SIGGRAPH ASIA 2024 東京

Conference | 3–6 December 2024
Exhibition | 4–6 December 2024
Venue | Tokyo International Forum, Japan

Rendering Course: Recent Advances in Photorealistic Cloth Rendering

Lingqi Yan	University of California, Santa Barbara
Matt Chiang	Meta Reality Labs
Junqiu Zhu	University of California, Santa Barbara
Zahra Montazeri	University of Manchester

Sponsored by

Organized by koelnmesse

Learning Objectives

- Cloth rendering basics for beginners
- Types of cloth representations and their advantages/disadvantages
- Previous solutions
- State of the art solutions (academia and industry)

Target Audiences

- Students / Beginners
- Industrial friends
- Researchers
- Basically everyone!

Speakers

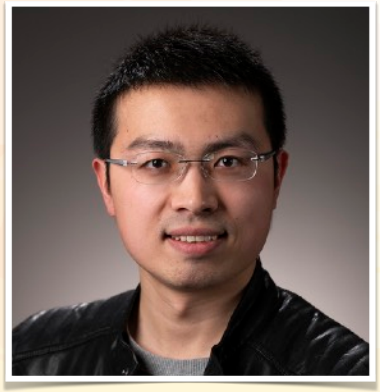

Lingqi Yan

Associate Professor

University of California Santa Barbara

Matt Chiang

Research Engineer

Meta Reality Labs

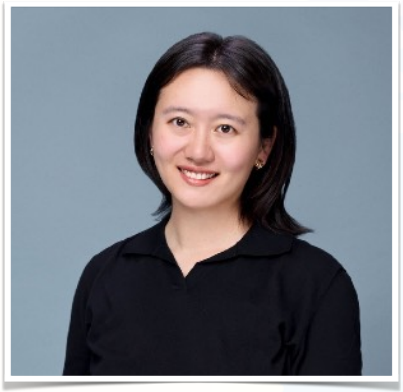

Junqiu Zhu

Postdoctoral Researcher

University of California Santa Barbara

Zahra Montazeri

Assistant Professor

University of Manchester

Arrangements

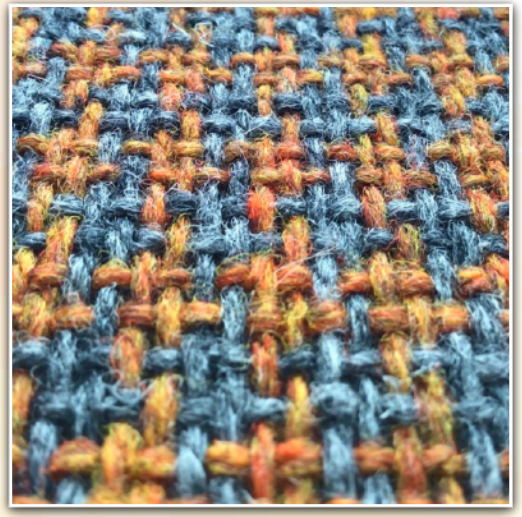

Cloth 101

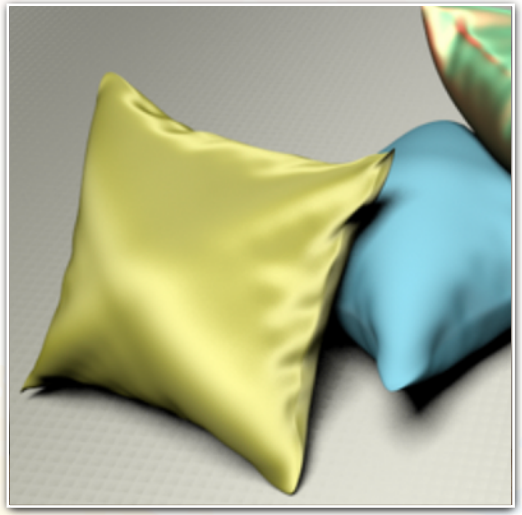

Cloth Rendering Models

Modern Cloth Rendering

Why still Cloth Rendering in 2024?
Isn’t that solved already?

Far From Photorealistic — As We Can See

By Alex Alvarado

By Pauline Boiteux

Far From Photorealistic — Cloth in Artists' Eyes

"We use 0.4 roughness,
0.3 **metallic**,
1.0 anisotropy,
a rather high **IOR**.
That's the cloth we got,
but I really don't know
what I am looking at."

— An anonymous friend
from the industry

Cloth modeling & rendering by Linctex Digital Technology

Key to Cloth Rendering — Cloth Structure

Determine

Structure

(Optical) Property

Cloth Structure

- Different levels of cloth geometry

Woven cloth

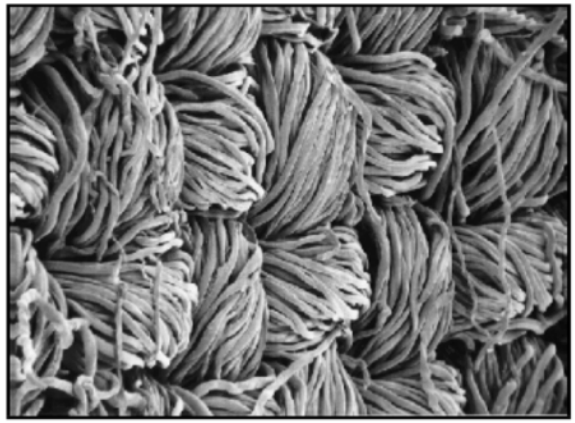

Underlying structure

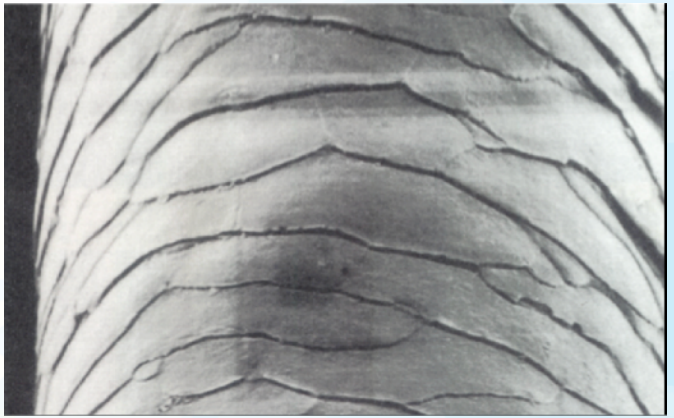

Sigle fiber

Cloth Structure

- Most common structure
- Three-level hierarchy
    - Fiber -> Ply -> Yarn

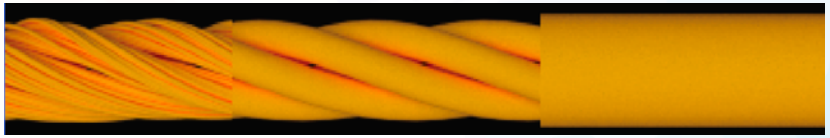

Fiber Ply Yarn

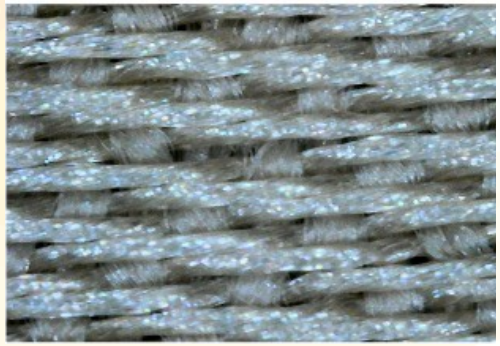

Woven cloth

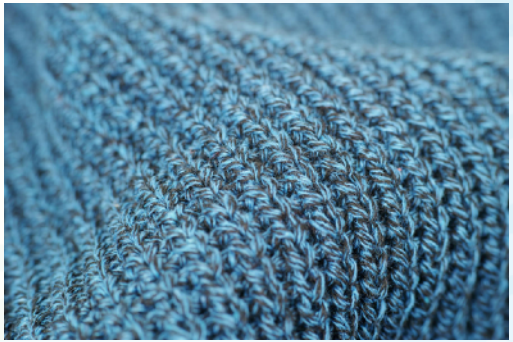

Knitted cloth

Cloth Structure

- Yarn to cloth?
- Woven or knitted

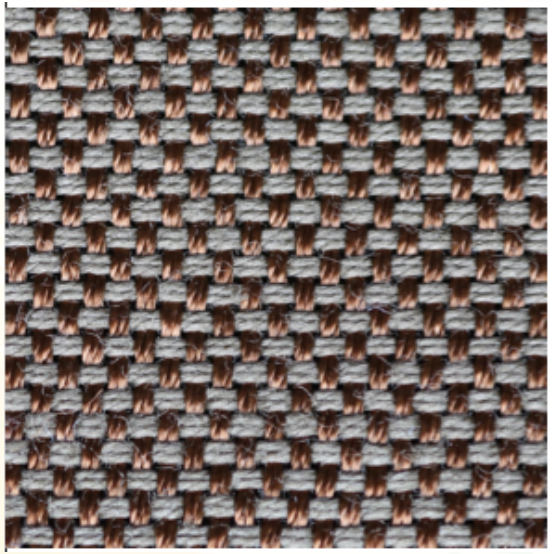

Woven

Knitted

Cloth Structure — Woven

- Use repetitive patterns to define a piece of cloth
- Different patterns -> Significantly different optical properties

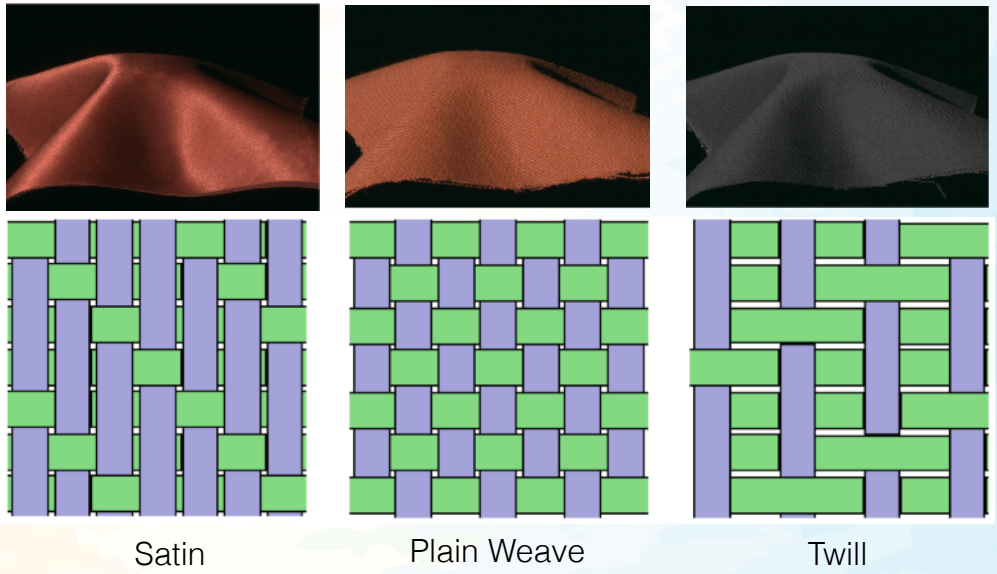

Cloth Structure — Knitted

- More complex patterns
  - Sweater
  - Shirt
- Yarns are usually thicker (though not necessarily)
  - So, the thickness of knitted cloth matters

Sweater Shirt

Desired Effects from Real Cloth

- Sheen

Desired Effects from Real Cloth

- Sheen
- Fly-away fibers

Desired Effects from Real Cloth

- Sheen
- Fly-away fibers
- Shadowing-masking

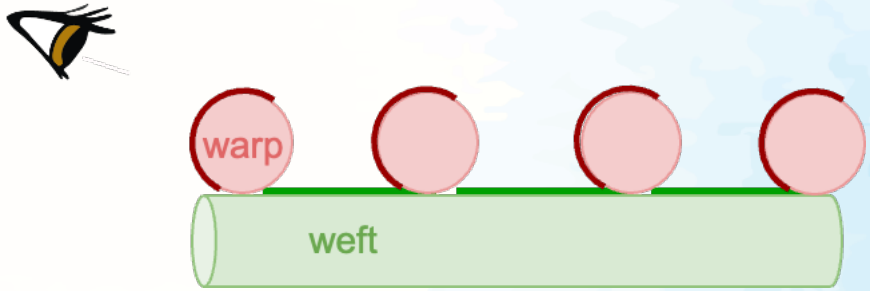

An example of masking

Desired Effects from Real Cloth

- Sheen
- Fly-away fibers
- Shadowing-masking
- Anisotropy

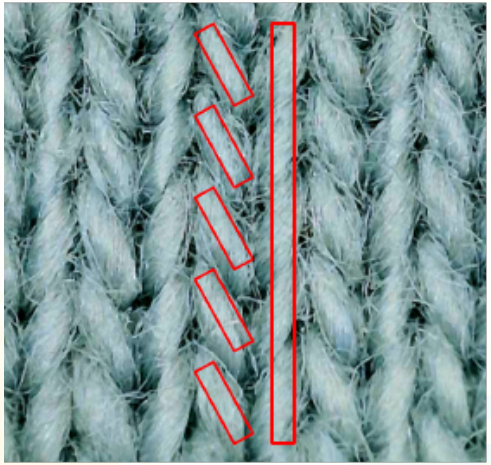

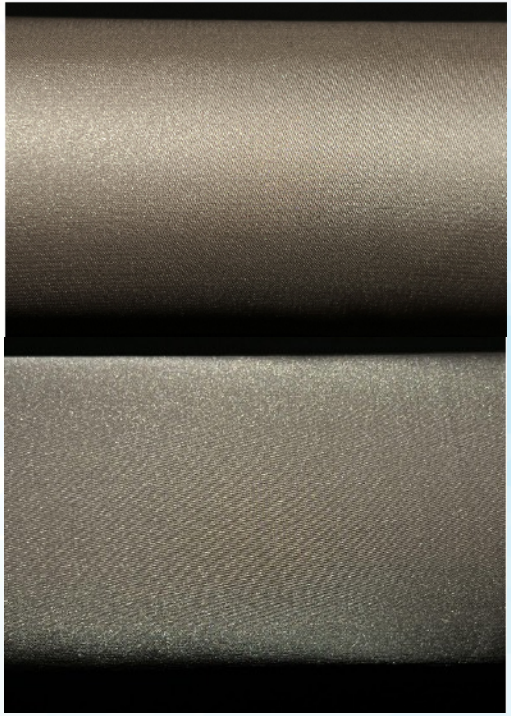

Desired Effects from Real Cloth

- Sheen
- Fly-away fibers
- Shadowing-masking
- Anisotropy
- Transparency

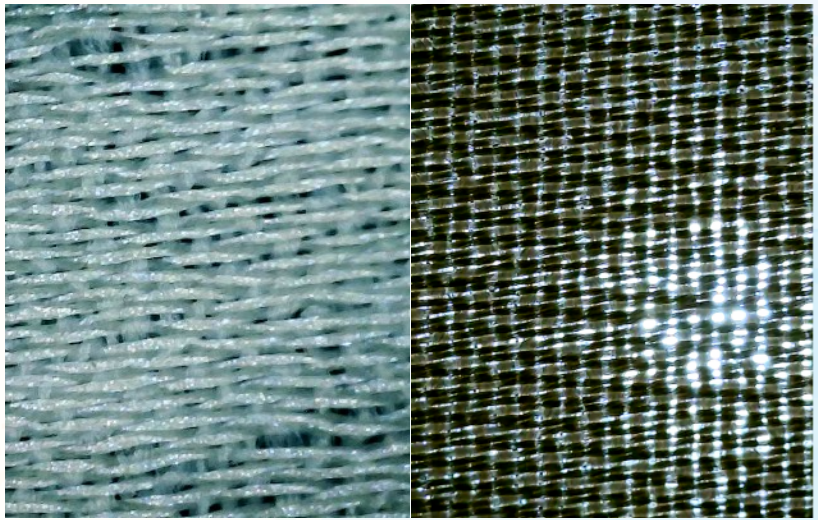

Desired Effects from Real Cloth

- Sheen
- Fly-away fibers
- Shadowing-masking
- Anisotropy
- Transparency
- Detail / Variation

