Recent Advances in Realistic Cloth Rendering

Junqiu Zhu
University of California, Santa Barbara
USA
junqiuzhu@ucsb.edu

Zahra Montazeri
The University of Manchester
UK
zahra.montazeri@manchester.ac.uk

Ling-Qi Yan
University of California, Santa Barbara
USA
lingqi@cs.ucsb.edu

Matt Jen-Yuan Chiang
Meta Reality Labs Research
USA
mattchiang@meta.com

Figure 1: Advanced cloth rendering models. Left: results of our ply-based cloth model [6]. Middle: a frame from the movie "Avatar: The Way of Water," which utilized our yarn-based cloth rendering techniques [18]. Right: rendering of various curtain fabrics using our surface-based methods [17]. These advanced cloth rendering models achieve both realism and high performance.

Abstract

Fabrics play a significant role in our daily lives, making their digital modeling crucial for applications in online retail, the textile industry, and entertainment, such as games and movies. Despite its importance, there have been significant advances in fabric rendering techniques over the past decade to highlight the state-of-the-art advancements. Rendering fabrics is a complex task due to their intricate microgeometry and anisotropic appearance. Conventional approaches either model fabrics as 2D sheets, which are efficient but lack high fidelity in close-ups, or model them at the fiber level, trading high fidelity with significant computational costs. Recent models propose aggregation-based techniques that combine the strengths of both approaches.

Cloth rendering is challenging due to the need for high performance, detailed representation, handling repetitive patterns, and efficient multi-scale rendering. These complexities make it a hard problem in both academia and industry. The industry, in particular, has a strong interest in advancing cloth rendering to meet the high demand for better image quality. This course is essential for bridging the gap between recent research and practical applications.

Participants will gain insights into basic cloth properties, the difficulties inherent in cloth rendering, and a detailed review of micro-appearance and surface-based models. Additionally, we will cover innovative appearance aggregation techniques that enhance rendering efficiency without compromising visual quality. Through comparative analysis under different patterns and lighting conditions and evaluation against photographs, attendees will develop a deep understanding of state-of-the-art cloth rendering methods. The course will also explore open problems and future research directions, providing a comprehensive overview of the current state and potential advancements in the field.

In this course, we plan to:

- explain the basic knowledge of cloth, including classifications and unique effects.
- discuss the difficulties in cloth rendering due to geometric and optical complexities.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

ACM ISBN 978-1-4503-XXXX-X/18/06
https://doi.org/XXXXXXX.XXXXXXX

- review micro-appearance cloth models, their approaches, and their advantages and disadvantages.
- review surface-based cloth models, their approaches, and their advantages and disadvantages.
- introduce the concept of cloth appearance aggregation to enhance efficiency without compromising quality.
- describe the new advanced curve-based model using the aggregation-based idea. Introduce advanced surface-based models to ultimately render fabrics efficiently in high fidelity.
- explore open problems and future directions in cloth rendering and facilitate further discussion.

CCS Concepts

• **Computing methodologies → Reflectance modeling**.

Keywords

Cloth modeling, cloth rendering, cloth BSDF

ACM Reference Format:

Junqiu Zhu, Zahra Montazeri, Ling-Qi Yan, and Matt Jen-Yuan Chiang. 2018. Recent Advances in Realistic Cloth Rendering. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation emai (Conference acronym 'XX).* ACM, New York, NY, USA, 4 pages. https://doi.org/XXXXXXX.XXXXXXX

1 COURSE FORMAT & PREREQUISITES

We propose a short course (1 hour and 45 minutes) at an intermediate difficulty level. It is expected that attendees have a basic understanding of appearance rendering and calculus. Familiarity with the rendering equation, microfacet theory, Monte Carlo integration, importance sampling, and the fundamentals of the Bidirectional Reflectance Distribution Function (BSDF), particularly for hair and cloth, is assumed. However, this course will provide a brief review of these concepts as gradually introduce the latest techniques in cloth rendering.

