sum_{\text{face } t} A_t\, D_t(\mathbf{X}) \\ \text{s.t.} \quad & \text{Cyl}(\mathcal{C}_i), && \forall \mathcal{C}_i \\ & \text{Str}(\mathcal{C}_i), && \forall \mathcal{C}_i \\ & \text{Int}(\mathcal{C}_P,\mathcal{C}_Q), && \forall \mathcal{C}_P,\mathcal{C}_Q \text{ neighbors.} \end{aligned} \tag{5}$$

Eliminating degrees of freedom. The constraints in (5) are all sparse linear homogeneous equalities, and they are overdetermined. For example, it is unnecessary to require all of the top and bottom boundaries to be on straight lines: it is sufficient to require this only for half of them, and the other half then must lie on straight lines due to the Int constraints. Similarly, the y part of (5) is redundant due to the Str constraints. We remove these redundant constraints using Gaussian elimination.

Initialization. Our optimization is based on Newton's method and requires a feasible starting point with no triangle flips. We use Tutte's embedding with uniform weights, which guarantees bijectivity if the boundary is convex. We can therefore map each cylinder to a rectangle in the plane, where the top and bottom boundaries are parallel to the x axis, in order to satisfy Str. We set the height of each rectangle to have the length of the cylinder boundary in order to get a more isometric initial guess. In order to satisfy Int we require each edge of a transition boundary to have the same length in its parameterization. We can do the same for the cylinder boundary edges, which then completely determines the boundary. However, we note that we can in fact use Cyl, as they are in the form of an orbifold Tutte's embedding (see [Aigerman and Lipman 2015]), instead of specifying vertex positions directly, which results in a slightly less distorted initial guess.

Optimization. We use a modified Newton's [Shtengel et al. 2017] method with linear constraints to solve (5). We use the line search method suggested in [Smith and Schaefer 2015], which guarantees that no triangle flips are introduced during optimization. Once the optimization converges, we continue to the next step, which is the generation of the spirals. We show an example of the parameterization of the T shape in Fig. 9.

3.3 Spiraling curve design

With the global seamless parameterization of the cylindrical decomposition available, the next stage in our algorithm is to generate a spiraling curve on the shape by drawing straight lines on the flattening and lifting them back into 3D by the inverse mapping. A spiral on a single $\mathcal{C}_i$ can be created as discussed in the beginning of Sec. 3, by drawing a straight line on the cylinder's parameterization. For the case of multiple $\mathcal{C}_i$'s, in order to obtain a single continuous curve, we must make sure that the spirals of the individual $\mathcal{C}_i$'s connect. We make the following simplifying assumptions:

1. The curve starts and ends at an open boundary, passes through each transition boundary and traverses all other open boundaries via Fermat spirals.
2. Each $\mathcal{C}_i$ has one transition boundary and one open boundary (see Sec. 3.1).
3. The curve enters and exits each $\mathcal{C}_i$ exactly once.

Note that a Fermat spiral requires drawing *two* lines on a cylinder (see Fig. 10). The assumptions above mean w.l.o.g. that a curve must start at an open boundary in $\mathcal{C}_1$, then cross a transition boundary and travel to the adjacent $\mathcal{C}_2$. Then touch its open boundary and

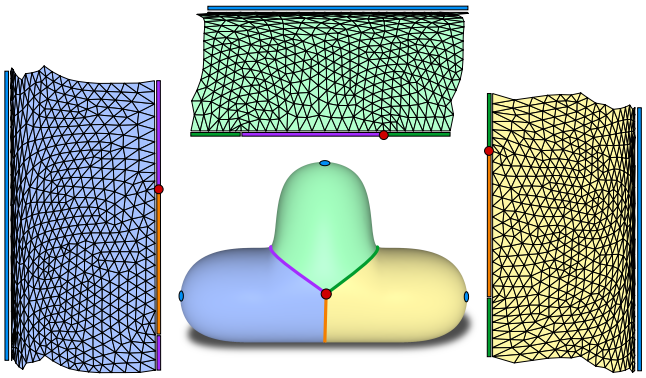

Figure 9: *The parameterization of the three topological cylinders of the T shape. See Fig. 6 for another perspective of the same segmentation. We mark copies of the same transition boundaries by matching colors, and the red dot represents one of the two points where all cylinders meet. We remark that this point has no particular significance and is only there as a visual guide. The blue sides of each flattening and the corresponding ellipses represent the open boundaries, while the unmarked sides represent the constrained cylinder seams.*