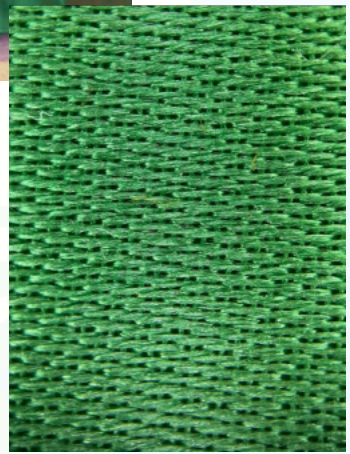

Key Challenge — The Tradeoff Between …

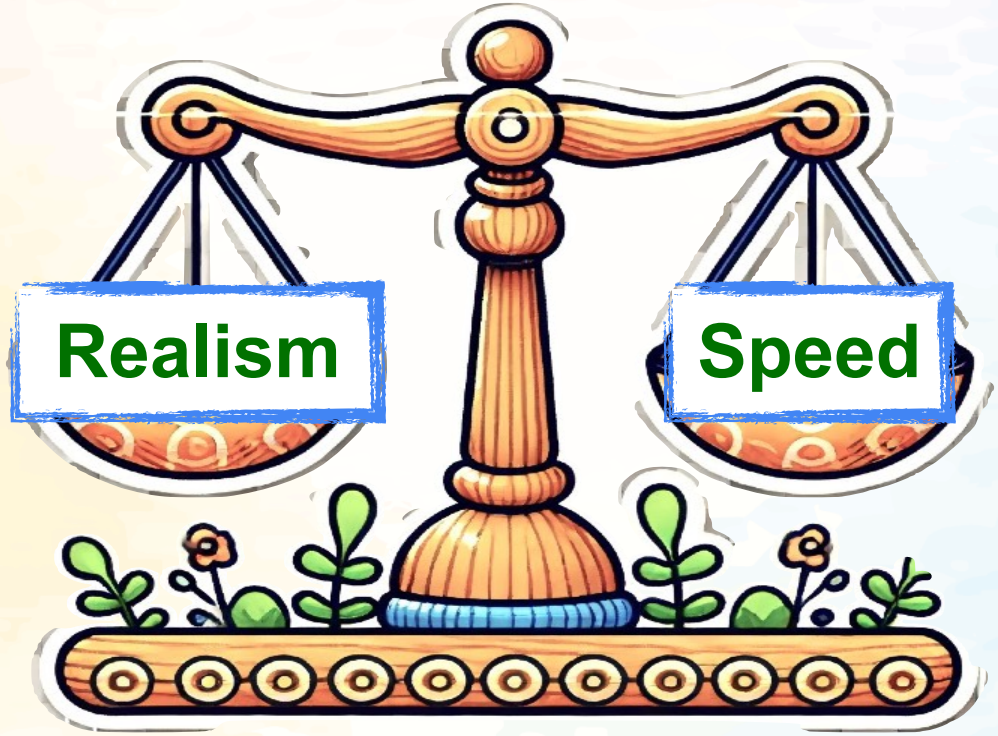

Cloth Rendering — Different Representations

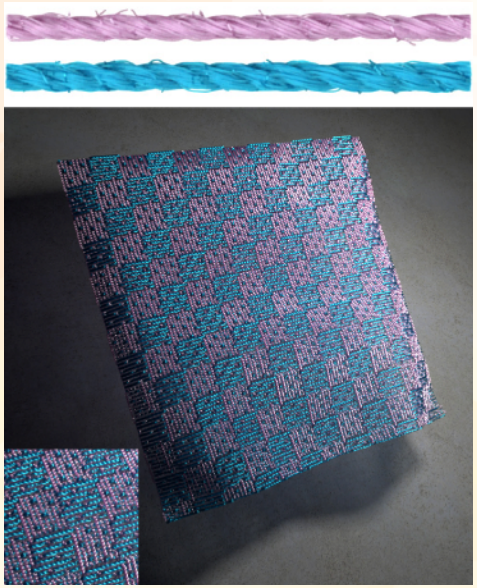

Fiber based

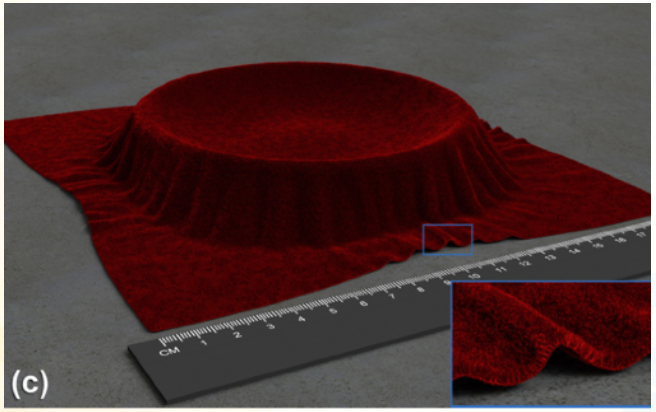

Volume based

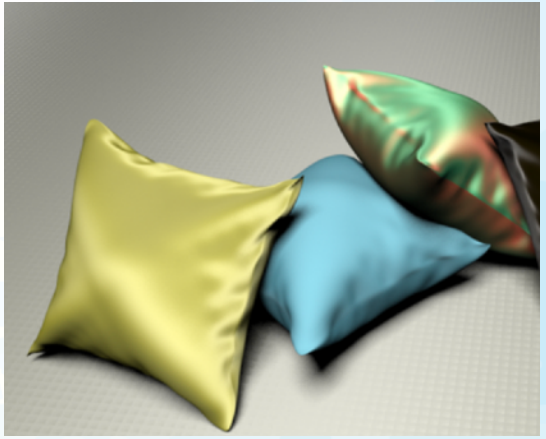

Surface based

Cloth as Fibers

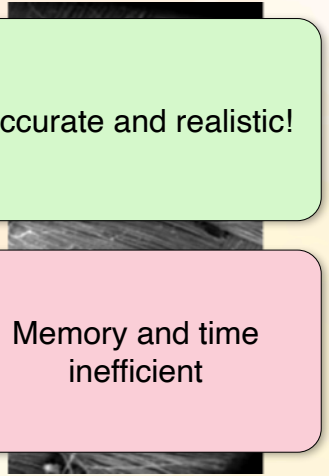

Accurate and realistic!

Memory and time inefficient

[Zhao et al. 2016]

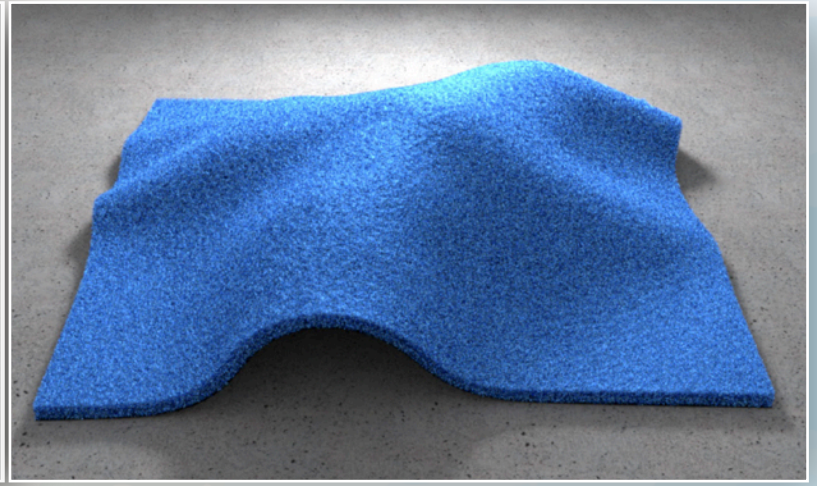

[Khungurn et al. 2015]

Cloth as Fibers — Key Idea

- BCSDF (Bidirectional Curve Scattering Distribution Function)
  - Evaluation:
    how a fiber interacts with the light
  - Importance sampling:
    how the light bounces off a fiber
    and scatters between fibers

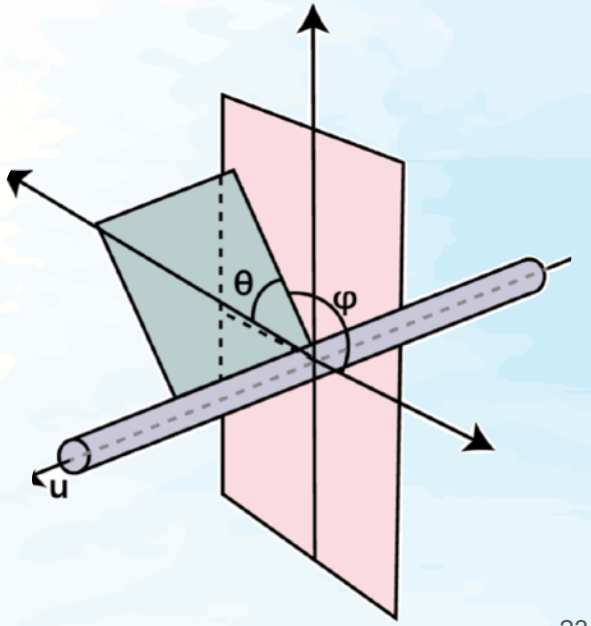

Cloth as Fibers — Key Problem

- The number of fibers is too high

300 M

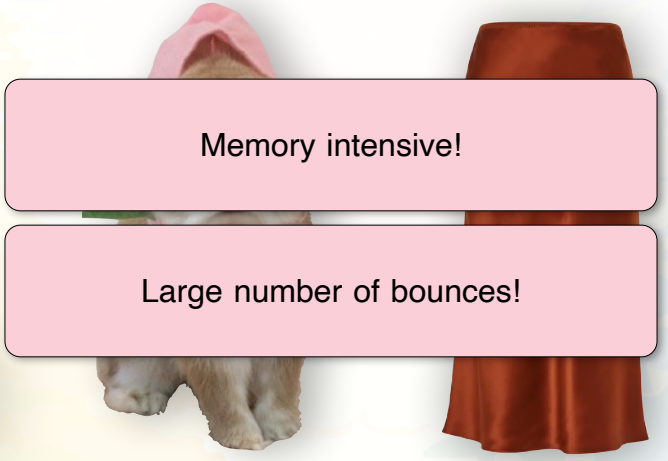

Memory intensive!

Large number of bounces!