2 TARGET AUDIENCE

This course is designed for students, researchers, and rendering engineers interested in realistic cloth rendering, whether they are unfamiliar with existing techniques or recent advancements. It aims to provide insights into the challenges of cloth rendering, familiarize attendees with recent developments for achieving optimal performance and accuracy, and highlight open problems that could inspire future research. This will begin with a brief review of basic cloth rendering models, comparing lightweight, less accurate surface-based models with more detailed, accurate micro-appearance models. The focus will then shift to advanced models proposed over the past decade that combine the strengths of both approaches. Attendees should have at least a basic understanding of BSDF and importance sampling, such as that gained from a graduate course in computer graphics or image synthesis.

3 WHY A SIGGRAPH ASIA COURSE IN 2024?

Since the introduction of the microcylinder model a decade ago [8], it has become widely used in the industry due to its speed, despite lacking detailed realism. Meanwhile, the research community has been focusing on more detailed, micro-appearance models [14] that are slower but more realistic, creating a divide between industry and research. In the past five years, however, we have been dedicated to developing cloth models that are both fast and realistic. Our team has published 10 papers to develop this problem [4–7, 10–12, 16–18]. Our current work bridges the gap between research and industry, achieving photo-realistic results in real-time. Our success extends to practical applications in the industry. Zhu et al. [18] was used in the film "Avatar: The Way of Water," and [16] has become Meta's official model. This is a significant advancement, and we now feel confident presenting a usable, understandable, and theoretically sound update on the state-of-the-art in cloth rendering. Contemporaneously, we have received numerous inquiries about our work, especially following our presentations, sparking significant interest from both the research and industrial communities. Many graduate students have also shown keen interest in our work and consider the implementation of our model as their final project [10]. This variety and quantity of interest suggest that a coherent update on the state-of-the-art in cloth rendering would be welcomed by the rendering community. It would help accelerate further research and adoption and also aid in achieving fast, high-quality cloth rendering in games, VR, and movies.

4 PROPOSED SYLLABUS (SHORT)

5 min: Welcome and introductions. *Lingqi Yan*

A course overview, learning objectives, and speaker introductions.

10 min: Basic Knowledge of Cloth. *Lingqi Yan*

This will clarify why it is necessary to investigate cloth rendering. This starts with the basic knowledge of cloth, including its classifications (woven, knitted, felt), its complex hierarchical structure (single fiber → ply → yarn → weave pattern), and the unique effects of cloth (sheen, anisotropic specular reflection, transparency).

10 min: Difficulties in Cloth Rendering. *Matt Jen-Yuang Chiang*

Rendering cloth is challenging due to its geometric and optical complexities. Geometrically, modeling down to the fiber level can lead to issues with storage and geometric trace speed. Optically, fibers are more complex than surface appearances as they are cylindrical and require BCSDF to describe their behavior, including complex reflection and transmission. This is compounded by the complex hierarchical structure that introduces multiple scattering. This clarifies the difficulties of cloth rendering to prepare for discussing specific methods. We will review the traditional models of cloth rendering, categorizing them into two types based on rendering speed and effect: micro-appearance and surface-based.

10 min: Micro-appearance Cloth Models. *Matt Jen-Yuang Chiang*

First, we introduce micro-appearance models. A large number of works define cloth as an assembly of fibers, either through explicit geometry [15] or as a heterogeneous volumetric representation [9, 13]. The main advantage of curve-based cloth models is the accurate representation of fiber-level geometric details at a microscopic scale. However, the storage requirements for 3D details in curve-based methods are substantial, and the complexity of multiple bounces in modeling curves makes efficient rendering challenging. In addition to introducing the main ideas, we will also present representative works, explaining their approaches and analyzing their advantages and disadvantages.

