

Mechanics-Aware Deformation of Yarn Pattern Geometry

Supplementary Document

GEORG SPERL, IST Austria, Austria

RAHUL NARAIN, Indian Institute of Technology, India

CHRIS WOJTAN, IST Austria, Austria

CONTENTS

S1	Removal of Short Yarn Fragments	1
S2	Rendering Details	1
S2.1	Transformation of Edge Normals	1
S2.2	Geometry Tessellation	1
S2.3	Ply Texture Mapping	2
S3	Other Implementation Details	3
	References	3

S1 REMOVAL OF SHORT YARN FRAGMENTS

After tiling the periodic yarn pattern over the mesh as in Section 4, yarns near mesh boundaries may be very short. This may generate tiny floating yarn fragments as illustrated in Figure S1, which we want to remove.

The yarn patterns used in the optimization of Section 3 also store rest lengths \bar{l}_i of edges i . During tiling, we copy these rest lengths along with the other pattern data. After tiling, we can now traverse connected yarns and accumulate these rest lengths to compute a parameter

$$a_0 = 0, \quad a_i = \sum_{j=1}^i \bar{l}_{j-1}. \quad (\text{S1})$$

Here, we use the final value of a on a yarn to remove yarns that are smaller than some user-specified minimum length. We also use a during rendering of ply twists, since it provides a parametrization that is independent of local stretching (and so stretching a yarn segment would not change the amount of twisting). Note that in general the rest lengths are different from the edge lengths in the pattern geometry corresponding to $\mathbf{I} = \text{Id}$, $\mathbf{\Pi} = \mathbf{0}$, where the yarns may locally still exhibit strains, e.g. tension balancing out collisions.

S2 RENDERING DETAILS

In this section, we provide detail on world-space mapping of edge normals, the tessellation of yarn centerlines into cylindrical geometry including yarn twists, our volume-preserving radius-scaling heuristic, and on the twistable ply- and fiber-level detail texture maps.

S2.1 Transformation of Edge Normals

Yarn twists are defined with a twist variable θ and reference directors $\underline{\mathbf{d}}_1$ (edge normals) per edge. In Section 4, we displace quantities from a material-space configuration \mathbf{Q} to a deformed material-space configuration $\hat{\mathbf{Q}}$ and then map that to world-space values \mathbf{q} along with the mesh. The twist variables transform trivially (simply copying values), since the edge normals take care of the actual orientation




Fig. S1. Tiled yarn geometry before (left) and after (right) removing short yarn fragments.

of the twisting frame. In principle, this means that we would have to first transform $\underline{\mathbf{d}}_1$ along with the vertex-displacements ($\hat{\mathbf{X}} = \mathbf{X} + \Delta\mathbf{X}$) and then along with the world-space mapping \mathbf{x} from (9), which may be non-linear (due to Phong deformation [James 2020]). Similar to how normals transform with normal matrices, $\underline{\mathbf{d}}_1$ get transformed by the inverse transpose of the Jacobian $\mathbf{J}^{-\top}$ of the total mapping $\mathbf{X} \rightarrow \mathbf{x}$. We efficiently approximate the transformation with

$$\mathbf{J} = [\mathbf{F}, \quad \mathbf{n}], \quad (\text{S2})$$

where \mathbf{F} is the underlying triangle's 3×2 deformation gradient and \mathbf{n} is the interpolated mesh-vertex normal. We apply this transformation to the initial values of $\underline{\mathbf{d}}_1$ from the pattern tiling. This approximation is cheap compared to the full Jacobian including Phong Deformation and local displacements. During rendering, we simply reorthonormalize the transformed edge normals with respect to yarn centerline tangents.

S2.2 Geometry Tessellation

We use a geometry shader to tessellate cylindrical yarn meshes. The inputs are a global list of yarn vertices and index lists for each yarn. Each yarn vertex has the following data: a position \mathbf{x} , the twist θ of its outgoing incident edge, a transformed normal $\mathbf{n} = \mathbf{J}^{-\top} \underline{\mathbf{d}}_1$ of the same edge using (S2), the cumulative rest length parameter a up until that vertex, as well as a copy of the original undeformed material-space coordinates for cloth-level texturing. Since the material space of the mesh is encoded as uv-coordinates and the yarn pattern is tiled over that space, this is a natural way of texturing or coloring a garment.

The vertices are passed to the geometry shader stage as primitives of four consecutive yarn vertices, from which we consistently tessellate the middle edge such that the adjacent primitives seamlessly connect, as illustrated in Figure S2. To do so, we push the edge-based quantities onto the vertices with rest-length-weighted averaging. The rest lengths can be computed from a or stored explicitly per edge. We average vertex tangents \mathbf{t}_{v_i} , vertex normals \mathbf{n}_{v_i} , and vertex twists θ_{v_i} for the inner two vertices $v_i = v_1$ and $v_i = v_2$ of the four-vertex primitive. We orthogonalize the \mathbf{n}_{v_i} with respect to \mathbf{t}_{v_i} ,




Fig. S2. Top: Four-vertex input primitive for the geometry shader. A yarn segment is tessellated for the middle edge, using averaged per-vertex tangents t_{v_i} and normals n_{v_i} . Bottom: Cylinders are generated from tessellated cross-sections per centerline vertex v_i , with tangential frames $M_j = [t_{v_i} \times n_j, t_{v_i}, n_j]$.