45 M

4.6 M

17 M

Cloth as Volume

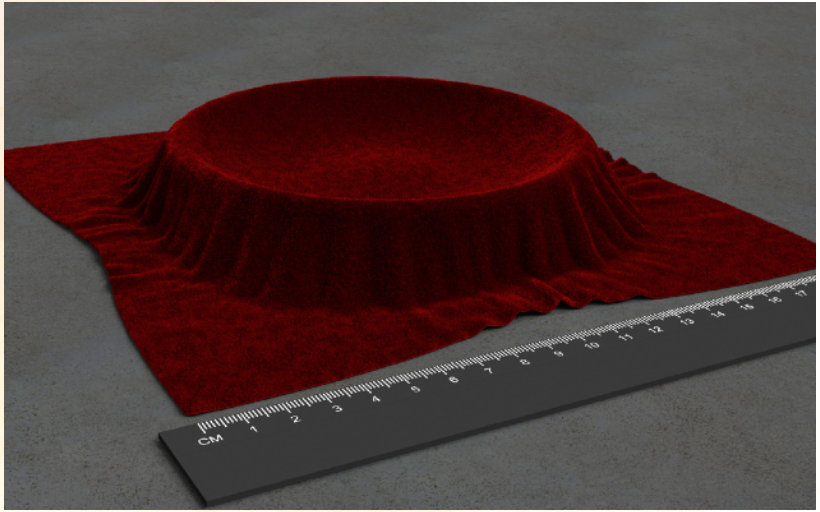

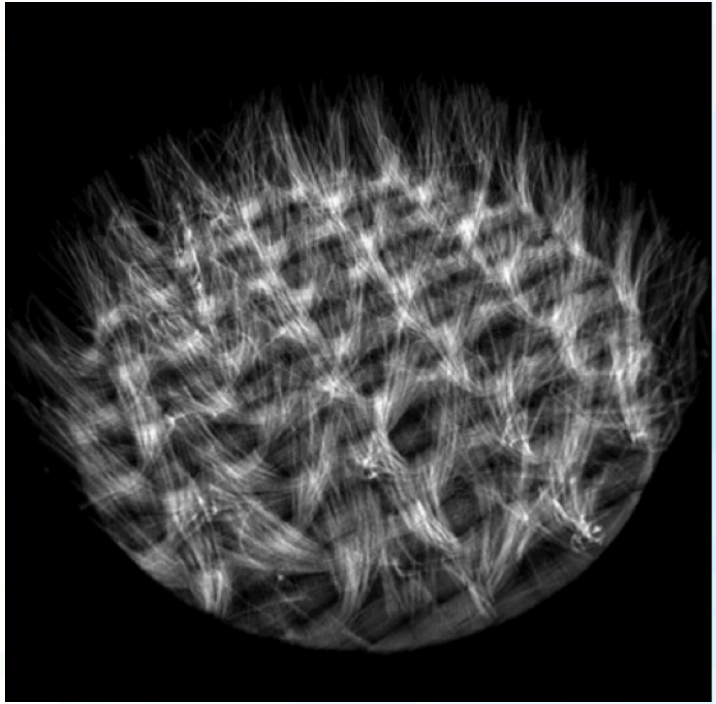

Cloth as Volume — Key Idea

- Scan/convert/represent cloth into volumes
- Use volumetric ray tracing thereafter
- Density and phase functions
- Classic, Microflake, SGGX

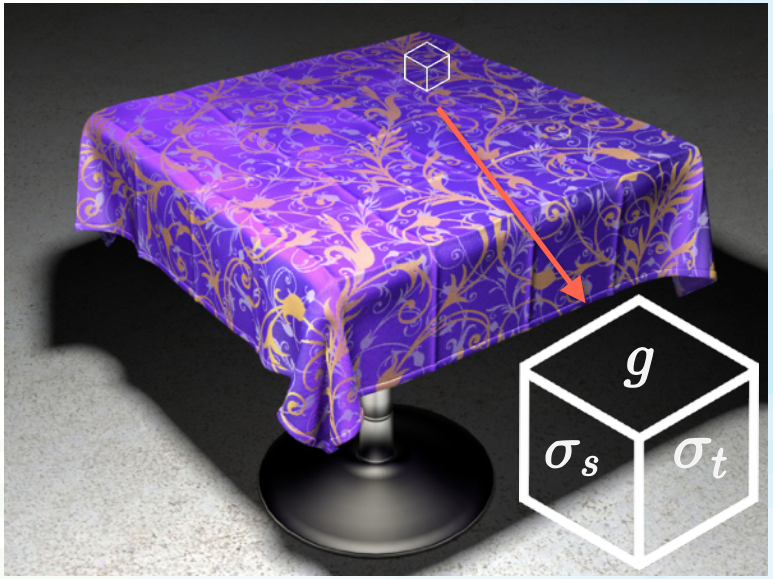

Cloth as Volumes — Pros and Cons

- Pros
  - Easy for renderer integration
  - Some effects are automatically handled
  - Potentially easier to support multi-scale rendering
- Cons
  - Not much faster than tracing fibers — optically dense
  - Still, usually impossible for real-time rendering

Cloth as Surfaces

Memory and time efficient!

Not realistic

Only thin cloth

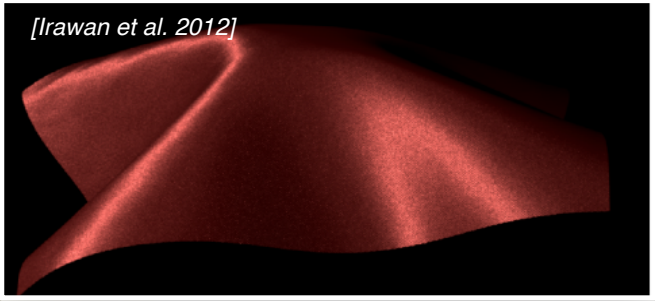

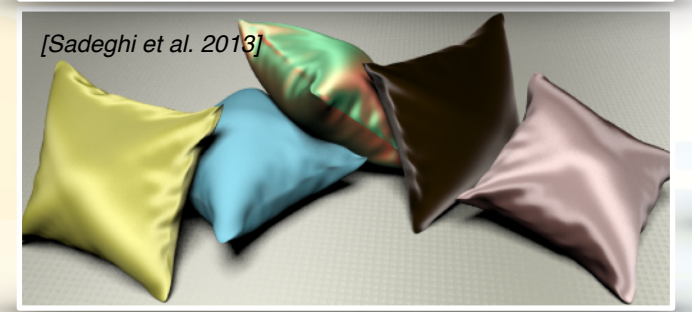

[Jin et al. 2022]

Key Idea 1 — Per-point Appearance

- Bidirectional Reflectance Distribution Functions (BRDFs)
  - Tells how a point reacts to the light
  - Reflections only
  - Recently extended to BSDF (BRDF + BTDF)

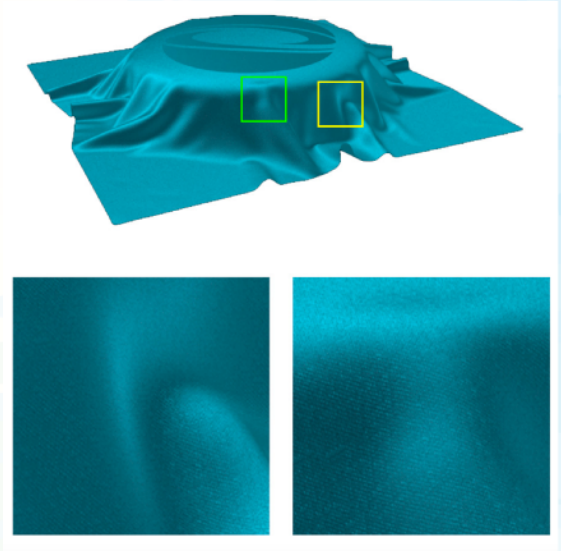

Problem 1 — Per-point Spatially-Varying BSDF

- Hit point (uv) -> fiber statistics
- Simulate the BCSDF on surface

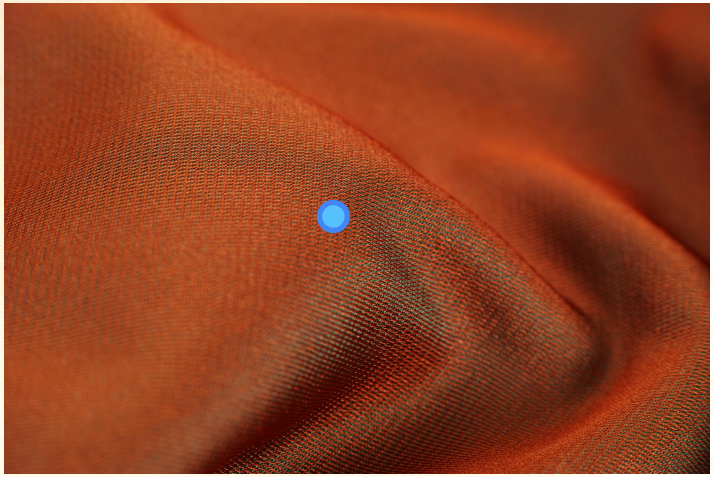

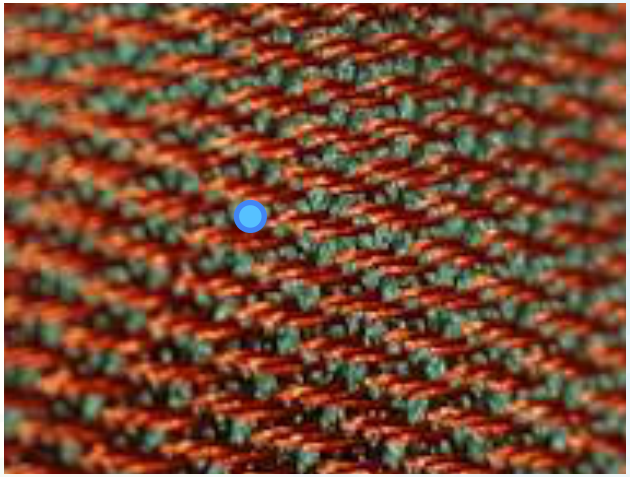

Key Idea 2 — Shading a point vs. Shading a pixel

- Modern solution cares about a pixel's appearance
  - Automatically anti-aliased
  - Noise-free
  - Automatically support multi-scale rendering