10 min: Surface-based Cloth Models. *Junqiu Zhu*

Surface-based methods have been used for further practicality. As the name indicates, cloth is defined as 2D thin sheets, usually on the surface of a 3D mesh. In this way, the appearance of cloth can be easily represented by Bidirectional Reflectance Distribution Functions (BRDFs) [1–3, 8]. Surface-based models are popular for their fast performance; however, these models often lack realism due to missing yarn/ply/fiber details (often compensated in an ad-hoc texture mapping process), energy transmission (often misunderstood as transparency/alpha), and shadowing-masking (either missing or homogeneous regardless of local structures). This group of methods is more widely used in the industry. Similarly, in addition to introducing the main ideas, we will also present representative works, explaining their approaches and analyzing their advantages and disadvantages.

10 min: Cloth Appearance Aggregation. *Junqiu Zhu*

Surface-based cloth rendering techniques often appear non-realistic, while micro-appearance models, though realistic, are costly and slow. The challenge is achieving both realistic accuracy and real-time speed. The solution may lie in **appearance aggregation**: condensing details into a simplified representation to enhance efficiency without sacrificing visual quality. This will introduce the concept of appearance aggregation [19], using vivid examples to explain why and how we need to aggregate the appearance.

15 min: Advanced Curve-based Cloth Models. *Zahra Montazeri*

This will introduce our advanced curve-based cloth rendering method [4–6, 10, 11, 18], which addresses the practical issues of traditional micro-appearance methods by exclusively modeling the ply/yarn geometry. By deriving an accurate aggregated Bidirectional Curve Scattering Distribution Function (BCSDF) for the ply/yarn, this part will detail the derivation process of the aggregated BCSDF, considering physically-based parameters such as twist, density, and the number of fibers in a ply/yarn.

15 min: Advanced Surface-based Cloth Models. *Zahra Montazeri*

Surface-based methods remain significantly faster than curve-based methods. We will introduce our surface-based work [16**?** , 17], detailing how we aggregate the complex light scattering on surfaces and perform Level of Detail (LoD) calculations to achieve real-time performance. Additionally, we will present our advanced sheen model [12].

5 min: Open problems and future directions *All*

We now have a fairly comprehensive understanding of cloth rendering, but there are still some open questions that need to be addressed. Montazeri will moderate a panel discussion where authors will highlight directions that are interesting for future work. This will transition into an audience Q&A session.

Table 1: High-level syllabus for cloth rendering. See text for detailed descriptions.

Duration	Topic	Speaker
5 min	Welcome and introduction	Yan
10 min	Basic Knowledge of Cloth	Yan
10 min	Difficulties in Cloth Rendering	Chiang
10 min	Fiber-based Cloth Models	Chiang
10 min	Surface-based Cloth Models	Zhu
15 min	Cloth Appearance Aggregation	Zhu
15 min	Advanced Curve-based Cloth Models	Montazeri
15 min	Advanced Surface-based Cloth Models	Montazeri
5 min	Open problems and future directions	All
10 min	Audience Q & A	All

10 min: Q&A *All*

5 ABOUT THE CONTRIBUTORS

Junqiu Zhu is a postdoctoral fellow at the University of California, Santa Barbara, advised by Prof. Lingqi Yan. She received her Ph.D. degree from Shandong University, China, in 2022, under the supervision of Prof. Xiangxu Meng. Her industry projects include "Avatar: The Way of Water" (cloth rendering) and Meta's official cloth rendering model.

Zahra Montazeri is an assistant professor at the University of Manchester in the Department of Computer Science. Her field of research is physics-based rendering and appearance modeling for complex materials such as cloth, hair, and fur. She worked as a research consultant at Disney Research and before academia, she was an R&D at Industrial Light&Magic (ILM) and interned at Pixar Animation Studios, DreamWorks, and for two years at Luxion (makers of KeyShot). She holds a PhD and MSc from University of California, Irvine and a B.Sc. from the Sharif University of Technology in Iran.

Matt Jen-Yuan Chiang is a research engineer at Meta Reality Labs. Prior to this, he was a senior software engineer at Walt Disney Animation Studios, where he worked on Disney's Hyperion renderer. He also contributed to projects at the Graphics Lab of USC Institute for Creative Technologies, focusing on light stage and digital human-related projects.