leave it at another point to continue through the other transition boundary to $\mathcal{C}_3$. This continues until all cylinders are traversed, and the curve ends at the open boundary of the last $\mathcal{C}_i$. See Fig. 11 for an illustration.

There are several choices we let the user make. The first one is the traversal order of the $\mathcal{C}_i$'s. The second one is the location where the curve passes between $\mathcal{C}_i$ and $\mathcal{C}_{i+1}$. Finally, the user can prescribe the number of windings of the spiral on each $\mathcal{C}_i$. See Fig. 12 for an example of the different choices. These choices impact the final look of the curve and should be based mostly on artistic considerations. Having as-uniform-as-possible spacing between the windings appears to be the best for practical fabrication reasons. For a regular spiral, even spacing is ensured by the low-distortion parameterization. The Fermat spirals require a bit more thought, since two lines are needed per $\mathcal{C}_i$. Ideally, we would like to draw the two lines parallel, such that their copies are equally spaced. This means that the distance between their intersections with the boundary is half its width (Fig. 10). This is not always possible though, since the interfaces between $\mathcal{C}_i$'s might not allow it. To illustrate this, consider a single $\mathcal{C}_i$, which the curve traverses after $\mathcal{C}_{i-1}$ and before $\mathcal{C}_{i+1}$ (Fig. 11). The $\mathcal{C}_i$ has an open boundary, and shares two transition boundaries with $\mathcal{C}_{i-1}$ and

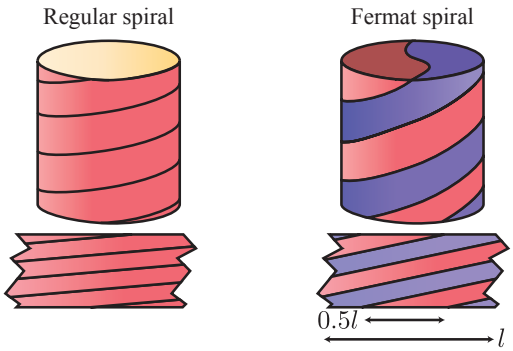

Figure 10: *Illustration of a Fermat spiral on a cylinder. The cap of the cylinder on the right represents the open boundary where the center of the Fermat spiral appears.*

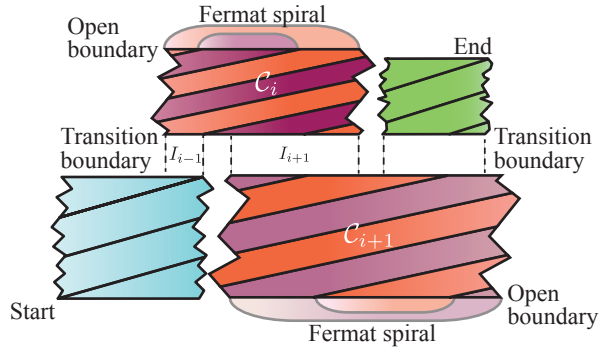

Figure 11: *Illustration of a spiraling curve traversing several cylinders. We mark the interfaces of $\mathcal{C}_i$ by I_{i-1} and I_{i+1} (see Sec. 3.3).*

$\mathcal{C}_{i+1}$ on its other boundary. The two boundaries are parameterized to two straight line segments of the same length l. For the open boundary, we are free to pick any two points with a distance $0.5l$ between them. Let us assume that the x coordinates of the opposite boundary are the interval $I = [0, l]$. Then the interfaces with $\mathcal{C}_{i-1}$ and $\mathcal{C}_{i+1}$ are sub-intervals $I_{i-1}, I_{i+1} \subset I$. We would like the curve to cross the boundary at two point x_1, x_2 such that $|x_2 - x_1| = 0.5l$. Clearly this is not always possible, but in practice, we did not find this to be an issue, since the user is given explicit control over these passing points. A possible future work is to solve an optimization problem that finds the best crossing points within those intervals when a problematic situation is encountered.