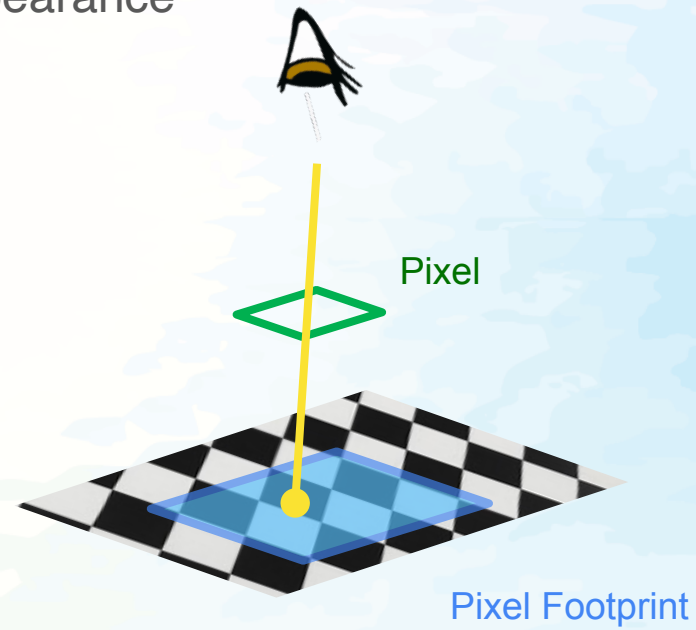

Problem 2 — Multi-scale Appearance Aggregation

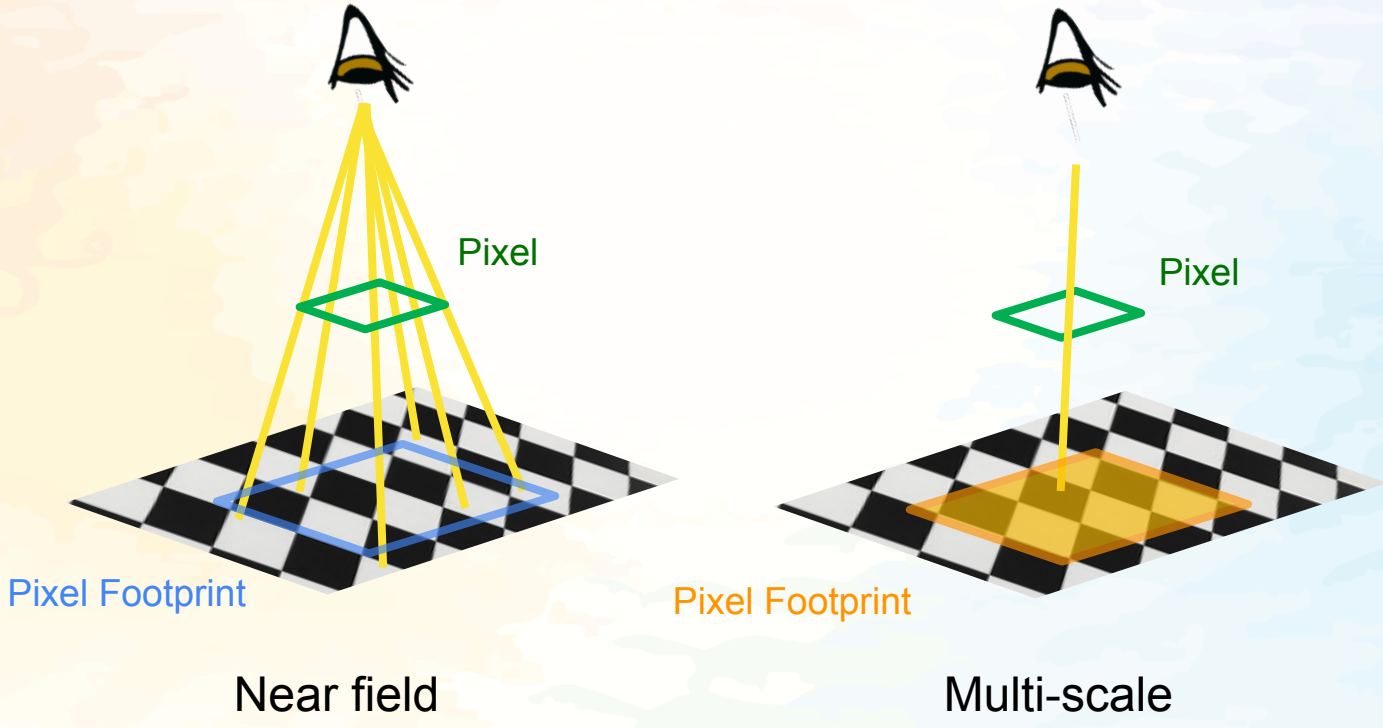

Problem 2 — Multi-scale Appearance Aggregation

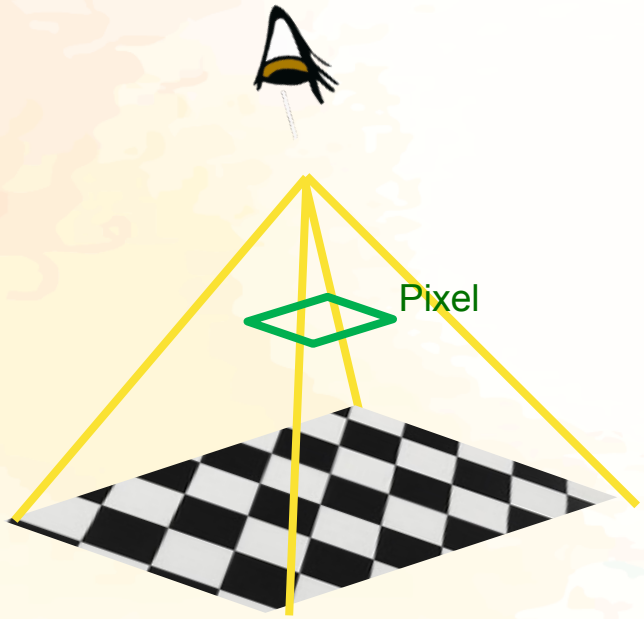

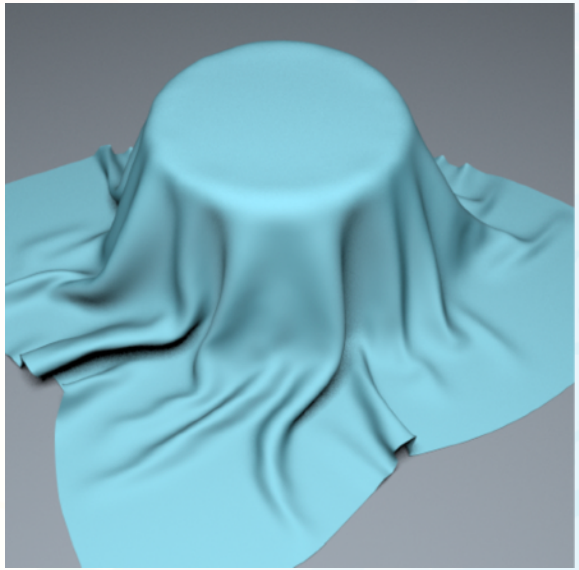

Far field

Next Part

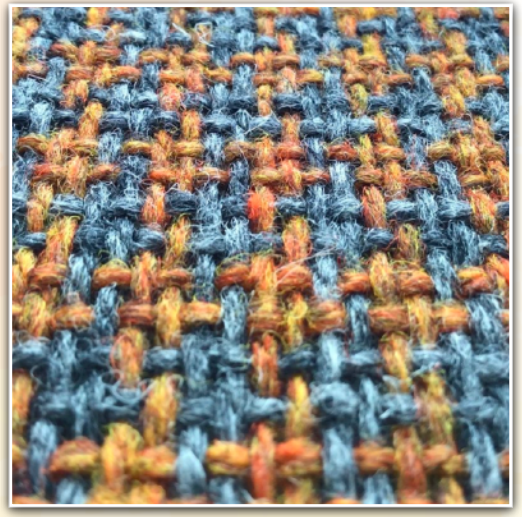

Cloth 101

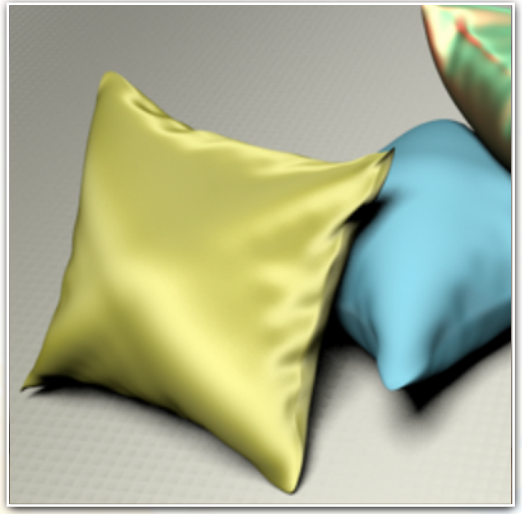

Cloth Rendering Models

Modern Cloth Rendering

Previous Cloth Rendering Models

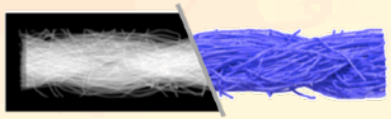

Fiber geometry

Rendering

Micro-appearance methods

Surface geometry

Rendering

Surface-based methods

Volumetric Representation

Input

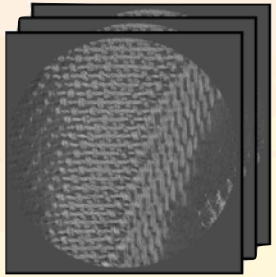

Micro CT Images

Real Cloth

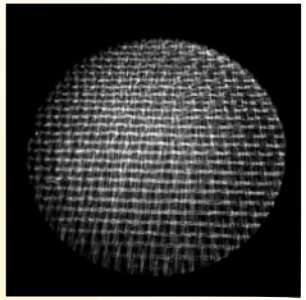

Reconstructed Density Field
and Orientation Field

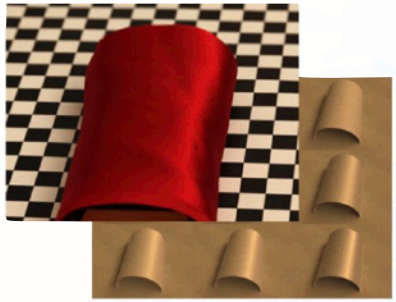

Appearance Matching

Output

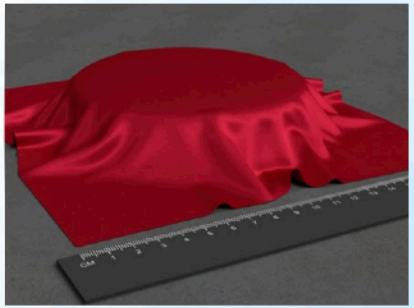

Rendered Results

Virtual Cloth

[Zhao et al. 2011]


[Zhao et al. 2011]
Satin / environment lighting

Volumetric Representation

Extremely realistic (models every single fiber and different kinds of cloth)

Requires **CT scan data** for preprocessing

Large data size (TB), long rendering times (hundreds of hours)

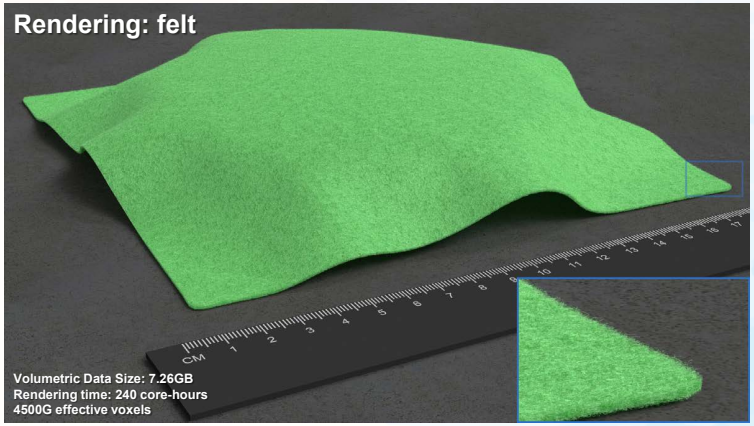

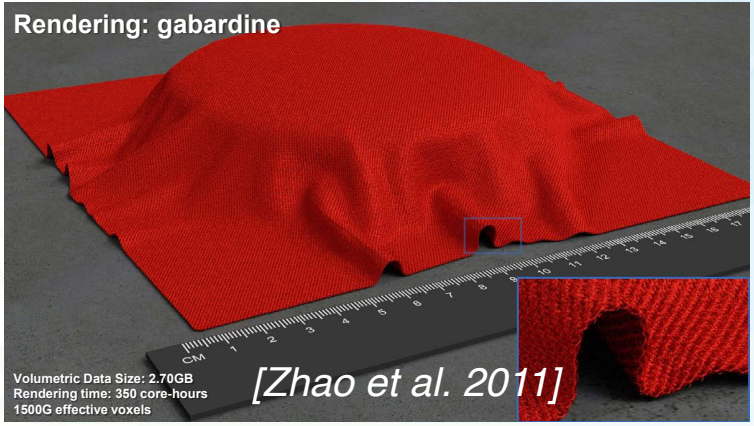

[Zhao et al. 2011]

Accelerated Volume Representation

- Main idea: Volume subdivided into coarse grids, with the statistical optical behavior within each grid averaged

Reduces storage, enabling larger fabrics rendering

No near-field

Rendering is still not efficient enough

[Schröder et al. 2011]

Volumetric Cloth Shading Models

- How to shade the volume cloth data:

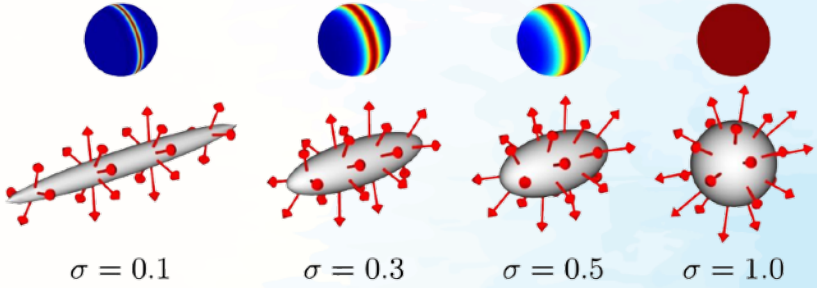

Micro-flake Model

[Jakob et al. 2010]

SGGX

[Heitz et al. 2015]

A more abstract mathematical description
Efficient and still realistic

Fiber-level Geometry Modeling

Modeling the fiber geometry using control parameters

- Twist, density, migration, fly-away fibers

Importance for realism

All fibers
in untied plies

Regular fibers

Flyaway fibers

[Zhao et al. 2016]

Fiber-level Geometry Modeling

Extremely realistic

No special data required. (Such as CT scans)

Large data size (TB), long rendering times

Difficult to match the real world

[Zhao et al. 2016]

Procedural Fiber-level Geometry Representation

- **Difficulty**: Converting procedural models into data-driven formats negates their compactness and makes them impractical for rendering.
- **Solution**: **Semi-implicit volumetric description** computes material properties on demand, treating cloth as an anisotropic medium.
- **Pros**: Efficient rendering with **reduced memory usage**, enabling practical handling of large procedural cloth models.
- **Cons:** Rendering inefficiency, still need to ray trace fiber-level geometry.

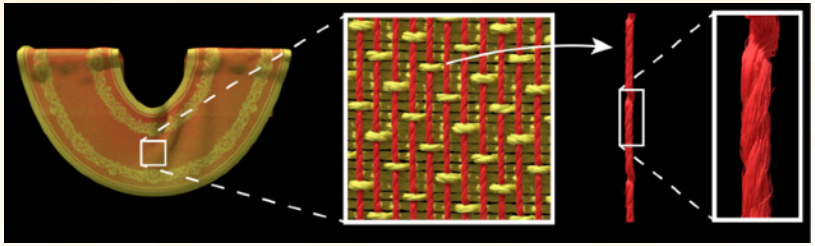

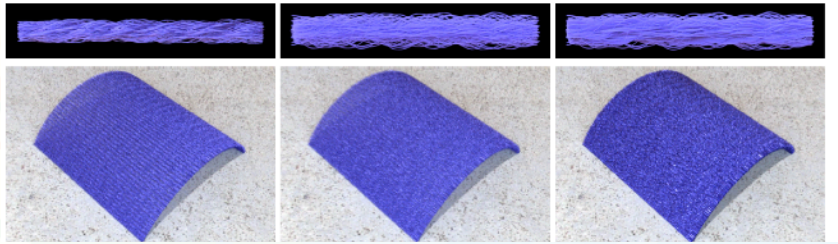

[Luan et al. 2017]

Micro-appearance Methods - Shading

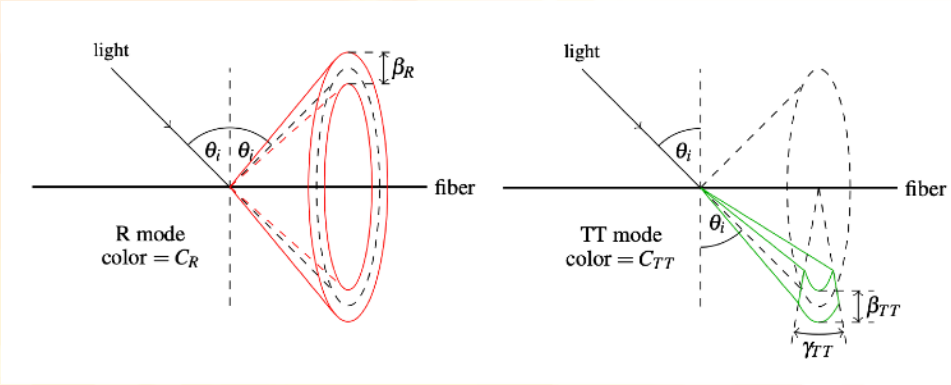

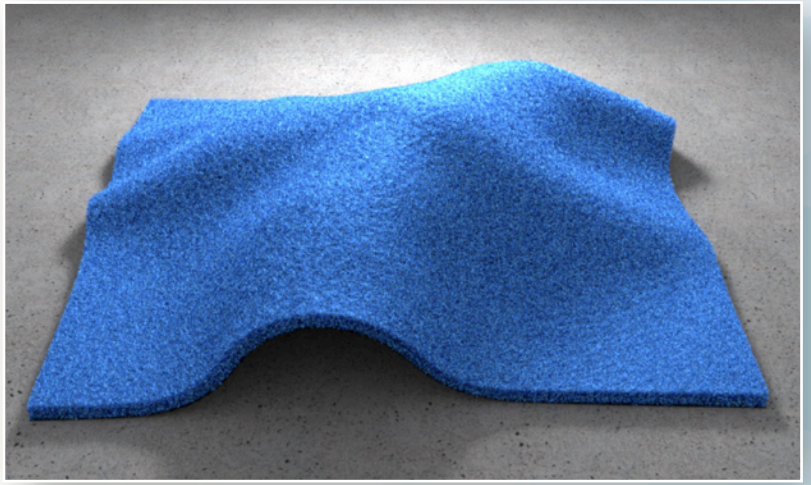

[Khungurn et al. 2015]

Micro-appearance Methods

Accurate and realistic!