Ling-Qi Yan Lingqi Yan is an Assistant Professor of Computer Science at UC Santa Barbara, a co-director of the MIRAGE Lab, and affiliated faculty in the Four Eyes Lab. Lingqi strives to achieve photo-realistic rendering. His research focuses on physically-based rendering in both micro and macro ends, encompassing areas such as appearance modeling and representation, physical light transport theory, and real-time ray tracing practices. Lingqi's outstanding contributions have been recognized through various accolades, including the SIGGRAPH 2019 Outstanding Doctoral Dissertation Award, a SIGGRAPH 2022 Best Paper Honorable Mention and an EGSR 2023 Best Paper Award.

References

[1] Neeharika Adabala, Nadia Magnenat-Thalmann, and Guangzheng Fei. 2003. Real-Time Rendering of Woven Clothes *(VRST '03)*. Association for Computing Machinery, New York, NY, USA, 41–47. https://doi.org/10.1145/1008653.1008663

[2] Piti Irawan and Steve Marschner. 2012. Specular reflection from woven cloth. *ACM Trans. Graph.* 31, 1 (2012), 1–20.

[3] Wenhua Jin, Beibei Wang, Milos Hasan, Yu Guo, Steve Marschner, and Ling-Qi Yan. 2022. Woven Fabric Capture from a Single Photo. In *SIGGRAPH Asia 2022 Conference Papers*. 1–8.

[4] Apoorv Khattar, Junqui Zhu, Emiliano Padovani, Jean-Marie Aurby, Marc Droske, Ling-Qi Yan, and Zahra Montazeri. 2024. A Multi-scale Yarn Appearance Model with Fiber Details. *arXiv preprint arXiv:2401.12724* (2024).

[5] Zahra Montazeri, Soren B Gammelmark, Henrik Wann Jensen, and Shuang Zhao. 2021. A Practical Ply-Based Appearance Modeling for Knitted Fabrics. In *Proceedings of Eurographics Symposium on Rendering 2021*.

[6] Zahra Montazeri, Søren B Gammelmark, Shuang Zhao, and Henrik Wann Jensen. 2020. A practical ply-based appearance model of woven fabrics. *ACM Trans. Graph.* 39, 6 (2020), 1–13.

[7] Zahra Montazeri, Chang Xiao, Yun Fei, Changxi Zheng, and Shuang Zhao. 2021. Mechanics-Aware Modeling of Cloth Appearance. *IEEE Transactions on Visualization and Computer Graphics* 27, 1 (jan 2021), 137–150. https://doi.org/10.1109/TVCG.2019.2937301

[8] Iman Sadeghi, Oleg Bisker, Joachim De Deken, and Henrik Wann Jensen. 2013. A practical microcylinder appearance model for cloth rendering. *ACM Trans. Graph.* 32, 2 (2013), 1–12.

[9] Kai Schroder, Reinhard Klein, and Arno Zinke. 2011. A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum* 30, 4 (2011), 1277–1286.

[10] Guan Yu Soh and Zahra Montazeri. 2023. Neural Yarn-Level Appearance Model for Cloth Rendering. arXiv:2311.04061 [cs.GR] https://arxiv.org/abs/2311.04061

[11] Bowen Xue, Shuang Zhao, Henrik Wann Jensen, and Zahra Montazeri. 2024. A Hierarchical Architecture for Neural Materials. arXiv:2307.10135 [cs.GR] https://arxiv.org/abs/2307.10135

[12] Tizian Zeltner, Brent Burley, and Matt Jen-Yuan Chiang. 2022. Practical Multiple-Scattering Sheen Using Linearly Transformed Cosines. In *ACM SIGGRAPH 2022 Talks*. 1–2.

[13] Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2011. Building Volumetric Appearance Models of Fabric Using Micro CT Imaging. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 30, 4 (2011), 44:1–44:10.