3.4 Cutting, remeshing and flattening

The goal of the final stage of our algorithm is to take the curve designed in the previous stage and look for a developable surface that has the curve as its boundary and approximates the shape well. This is a challenging task in general (see e.g. [Rose et al. 2007; Tang et al. 2016]), but it is somewhat simpler in the discrete setting. It is well known that a triangle mesh in 3D that has no internal vertices, i.e., all its vertices are boundary vertices, is developable. Finding a developable triangulation is still a difficult problem, where the challange lies in finding a meshing that appears smooth. Mitani and Suzuki [2004] proposed to use edge collapse and vertex removal operations until no internal vertices exist, and then to apply edge flip

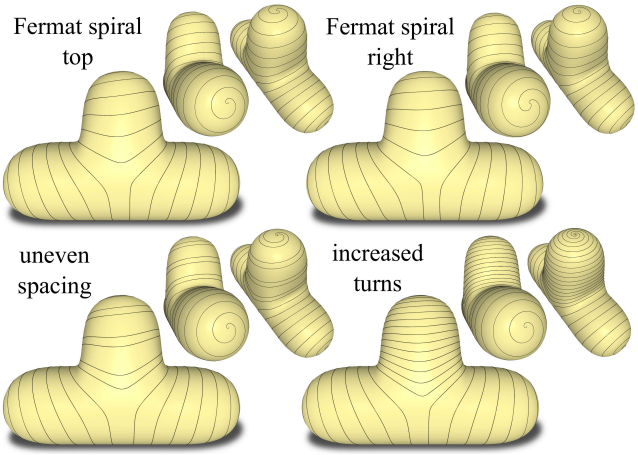

Figure 12: *We show several possible spiraling curves on the T shape. The design is up to the user's artistic choices.*

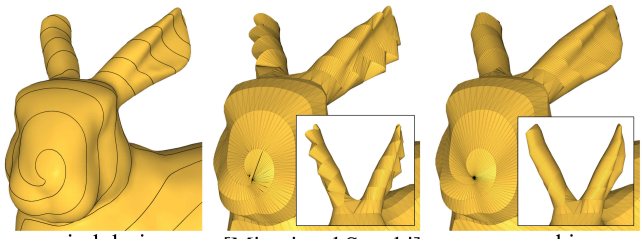

Figure 13: *A comparison of our simple ribbon meshing approach to the method in [Mitani and Suzuki 2004]: The greedy edge flipping step can generate non-optimal triangle fans (middle column), whereas ours results in a smoother and better approximation of the original surface. But it may also fail to create a good meshing (see Fig. 21).*

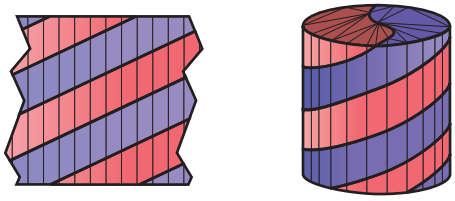

Figure 14: *Illustration of a remeshing to a developable ribbon. Points on adjacent lines with the same x coordinate are connected by an edge.*

operations in order to improve the smoothness of the triangulation. We have implemented their method, but observed that the greedy edge flipping can introduce triangle fans near narrow curve turns (Fig. 13). We instead propose a simple approach based on the parameterization we already have from previous stages. The idea is to define correspondences between points on adjacent windings of the spiral, which act as rulings of a developable surface. The simple correspondence we choose is based on the x coordinates of the lines in each $\mathcal{C}_i$'s parameterization (Fig. 14). We sample these lines uniformly and connect two samples in adjacent windings if their x coordinate is the same. We use Triangle [Shewchuk 1996] to complete the triangulation where there is no natural correspondence, that is, around the interfaces between cylinders. Although this simple approach is not always optimal (see Fig. 21), we found it to be sufficient in most cases. Once the ribbon is triangulated, we can unfold it to the plane (see Fig. 17). We cut the ribbon into few smaller pieces in order to utilize the cutting area better and resolve overlaps.