Memory and time inefficient

[Zhao et al. 2016]

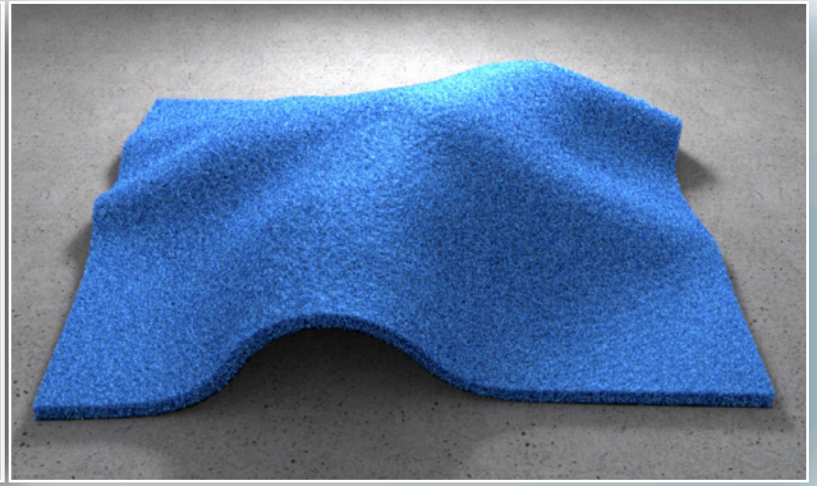

[Khungurn et al. 2015]

Cloth Rendering Models

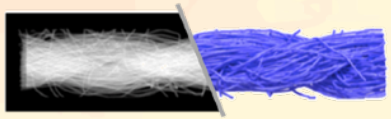

Fiber geometry

Rendering

Micro-appearance methods

+texture maps
+normal
+tangent

Surface geometry

Rendering

Surface-based methods

Industry

- Diffuse Cloth

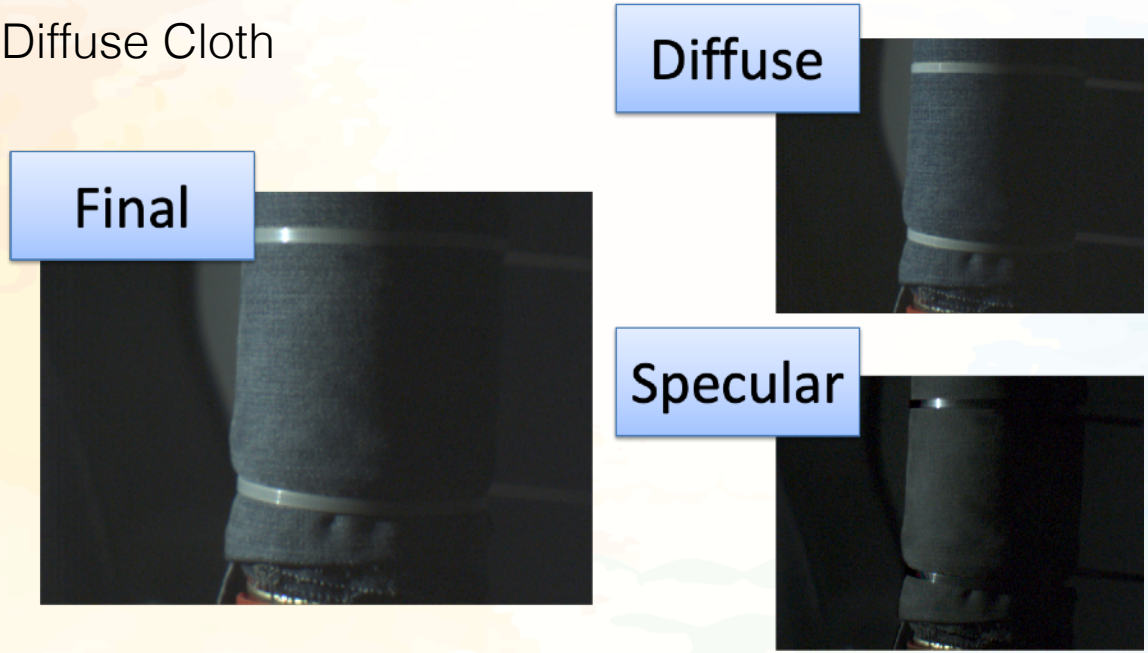

[Hable 2010]

Industry

- Shiny Cloth

Rim Lobe | **Inner Lobe** | **Remaining Diffuse** | **Rendering Result**

[Hable 2010]

Surface-based Methods

- Surface geometry + pattern parameters (uv tangent directions)
- Simulate the BCSDF on surface

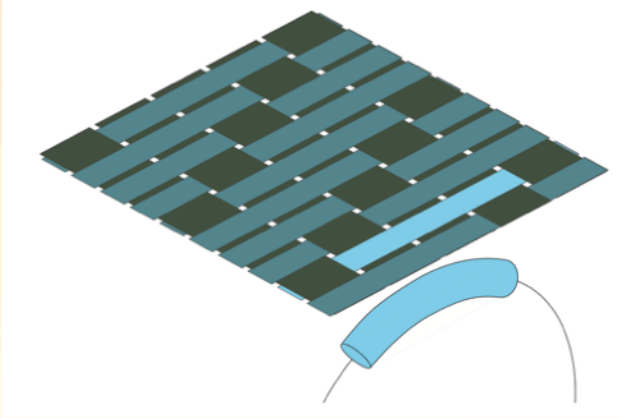

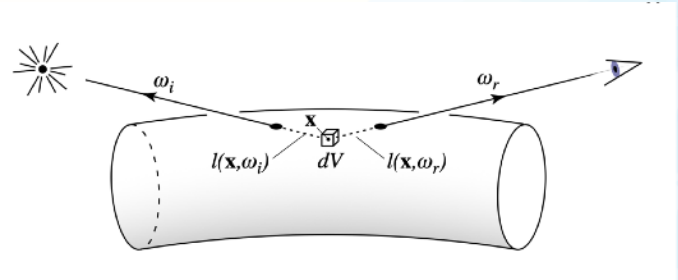

[Irawan et al. 2012]

Surface-based Methods

- Matching (surface reflectance only)

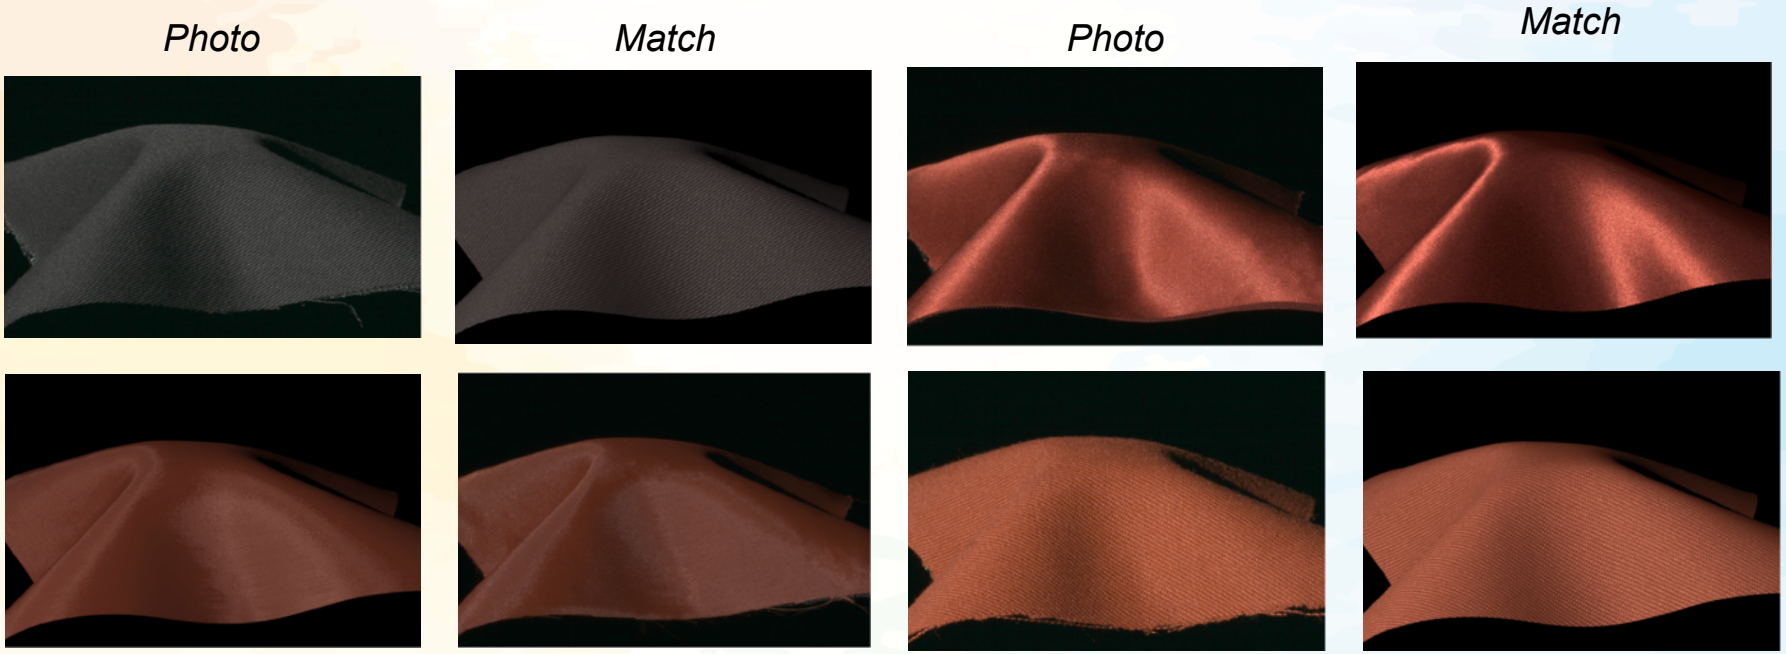

[Irawan et al. 2012]

Surface-based Methods

- Define the cloth as parameters
- Average the contribution (far-field only)

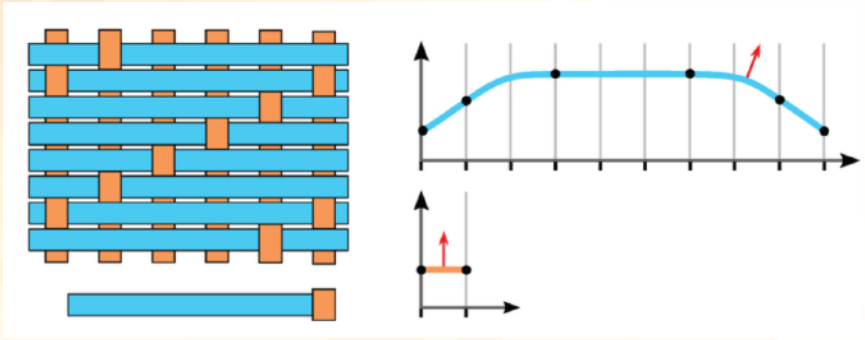

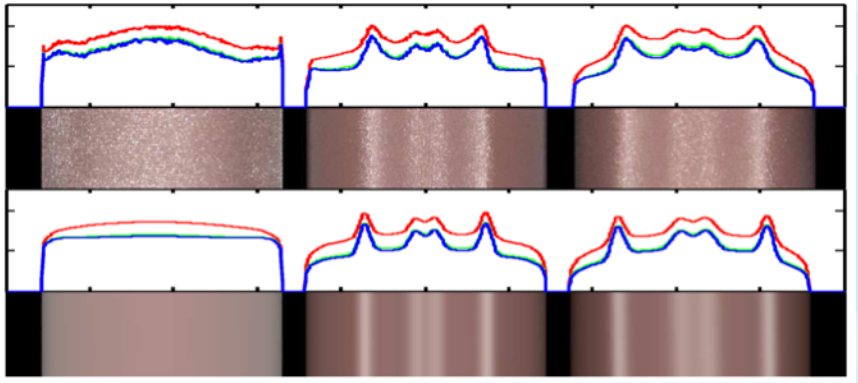

[Sadeghi et al. 2013]

Surface-based Model - Micro-cylinder

- Far-field only (no details)
- Surface reflectance only

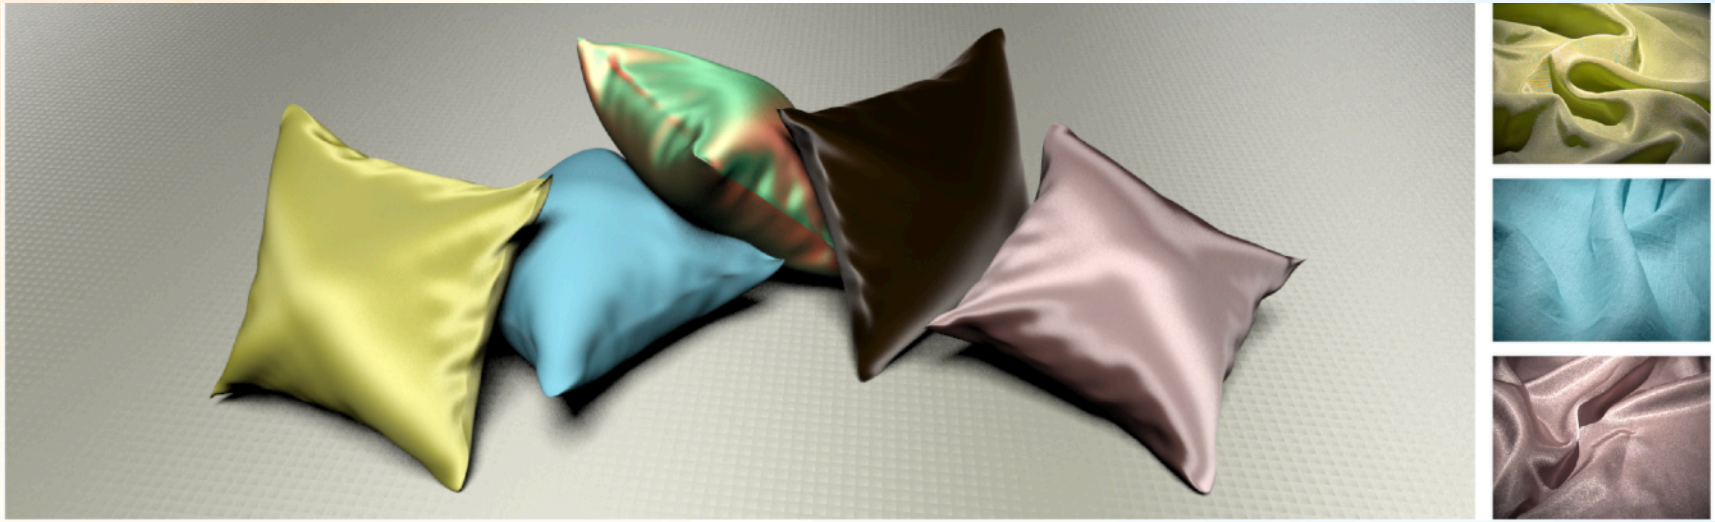

[Sadeghi et al. 2013]

Surface-based Methods

Memory and time efficient!