[14] Shuang Zhao, Wenzel Jakob, Steve Marschner, and Kavita Bala. 2012. Structure-aware synthesis for predictive woven fabric appearance. *ACM Trans. Graph.* 31, 4 (2012), 1–10.

[15] Shuang Zhao, Lifan Wu, Frédo Durand, and Ravi Ramamoorthi. 2016. Downsampling Scattering Parameters for Rendering Anisotropic Media. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 35, 6 (2016), 166:1–166:11.

[16] Junqiu Zhu, Christophe Hery, Lukas Bode, Carlos Aliaga, Adrian Jarabo, Ling-Qi Yan, and Matt Jen-Yuan Chiang. [n. d.]. A Realistic Multi-scale Surface-based Cloth Appearance Model. ([n. d.]).

[17] Junqiu Zhu, Adrian Jarabo, Carlos Aliaga, Ling-Qi Yan, and Matt Jen-Yuan Chiang. 2023. A Realistic Surface-based Cloth Rendering Model. In *ACM SIGGRAPH 2023 Conference Proceedings*. 1–9.

[18] Junqui Zhu, Zahra Montazeri, J Aubry, Lingqi Yan, and Andrea Weidlich. 2023. A Practical and Hierarchical Yarn-based Shading Model for Cloth. In *Computer Graphics Forum*, Vol. 42. Wiley Online Library, e14894.

[19] Junqiu Zhu, Sizhe Zhao, Lu Wang, Yanning Xu, and Ling-Qi Yan. 2022. Practical level-of-detail aggregation of fur appearance. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–17.