4 Results and discussion

Implementation details. We implemented all parts of our algorithm in C++, except the parameterization, which was implemented in MATLAB. The modified Newton's method converges within at most 15 iterations and and takes about 20 seconds for a mesh with 30k triangles on our Xeon CPU E5-2650 v2, 64 GB RAM machine.

Approximation capabilities. Naturally, the thinner the designed ribbon, the better its approximation power. However, this also results in a longer ribbon, prolonging the fabrication. The decision regarding the ribbon's length and width is left to the user. See Fig. 15 for different widths and approximations of the T shape.

Papercraft comparison Our method can also be used to create papercraft models. In general, it produces fewer initial pieces then

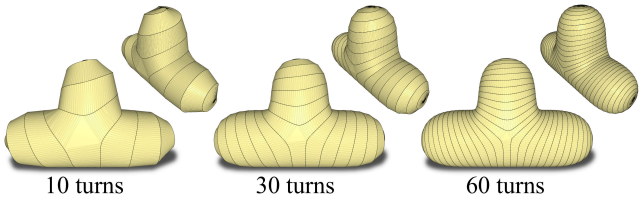

Figure 15: *Running our method on the T shape with different spiral densities.*

related methods and the assembly by gluing is straight forward and does not need elaborate instructions. See Fig. 16 for a comparison of our bunny result fabricated from paper with the method of [Mitani and Suzuki 2004].

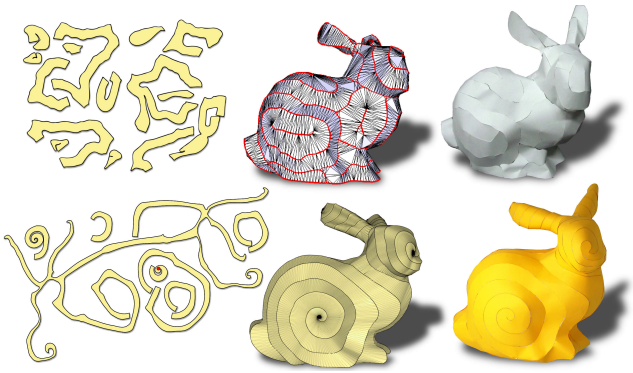

Figure 16: *Our laser cut plans of the bunny with 7 pieces (bottom row) for a single developable piece that can be assembled by linearly gluing its border, starting at the designated red point, compared to [Mitani and Suzuki 2004] (taken from their paper) with 15 pieces which require more detailed instructions to be glued together (top row).*

Fabrication process. We fabricated six of our designs in fabric, see Figures 2, 3, 5, 17, 19, 20. Fig. 17 shows all the steps of our method with our simplest model, the T shape. All fabricated results have a zipper length of 10 meters, except the star, which has 11 meters. The models have a maximum size of around 50 cm. We asked professional tailors to do the sewing work, and it took them between 5 to 6 hours for each model. Even though they had no experience with this special kind of fabrication, there were no problems in attaching the zippers to the cut fabric pieces thanks to the linearity of the assembly method. To guarantee correct alignment of the zipper to the fabric, we place markers every 5 cm both on the zipper and along the laser cut piece. We note that care must be taken to align the teeth correctly when attaching the slider. Misalignment leads to an offset in the boundary correspondence, which results in an incorrect zip-up. However, this alignment needs to be done only once, and the markers can be taken as a guidance.

Further examples. We experimented with our system and created several additional examples. In Fig. 22 we show a heart shape with a nontrivial zig-zag segmentation. In Fig. 23 we show a result obtained on a twisting tube that changes width.