Not real enough

Only surface reflectance

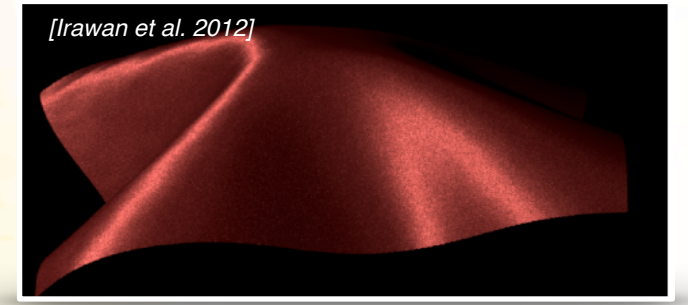

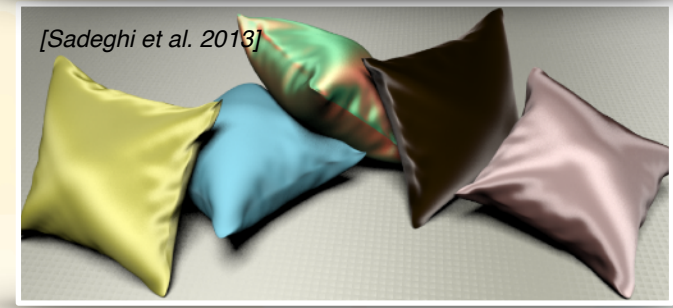

[Jin et al. 2022]

Real-time Fiber-level Cloth Rendering

- Simplify fiber geometry while preserving the details
- Level-of-detail Rendering
- Focus on the Knitted Cloth

Core Fibers LoD Self-shadows Ambient Occlusion

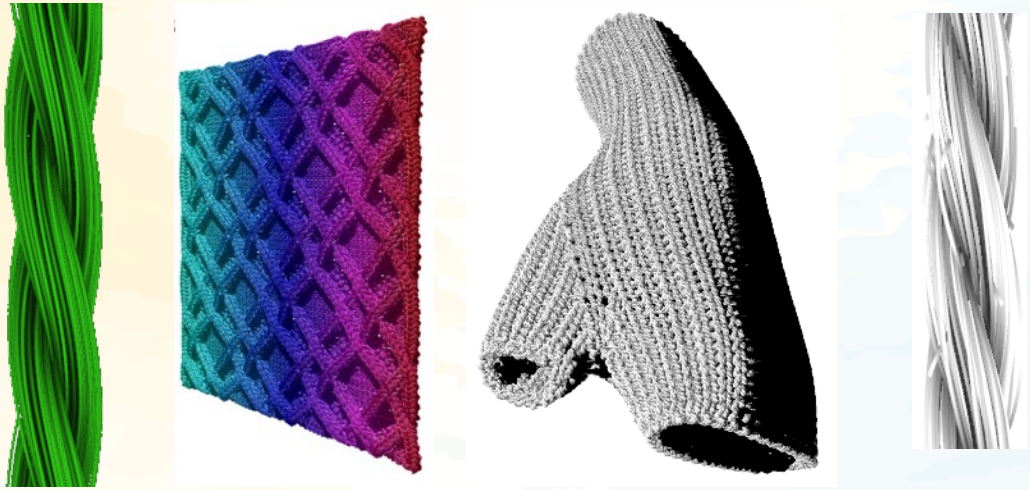

[Wu et al. 2017]

Real-time Fiber-level Cloth Rendering

Real-time Rendering

Fiber-level Details

Not Focusing on the Cloth Shading Model

[Wu et al. 2017]

Hybrid Cloth Rendering Method

Real-time Rendering

W/ Fly-away Fibers

Only surface reflectance

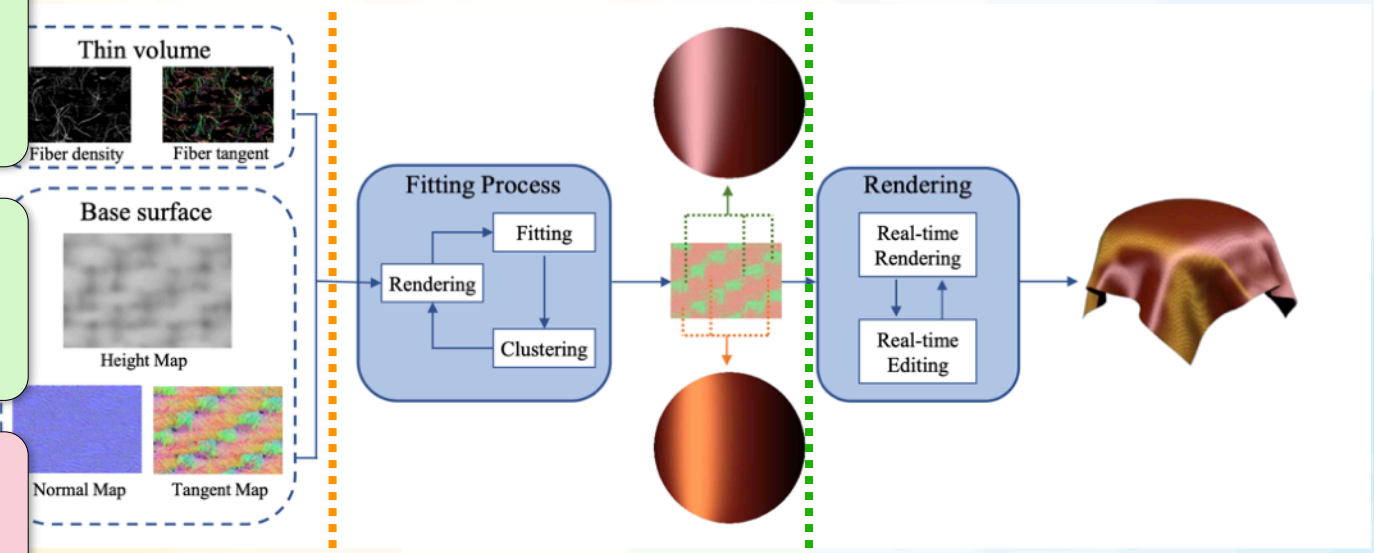

on and Separation

Clustering and Fitting

Rendering

[Xu et al. 2017]

Ideal Rendering Model

- Accurate / realistic (physically-based)
- Memory efficient
- Fast to render, better real-time
- General to render all types of cloth

Reduce #fibers!

Aggregate the appearance!

Appearance Aggregation

Level of details (LOD)

Why Aggregation?

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

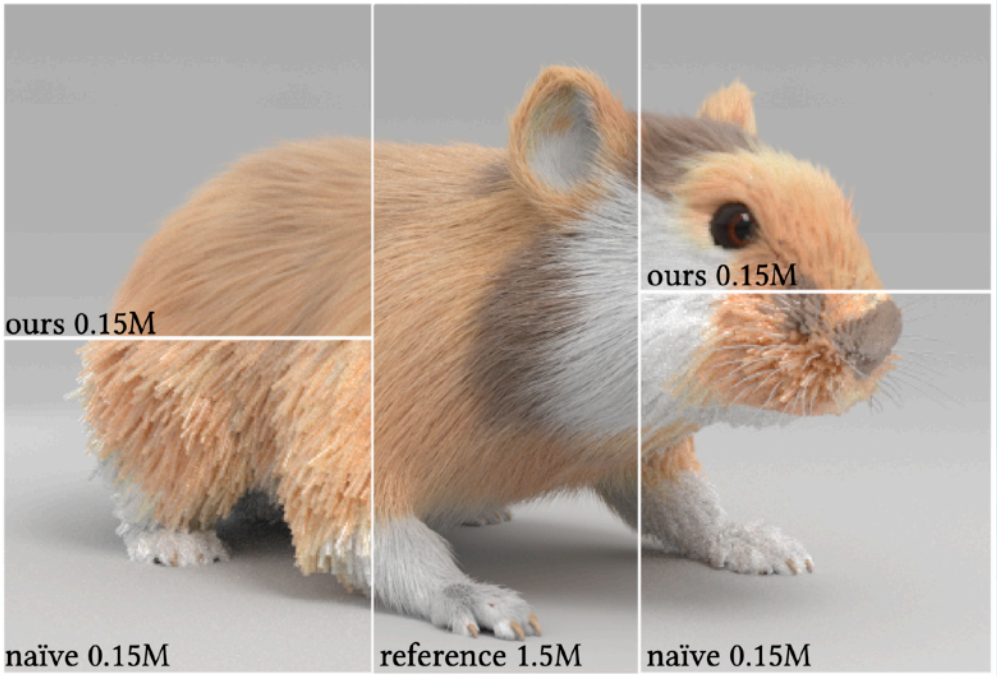

Outline

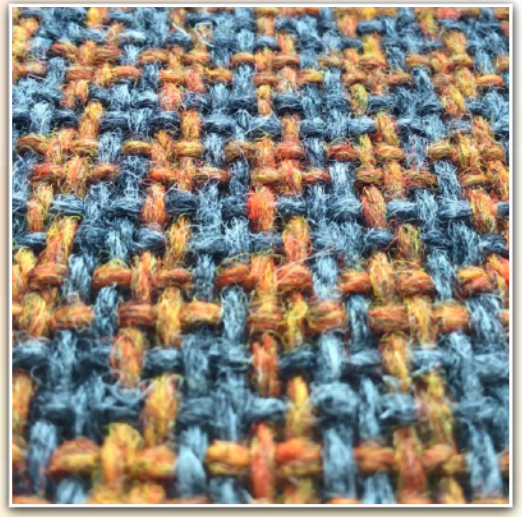

Cloth 101

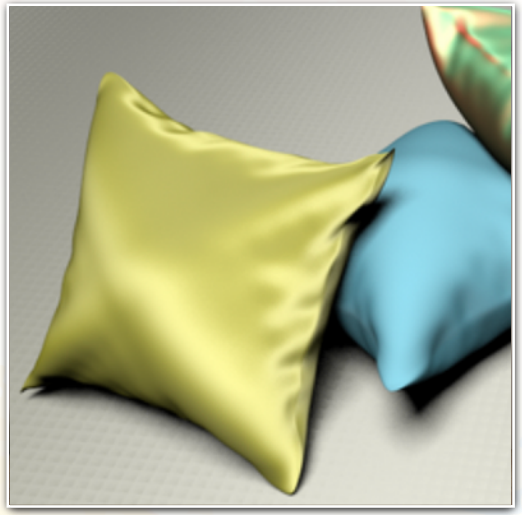

Cloth Rendering Models

Modern Cloth Rendering
(Recent 5 Years)

Ply-based Cloth Model

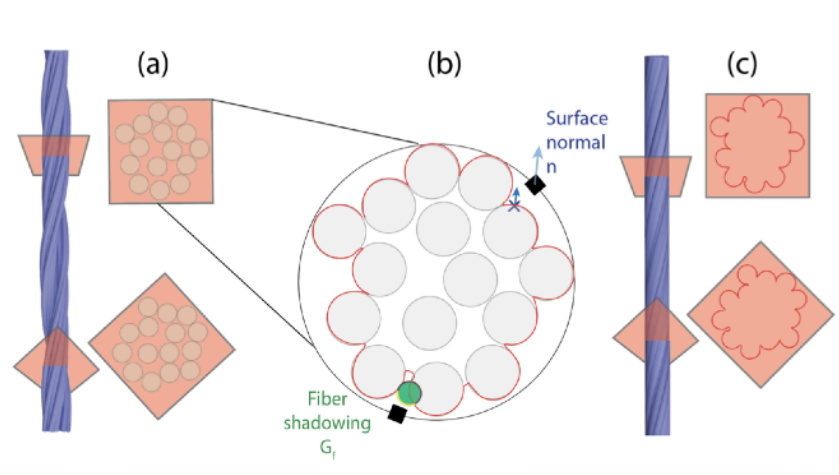

Ply geometry with fiber silhouette

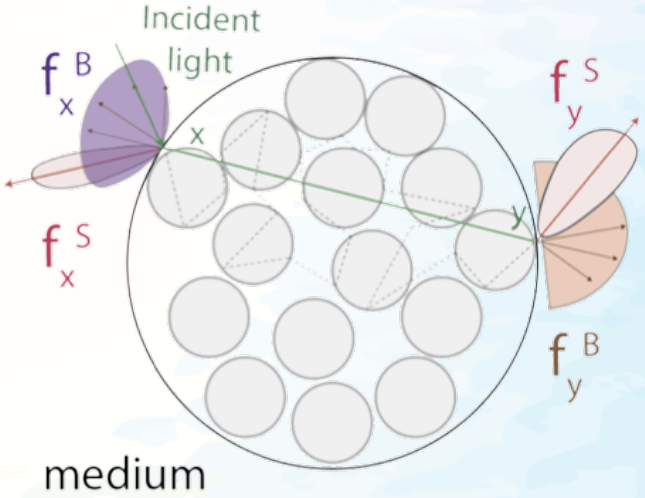

Ply-based shading model

[Montazeri et al. 2020]