Outline

Cloth 101

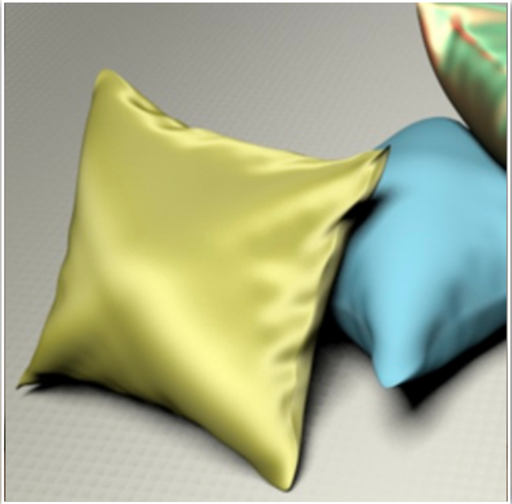

Cloth Rendering Models

Modern Cloth Rendering

Cloth 101

- Different level of cloth geometry

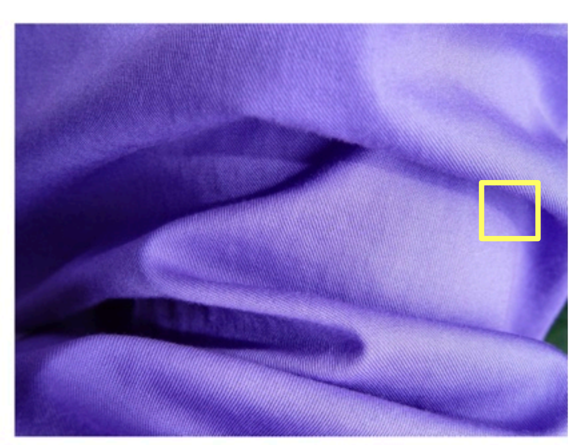

Woven cloth

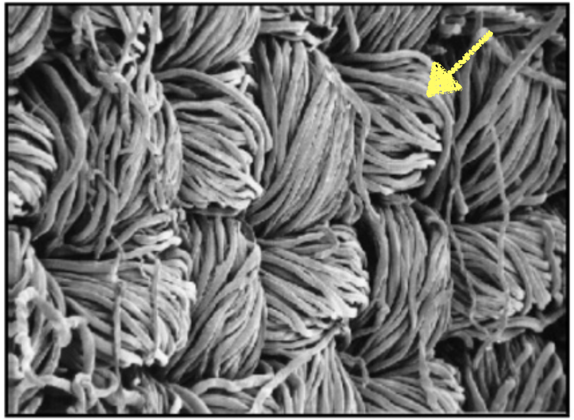

Underlying structure

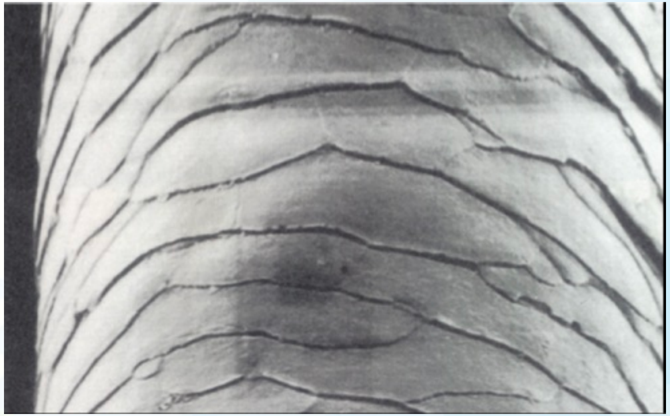

Sigle fiber

Cloth 101

- Common structure to form a yarn

- Common structure to form a piece of cloth

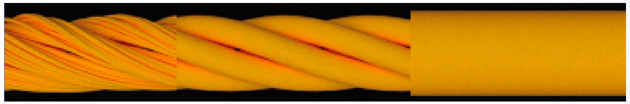

Fiber Ply Yarn

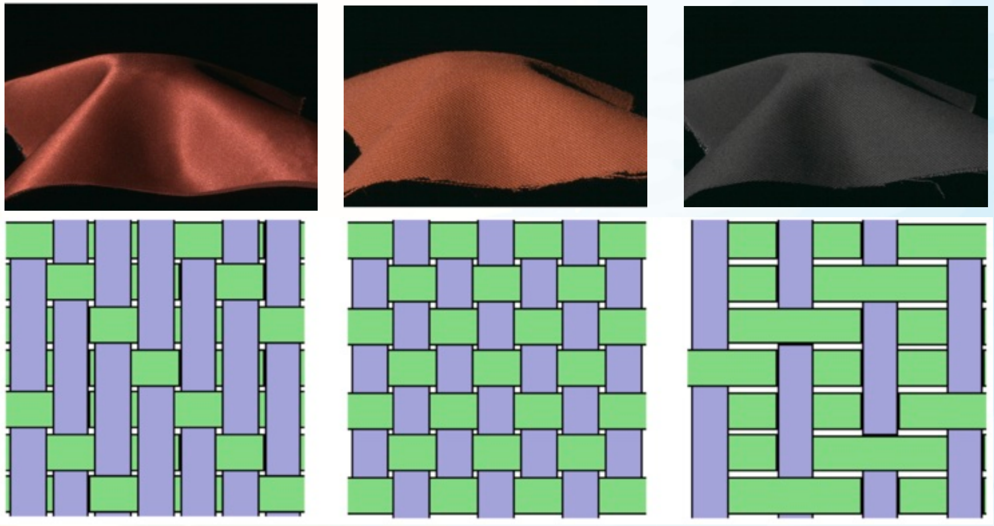

Satin Plain Weave Twill

Cloth 101

- Complex and Unique Visual Effects

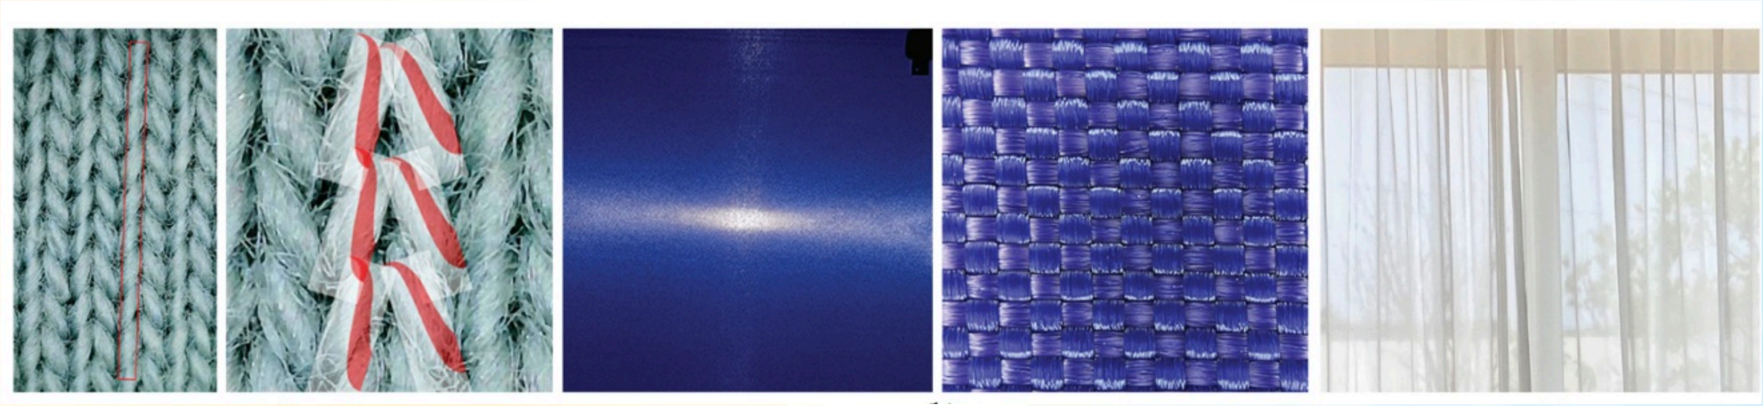