5 Conclusion

We presented an approach for shape representation using a single developable piece of fabric. We show several examples to demonstrate the power and generality of our approach. Currently, we do not

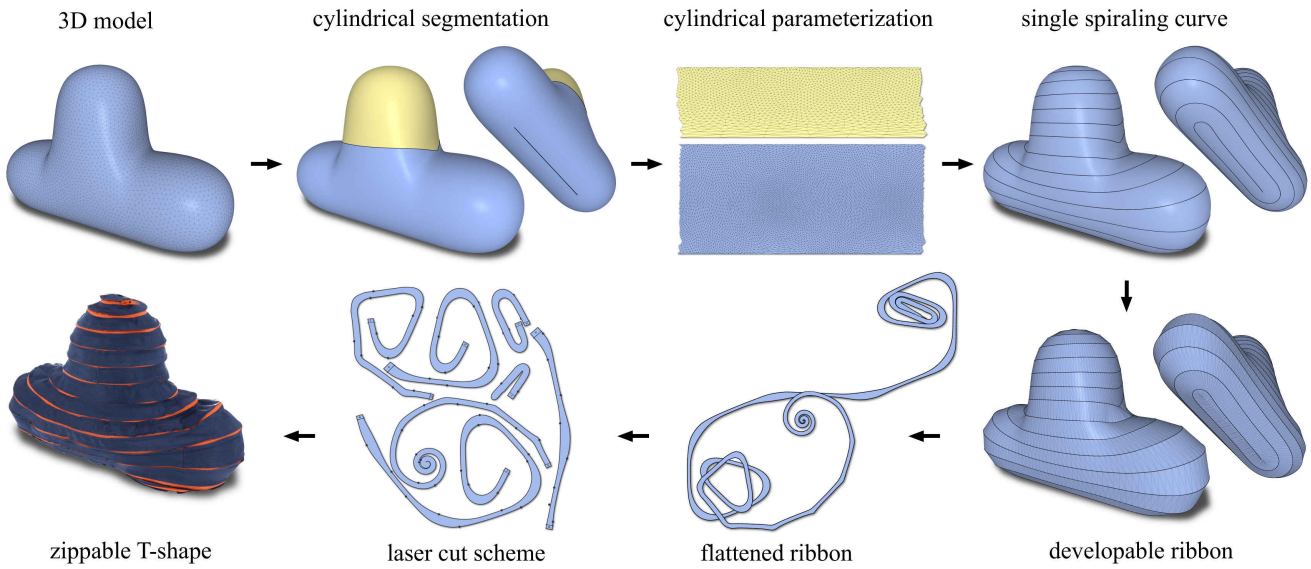

Figure 17: *Overview of our pipeline. We begin by segmenting a 3D model to cylinders, followed by a global cylinder parameterization. Using the parameterization, we trace a spiraling curve on the shape. The shape is then cut along the curve and approximated by a developable ribbon. The ribbon is then unfolded to the plane. We proceed by packing the design in order to create a cutting program for a laser cutter. Finally, we cut the design from a piece of fabric and sew a zipper along its boundary. When zipped-up, the ribbon reproduces the original shape.*

Figure 18: *A display of all our fabricated results (not including the simple T shape).*

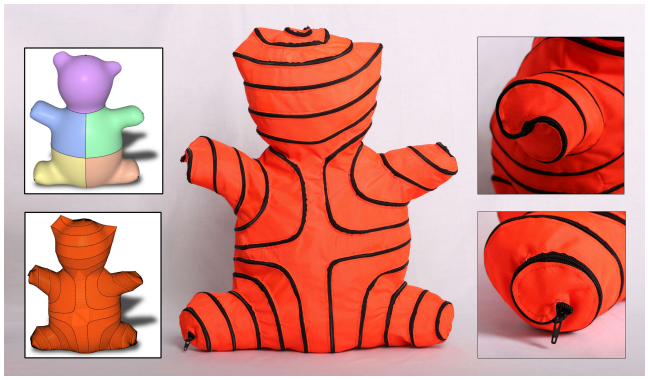

Figure 19: *A zippable teddy bear with Fermat spirals for the head and the two arms.*

attempt to align the zipper curve to the input shape's features. This means that regions with sharp corners may not be well represented by the assembled ribbon, unless the user manually specifies it. We plan to tackle this issue in the future by using feature detection and incorporating it into the optimization. Additionally, we are interested in targeting more global objectives, such as symmetry. Furthermore,

Figure 20: *The Standford bunny from the teaser, zipped-up and shown from different perspectives.*

we would like to combine the curve design with the final stage of remeshing to get a developable surface. Currently these steps are strictly decoupled, and we expect to get a better approximation by optimizing both parts simultaneously.