Fig. S3. Tessellated yarn geometry without (left) and with (right) volume-preserving radius scaling. In our experiments, we found the effect to be small such as shown here.

normalize both vectors and compute vertex binormals b_{v_i} as the cross-product

$$\mathbf{b}_{v_i} = \mathbf{t}_{v_i} \times \mathbf{n}_{v_i}. \quad (\text{S3})$$

Volume-Preserving Heuristic. At this point, we also compute local yarn radii such that the deforming yarns preserve volume compared to their rest shape. This is a fast approximation that appears to work well as shown in Figure S3. It would be possible to instead integrate the method of Montazeri et al. [2019] for a data-driven solution to adapting yarn cross-sections to stretching, or to use a different thickness heuristic.

For each of the three edges e_i in the geometry primitive, we compute volume-preserving radii r_{e_i} from their current length l_{e_i} , their rest lengths \bar{l}_{e_i} , and the base radius R :

$$r_{e_i} = R \sqrt{\bar{l}_{e_i} / l_{e_i}} \quad (\text{S4})$$

We additionally clamp the ratio of the modified vs. original radii such that their ratio to the original yarn radius is between 0.25 and 2.0, although this almost never takes effect for “reasonable” deformation within the realm of the strains sampled for our precomputed physical




Fig. S4. Adding random displacements to tessellated yarn cylinder vertices is an efficient way of generating a fuzzy look.

yarn data. We compute vertex radii r_{v_i} using the weighted average discussed above.

Tessellation. Finally, with the rotational frames $(\mathbf{n}_{v_i} \mathbf{b}_{v_i})$ and the twists θ_{v_i} at the two vertices of the middle segment, we can generate the cylindrical mesh by generating a discretized circle at each vertex and connecting them with triangular faces.

We compute the position \mathbf{x}_j of the j -th circle vertex at primitive vertex v_i as

$$s = \sin\left(\frac{2\pi j}{N-1} + \theta_{v_i}\right), \quad (\text{S5a})$$

$$c = \cos\left(\frac{2\pi j}{N-1} + \theta_{v_i}\right), \quad (\text{S5b})$$

$$\mathbf{n}_j = c \mathbf{n}_{v_i} + s \mathbf{b}_{v_i}, \quad (\text{S5c})$$

$$\mathbf{x}_j = \mathbf{x}_{v_i} + r_{v_i} \mathbf{n}_j, \quad (\text{S5d})$$

where N is the number of vertices per circle. \mathbf{n}_j is the normal to the cylinder at cylinder-vertex j and different from the edge normal \mathbf{n}_{v_i} of the centerline vertex v_i . We implement the elastic rods yarn twist here as an offset to the arguments of s and c . Note that this basically corresponds to the combined computation and tessellation of “material frames” from [Bergou et al. 2010]. It is possible to perform a naive level-of-detail optimization at this stage by reducing N depending on the depth value of vertex v_i . For normal mapping, we also emit a tangent-frame matrix M_j per vertex

$$M_j = [\mathbf{t}_{v_i} \times \mathbf{n}_j, \mathbf{t}_{v_i}, \mathbf{n}_j], \quad (\text{S6})$$

We render the tessellated geometry with “matcaps” (spherical environment maps) and screen-space ambient occlusion. In Figure S4, we briefly experimented adding fuzz detail by displacing vertices similar to [Zhong et al. 2001].

S2.3 Ply Texture Mapping

We add ply- and fiber-level detail onto the tessellated yarn geometry using *twistable* normal maps and ambient occlusion maps. The twists discussed here are separate from the discrete elastic rods twist θ , which we already incorporate into the cylinder tessellation. Instead, the twists in this section represent the helical geometry of yarn plies. We propose texture maps that can be tiled and twisted along the yarn in a parametric fashion. Notably, this twisting and tiling has an influence on the resulting normals.

The process with which we construct parametrized normals is as follows: we define ambient occlusion and normal textures for a single ply half-cylinder (Figure S5 top left), we stretch, tile and shear




Fig. S5. Top left: The normalized single-ply space from which we bake geometry (indicated by the blue half-cylinder) into normal and ambient occlusion textures. Top right: Transformation of single-ply into stretched/tilted/twisted ply space. Bottom left: Our baked combined normal and ambient occlusion texture with fiber geometry for a single ply. Bottom right: Yarn geometry rendered with our twistable texture maps.

these textures into an intermediate space (Figure S5 top right), and lastly with “standard” normal mapping we map the normal into the tangent frame of our textured yarn geometry (Figure S2).