Ply-based Cloth Model

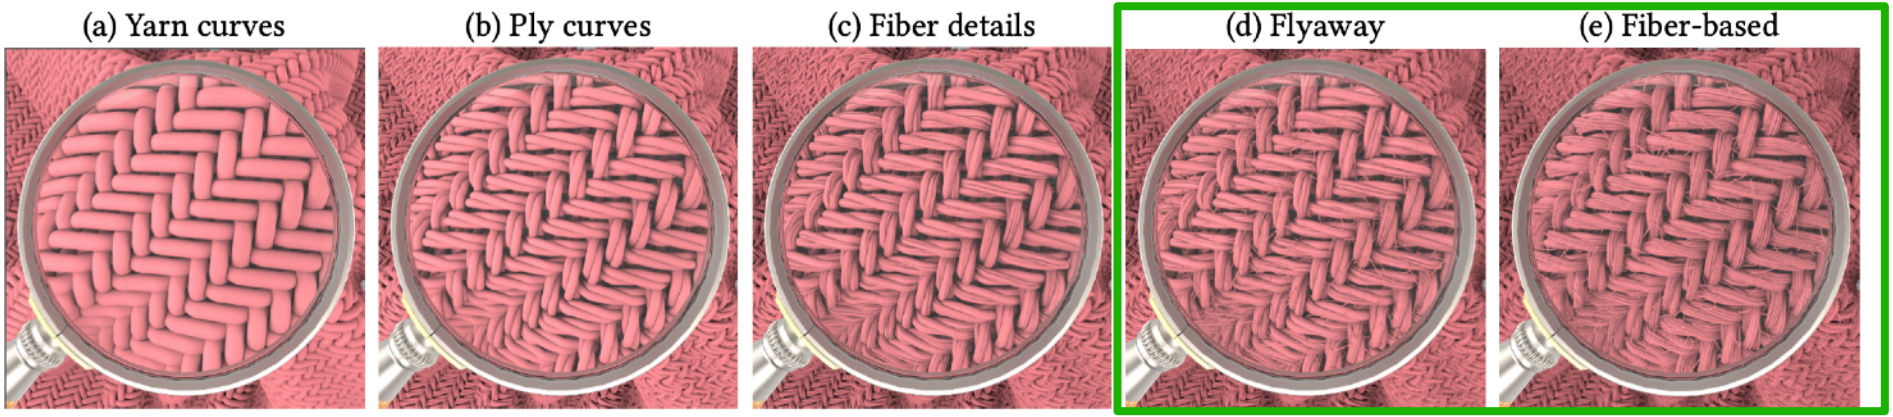

Ply-based Cloth Model

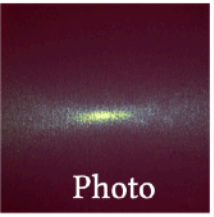

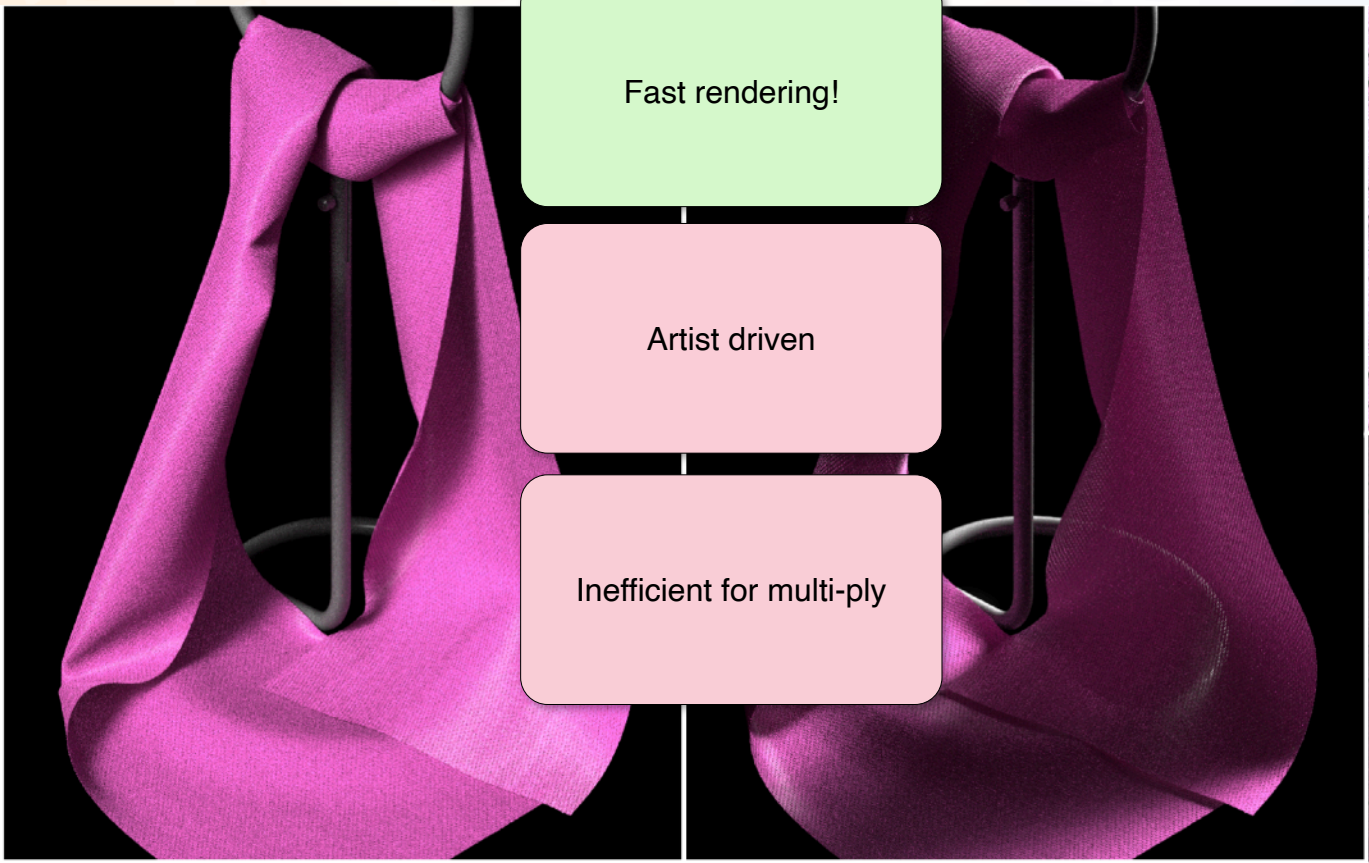

Fast rendering!

Artist driven

Inefficient for multi-ply

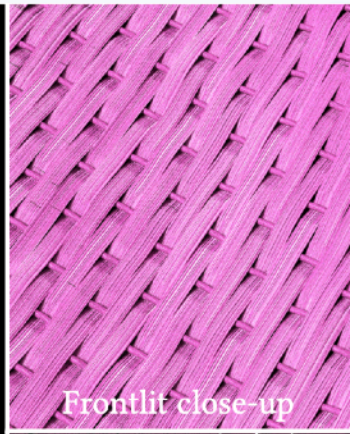

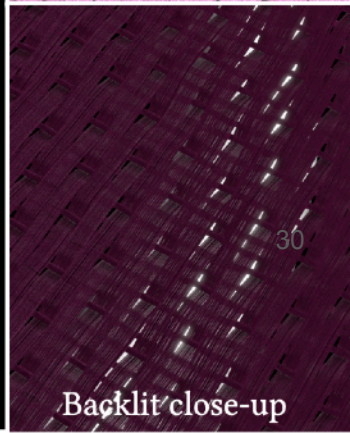

Yarn-based Cloth Model

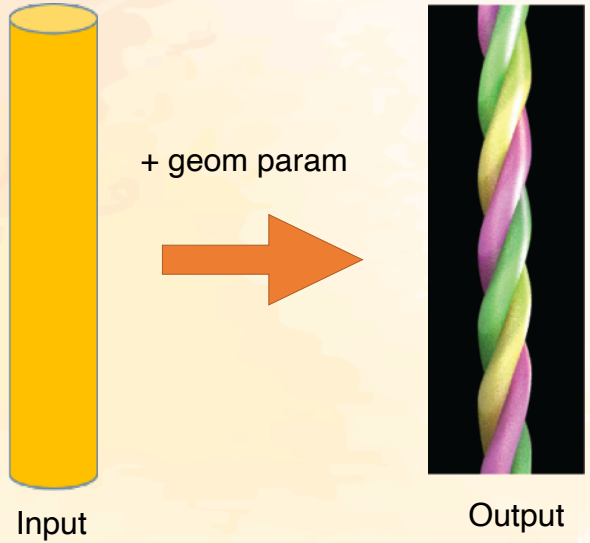

Geometric representation

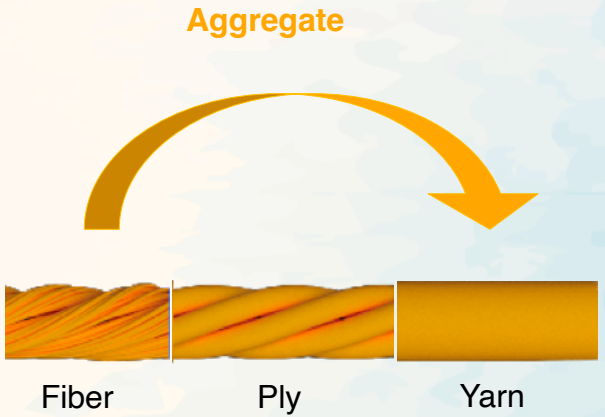

Shading

[Zhu et al. 2023]

Ply-level Shading Model — Aggregating Fibers

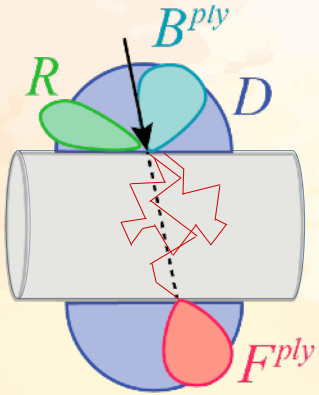

longitudinal

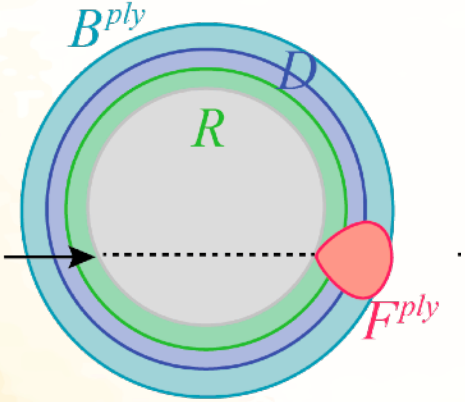

azimuthal

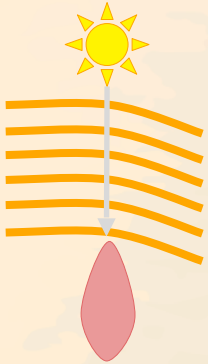

Forward scattering

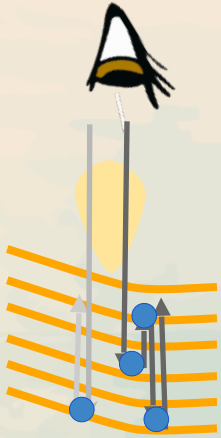

Backward scattering

Extended dual scattering

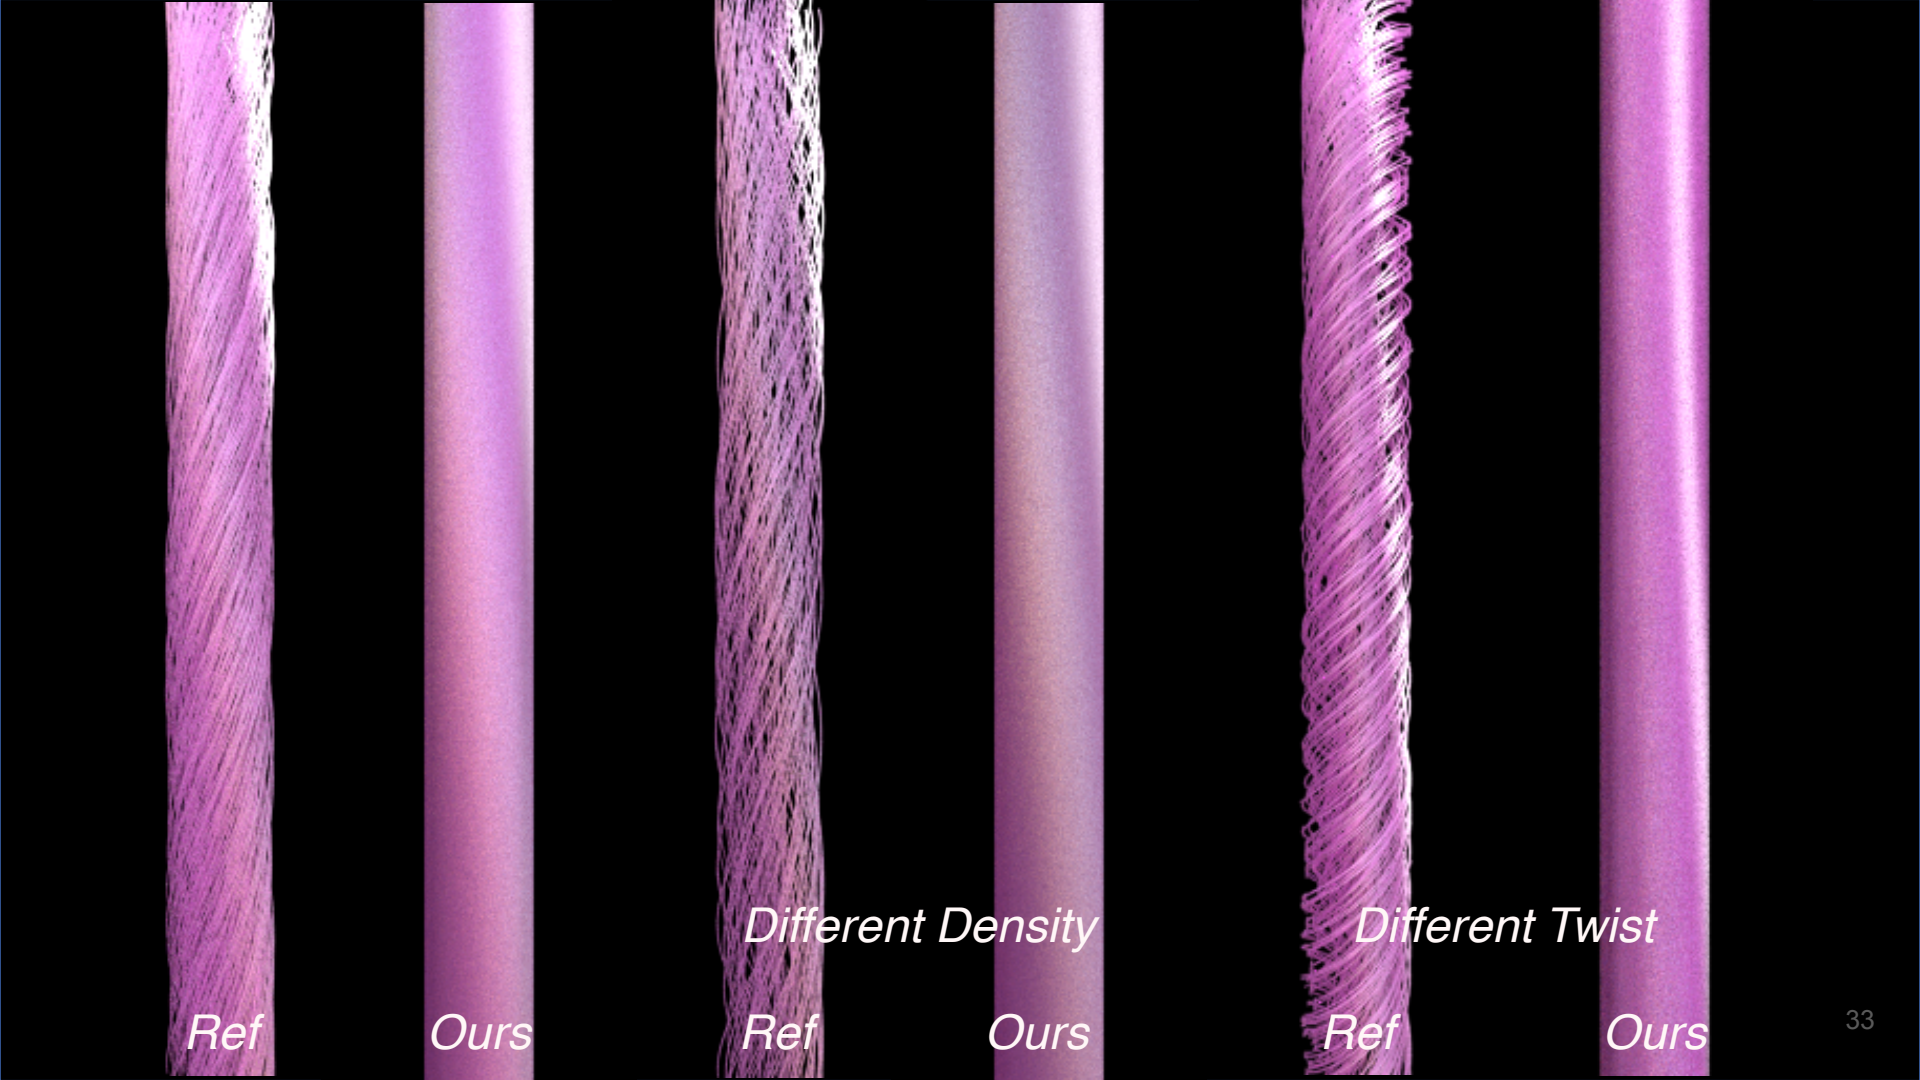
Different Density
Different Twist
Ref
Ours
Ref
Ours
Ref
Ours