In this paper we mostly discussed the design part of the process, and put less emphasis on the fabrication part. In the future, we plan to explore the possibility of automating the process, extend it to other types of fabrication methods, and search for different applications. One application that we are particularly interested in is *pipe cladding*, which is part of the process of insulating heated pipes with metal sheets, and is a labour intensive task. The process is similar to our zipping, but usually performed by and expert using a manual approach; our method could be potentially used to greatly simplify this task.

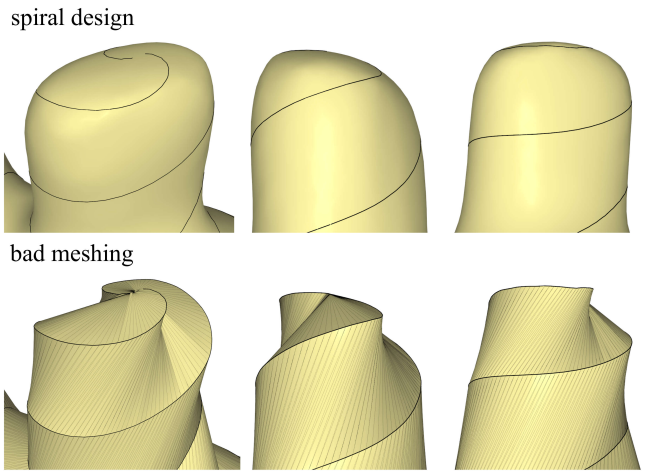

Figure 21: *Our naive meshing algorithm can sometimes fail to produce a good approximation of the original surface.*

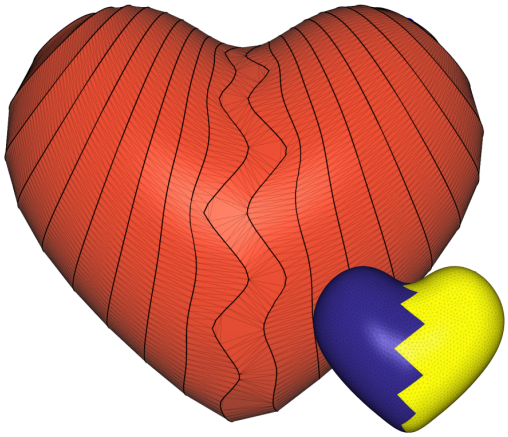

Figure 22: *The heart shape is segmented in a zig-zag way to enforce a curved ribbon in the middle of it.*

6 Acknowledgements

The authors would like to thank Isabelle von Salis for her support with the fabrication of the prototype. This work was supported in part by the ERC grant iModel (StG-2012-306877).

References

AIGERMAN, N., AND LIPMAN, Y. 2015. Orbifold tutte embeddings. *ACM Trans. Graph. 34*, 6 (Oct.), 190:1–190:12.

BICKEL, B., BÄCHER, M., OTADUY, M. A., LEE, H. R., PFISTER, H., GROSS, M., AND MATUSIK, W. 2010. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph. 29*, 4 (July), 63:1–63:10.

BOMMES, D., LÉVY, B., PIETRONI, N., PUPPO, E., SILVA, C., TARINI, M., AND ZORIN, D. 2013. Quad-mesh generation and processing: A survey. *Comput. Graph. Forum 32*, 6, 51–76.

CHANDRA, S., KÖRNER, A., KORONAKI, A., SPITERI, R., AMIN, R., KOWLI, S., AND WEINSTOCK, M. 2015. Computing curved-folded tessellations through straight-folding approximation. In *Proc. Symposium on Simulation for Architecture & Urban Design*.