Asymmetric highlight due to fibers twisting   *Cross-shaped highlight due to weave pattern*   *Transparency of thin cloth*

Cloth 101 - Cloth Rendering is Difficult

- The number of fibers is too high

Memory intensive!

Large number of bounces!

Violation of energy conservation!

300 M

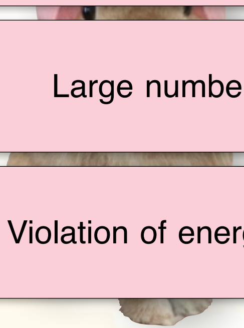

45 M

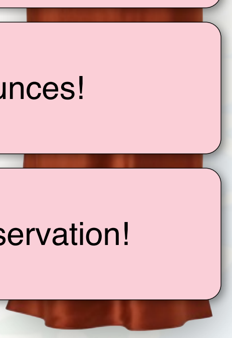

4.6 M

17 M

Cloth 101 - Cloth Rendering is Difficult

- Complex geometry structure

[Leaf et al. 2018]

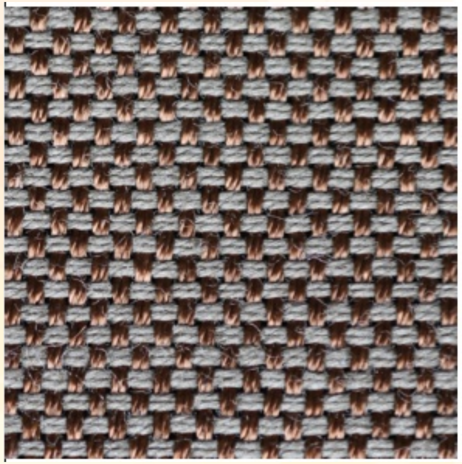

Woven sample

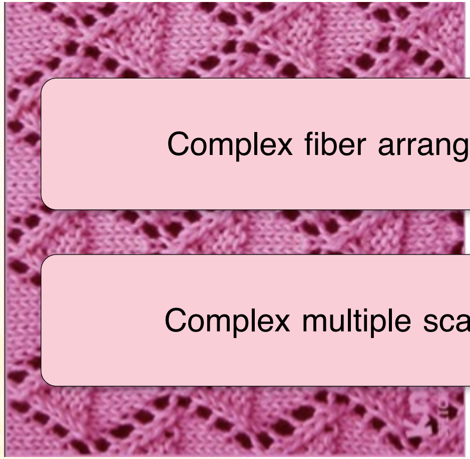

Knitted sample

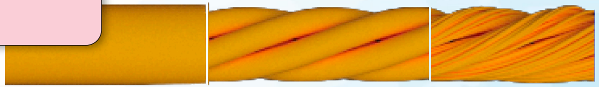

Complex fiber arrangement!