The projection of geometry in single-ply space can be described using a height field h

$$\mathbf{p}_{\text{ply}}(\hat{\alpha}, \hat{l}) = \begin{pmatrix} \hat{\alpha} \\ \hat{l} \\ h(\hat{\alpha}, \hat{l}) \end{pmatrix} \quad (\text{S7})$$

with the “across-ply” parameter $\hat{\alpha}$ and the “along-ply” \hat{l} in $[0, 1]$. This single-ply geometry is what we bake into a combined normal and occlusion texture (Figure S5 bottom left). Specifically, we bake $\left(\frac{n_x+1}{2}, \frac{n_y+1}{2}, \text{AO}\right)$ into the RGB channels, where n_x and n_y are the components of the normals to h , and AO is the ambient occlusion value.

Next we stretch, tile, and shear the single-ply geometry into an intermediate space (corresponding to a flattening of the tessellated yarn cylinder)

$$\mathbf{p}_{\text{int}}(\hat{\alpha}, a) = \begin{pmatrix} 2\pi R \hat{\alpha} \\ a \\ 2H h(N\hat{\alpha} - a/L \hat{V}, a/L) \end{pmatrix} \quad (\text{S8})$$

with cumulative rest length a along the yarn, N the number of tiled plies, R the yarn radius (not the volume-preserving one), H the transformed height, L the transformed length, and \hat{V} a twist factor.

We want a nice expression for the normal \mathbf{n}_{int} of \mathbf{p}_{int} using the baked single-ply texture values n_x and n_y . We redefine $L = \hat{L}R/N$ and $H = \hat{H}R$. Then with texture coordinates

$$\mathbf{w} = \begin{pmatrix} N\hat{\alpha} - a \hat{V}/(R\hat{L}) \\ a/(R\hat{L}) \end{pmatrix}, \quad (\text{S9})$$

the normal of \mathbf{p}_{int} is proportional to

$$\mathbf{n}_{\text{int}}(\hat{\alpha}, a) \propto \begin{pmatrix} \frac{\hat{H}N n_x(\mathbf{w})}{\pi} \\ \frac{2\hat{H}N}{\hat{L}} (n_y(\mathbf{w}) - \hat{V} n_x(\mathbf{w})) \\ \sqrt{1 - n_x^2(\mathbf{w}) - n_y^2(\mathbf{w})} \end{pmatrix}. \quad (\text{S10})$$

To apply our transformed normal maps, during tessellation we compute $\hat{\alpha} = \frac{2\pi j}{N-1}$ from (S5), and store \mathbf{w} as well as \mathbf{M}_j from (S6) per cylinder vertex. These quantities are interpolated in the fragment shader, where we then compute \mathbf{n}_{int} from (S10) and map it into the tangent frame as $\mathbf{n}_{\text{final}} = \mathbf{M} \mathbf{n}_{\text{int}}$.

S3 OTHER IMPLEMENTATION DETAILS

Preloading Animation Data. We experiment with cloth mesh animations stored as obj-file sequences. For meshes of moderate size, loading meshes each frame can become a bottleneck. Therefore, we preload mesh animations into memory before adding and animating yarn detail. Additionally, we convert and serialize several mesh animations into a binary format for faster loading.

GPU Buffers. As stated in the main paper, we do all per-yarn-vertex computation on the gpu, and perform cloth-mesh computations on the cpu. With our results based on non-remeshing cloth simulations, we can precompute material-space mesh quantities (such as shape matrices and topology-dependent interpolation weights) once on the cpu without becoming a bottleneck. For more performance and remeshing cloth simulations, we suggest to implement all mesh processing on the gpu.

For yarn vertices, we mainly use two gpu buffers: one for material-space data and one for world-space data. The material-space data basically corresponds to the initial flat tiling and triangle assignment, which we do once on the cpu. We then buffer this data onto the gpu, and allocate space for the world-space buffer. Afterwards, the yarn data remains on the gpu, getting passed directly into shaders for rendering, and thus avoiding the need for expensive data-passing between the cpu and gpu.

Mesh Interpolation Weights. We use modified Shepard weights [James 2020] for strain interpolation from mesh faces to mesh vertices. We also use the same weights to compute vertex normals. To compute these weights, we use the mesh-topology in material space. Specifically, we have two sets of faces for material-space and world-space topology respectively. This way we can prohibit interpolation over uv-seams, i.e. faces that are neighbors in world space but not in material space do not count as neighbors for face-to-vertex interpolation of strains and normals.

REFERENCES

- Miklós Bergou, Basile Audoly, Etienne Vouga, Max Wardetzky, and Eitan Grinspun. 2010. Discrete viscous threads. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 1–10.
- Doug L James. 2020. Phong deformation: a better C0 interpolant for embedded deformation. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 56–1.
- Zahra Montazeri, Chang Xiao, Yun Fei, Changxi Zheng, and Shuang Zhao. 2019. Mechanics-Aware Modeling of Cloth Appearance. *IEEE transactions on visualization and computer graphics* 27, 1 (2019), 137–150.
- Hua Zhong, Ying-Qing Xu, Baining Guo, and Heung-Yeung Shum. 2001. Realistic and efficient rendering of free-form knitwear. *The Journal of Visualization and Computer Animation* 12, 1 (2001), 13–22.