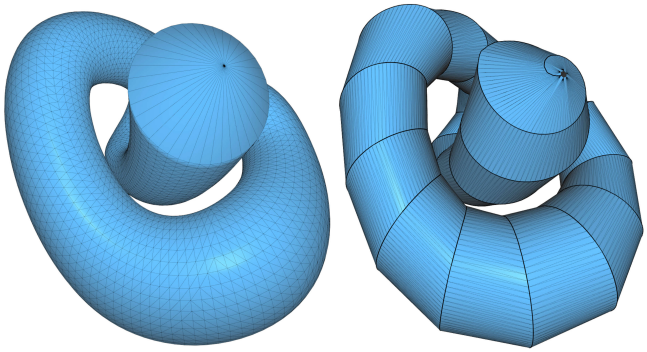

Figure 23: *The developable ribbon nicely wraps around the extruded cylinder even for the long and twisty shapes. The open boundaries are at the two end points of the cylinder knot.*

CHEN, X., ZHENG, C., XU, W., AND ZHOU, K. 2014. An asymptotic numerical method for inverse elastic shape design. *ACM Trans. Graph. 33*, 4.

EPPSTEIN, D., AND GOPI, M. 2004. Single-strip triangulation of manifolds with arbitrary topology. In *Proc. Symp. Computational Geometry*, 455–456.

HORMANN, K., LÉVY, B., AND SHEFFER, A. 2007. Mesh parameterization: theory and practice. In *ACM SIGGRAPH Courses*.

IGARASHI, Y., AND IGARASHI, T. 2008. Pillow: Interactive flattening of a 3D model for plush toy design. In *Proc. Smart Graphics*.

IGARASHI, Y., IGARASHI, T., AND SUZUKI, H. 2009. Interactive cover design considering physical constraints. *Comput. Graph. Forum 28*, 7, 1965–1973.

JULIUS, D., KRAEVOY, V., AND SHEFFER, A. 2005. D-Charts: Quasi-developable mesh segmentation. *Comput. Graph. Forum 24*, 3, 581–590.

KÄLBERER, F., NIESER, M., AND POLTHIER, K. 2011. Stripe parameterization of tubular surfaces. In *Topological Methods in Data Analysis and Visualization*. Springer, 13–26.

KILIAN, M., FLÖRY, S., CHEN, Z., MITRA, N. J., SHEFFER, A., AND POTTMANN, H. 2008. Curved folding. *ACM Trans. Graph. 27*, 3.

KNÖPPEL, F., CRANE, K., PINKALL, U., AND SCHRÖDER, P. 2015. Stripe patterns on surfaces. *ACM Trans. Graph. 34*, 4.

KOLMIANIČ, S., AND GUID, N. 2002. From geometric modeling to shape modeling. In *Proc. Workshop on Geometric Modeling: Fundamentals and Applications*, 35–46.

KOVALSKY, S. Z., GALUN, M., AND LIPMAN, Y. 2016. Accelerated quadratic proxy for geometric optimization. *ACM Trans. Graph. 35*, 4, 134.

LIU, Y., LAI, Y., AND HU, S. M. 2009. Stripification of free-form surfaces with global error bounds for developable approximation. *IEEE Trans. Automation Science and Engineering 6*, 4, 700–709.

LIVESU, M., ATTENE, M., PATANÈ, G., AND SPAGNUOLO, M. 2017. Explicit cylindrical maps for general tubular shapes. *Computer-Aided Design*. Accepted for publication.

MAHDAVI-AMIRI, A., WHITTINGHAM, P., AND SAMAVATI, F. 2015. Cover-it: an interactive system for covering 3D prints. In *Proc. Graphics Interface*, 73–80.

MASSARWI, F., GOTSMAN, C., AND ELBER, G. 2007. Papercraft models using generalized cylinders. In *Proc. Pacific Graphics*.

MITANI, J., AND SUZUKI, H. 2004. Making papercraft toys from meshes using strip-based approximate unfolding. *ACM Trans. Graph. 23*, 3, 259–263.

MORI, Y., AND IGARASHI, T. 2007. Plushie: an interactive design system for plush toys. *ACM Trans. Graph. 26*, 3.

MYLES, A., AND ZORIN, D. 2012. Global parametrization by incremental flattening. *ACM Trans. Graph. 31*, 4, 109:1–109:11.