Complex multiple scattering!

Outline

Cloth 101

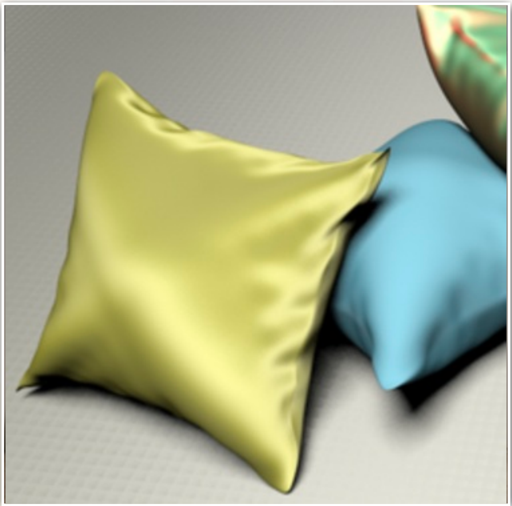

Cloth Rendering Models

Modern Cloth Rendering

Cloth Rendering Models

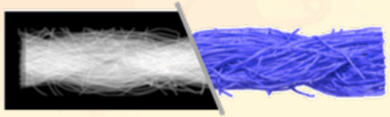

Fiber geometry *Rendering*

Micro-appearance methods

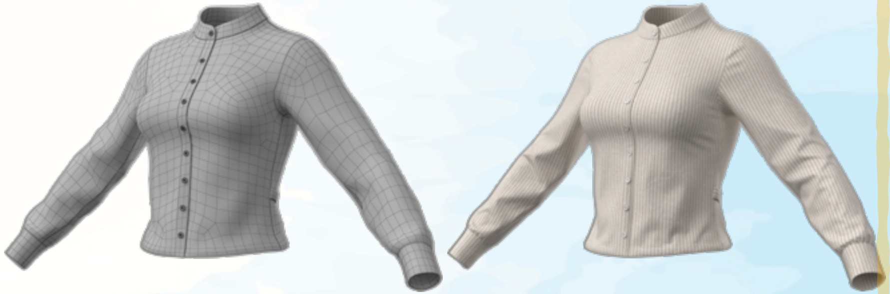

Surface geometry *Rendering*

Surface-based methods

Micro-appearance Methods - Geometry Modeling

Modeling the fiber geometry using control parameters

- Twist, density, migration, fly-away fibers

All fibers
in untied plies

Regular fibers

Flyaway fibers

[Zhao et al. 2016]

Why Aggregation?

Level of details (LOD)

Why Aggregation?

Peacock feather structure
(electron microscope)

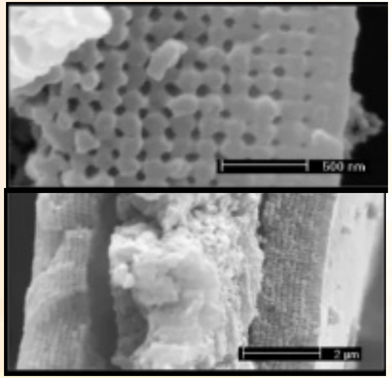

Peacock feather

Near

Far

Close-up:

primitive,

micro structure

Far away:

appearance

Why Aggregation?

- Use fewer fibers w/ **original fibers'** appearance ❌
    - Drier, harder and brighter!
- Use fewer fibers w/ **aggregated** appearance ✅
    - Realistic!