RABINOVICH, M., PORANNE, R., PANOZZO, D., AND SORKINE-HORNUNG, O. 2017. Scalable locally injective mappings. *ACM Trans. Graph. 36*, 2, 16:1–16:16.

RAY, N., LI, W., LÉVY, B., SHEFFER, A., AND ALLIEZ, P. 2006. Periodic global parameterization. *ACM Trans. Graph. 25*, 4.

ROSE, K., SHEFFER, A., WITHER, J., CANI, M.-P., AND THIBERT, B. 2007. Developable surfaces from arbitrary sketched boundaries. In *Proc. Symp. Geom. Processing*, 163–172.

ROSSIGNAC, J. 1999. Edgebreaker: Connectivity compression for triangle meshes. *IEEE Trans. Vis. Comput. Graph. 5*, 1, 47–61.

SHATZ, I., TAL, A., AND LEIFMAN, G. 2006. Paper craft models from meshes. *The Visual Computer 22*, 9-11, 825–834.

SHEWCHUK, J. R. 1996. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. In *Applied Computational Geometry: Towards Geometric Engineering*, vol. 1148 of *Lecture Notes in Computer Science*. 203–222.

SHTENGEL, A., PORANNE, R., SORKINE-HORNUNG, O., KOVALSKY, S. Z., AND LIPMAN, Y. 2017. Geometric optimization via composite majorization. *ACM Trans. Graph. 36*, 4 (July), 38:1–38:11.

SKOURAS, M., THOMASZEWSKI, B., BICKEL, B., AND GROSS, M. H. 2012. Computational design of rubber balloons. *Comput. Graph. Forum 31*, 2, 835–844.

SKOURAS, M., THOMASZEWSKI, B., COROS, S., BICKEL, B., AND GROSS, M. 2013. Computational design of actuated deformable characters. *ACM Trans. Graph. 32*, 4, 82:1–82:10.

SKOURAS, M., THOMASZEWSKI, B., KAUFMANN, P., GARG, A., BICKEL, B., GRINSPUN, E., AND GROSS, M. H. 2014. Designing inflatable structures. *ACM Trans. Graph. 33*, 4.

SMITH, J., AND SCHAEFER, S. 2015. Bijective parameterization with free boundaries. *ACM Trans. Graph. 34*, 4, 70:1–70:9.

STRAUB, R., AND PRAUTZSCH, H. 2011. Creating optimized cut-out sheets for paper models from meshes. Tech. rep., Karlsruhe Institute of Technology.

TAKAHASHI, S., WU, H., SAW, S. H., LIN, C., AND YEN, H. 2011. Optimized topological surgery for unfolding 3D meshes. *Comput. Graph. Forum 30*, 7, 2077–2086.

TANG, C., BO, P., WALLNER, J., AND POTTMANN, H. 2016. Interactive design of developable surfaces. *ACM Trans. Graph. 35*, 2, 12:1–12:12.

TARINI, M. 2012. Cylindrical and toroidal parameterizations without vertex seams. *J. Graphics Tools 16*, 3, 144–150.

WANG, C. C. L. 2010. From designing products to fabricating them from planar materials. *IEEE Computer Graphics and Applications 30*, 6, 74–85.

ZENG, L., LIU, Y., CHEN, M., AND YUEN, M. M. 2012. Least squares quasi-developable mesh approximation. *Computer Aided Geometric Design 29*, 7, 565–578.

ZHAO, H., GU, F., HUANG, Q., GALICIA, J. A. G., CHEN, Y., TU, C., BENES, B., ZHANG, H., COHEN-OR, D., AND CHEN, B. 2016. Connected Fermat spirals for layered fabrication. *ACM Trans. Graph. 35*, 4, 100:1–100:10.

ZHOU, Y., YIN, K., HUANG, H., ZHANG, H., GONG, M., AND COHEN-OR, D. 2015. Generalized cylinder decomposition. *ACM Trans. Graph. 34*, 6, Article 171.

ZIPIT, 2017. zipit store online. http://www.zipitstore.com/. Accessed: 2017-02-01.