Yarn-level Shading Model — Aggregating Plies

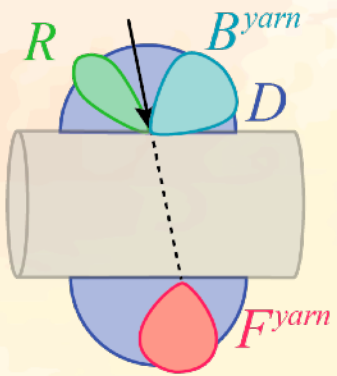

longitudinal

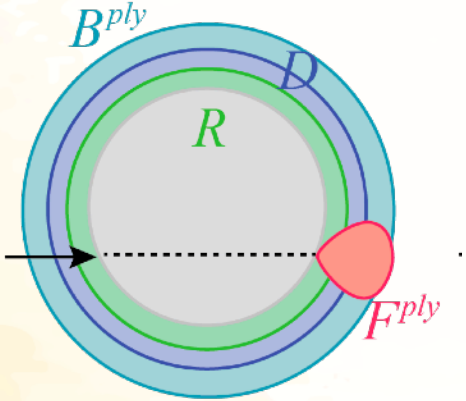

azimuthal

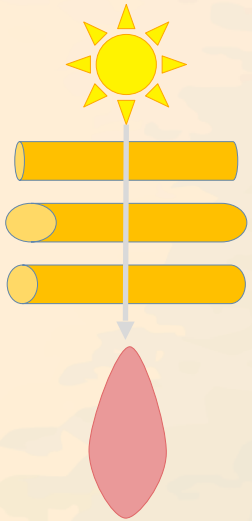

Forward scattering

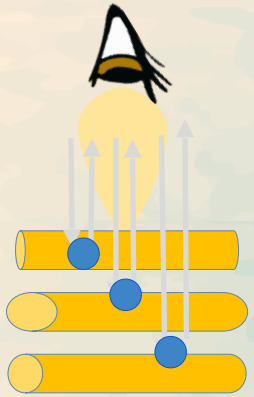

Backward scattering


Ref
[Zhu et al. 2023]
Different Density
Ref
[Zhu et al. 2023]
Different Twist
Ref
[Zhu et al. 2023]

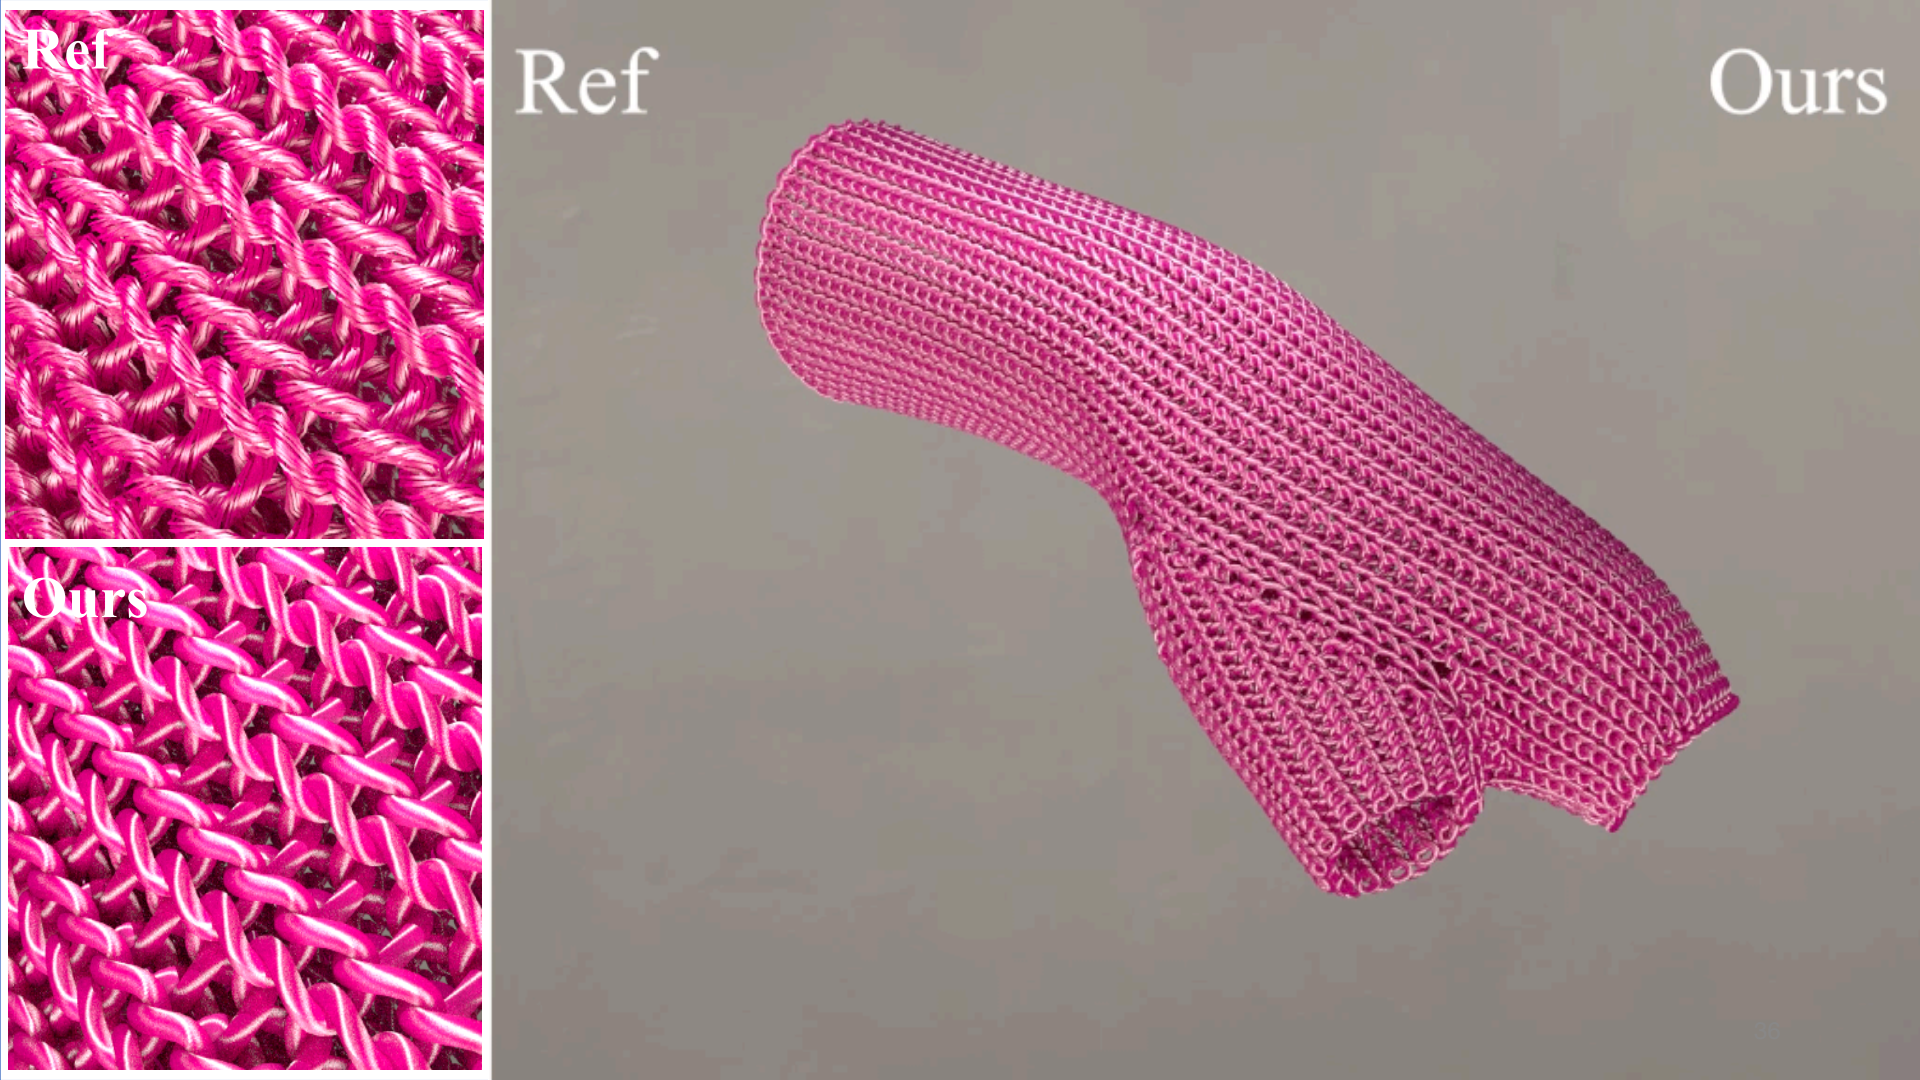
Ref
Ours
Ref
Ours

Yarn-based Cloth Model

Complete shading remake

- Sampled secondary ray, and
- (More) physically-based shading (rather than dual scattering)

Previous

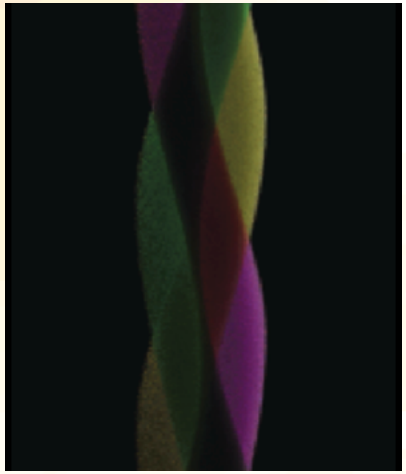

[Khattar et al. 2024]

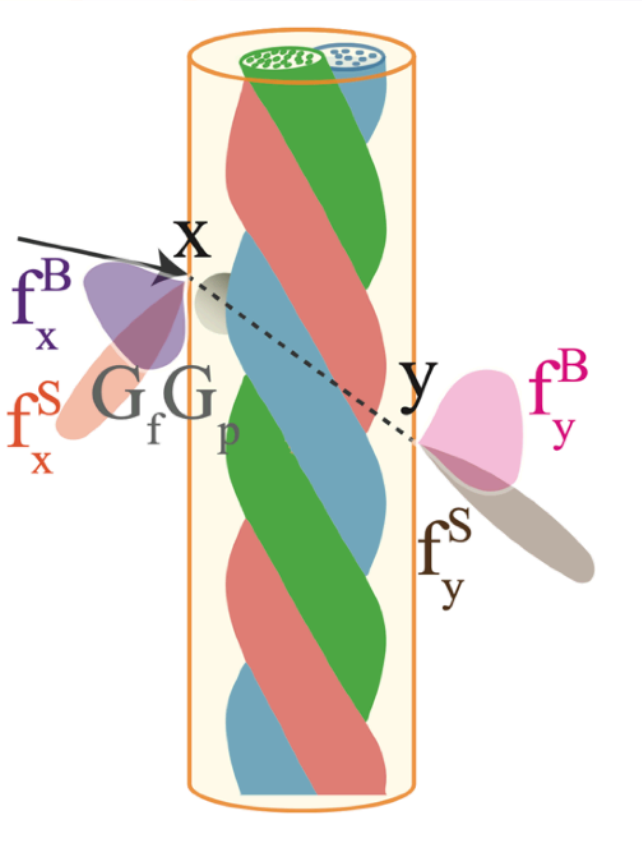

New

Conclusion

Accurate!

Fast rendering!

Memory efficient!

Physically-based!

General model!

Easy to edit!

Can we do better in performance?

Overview

- State-of-the-art surface-based cloth model
  - Realistic appearance modeling
  - Level of detail
    - Precomputation-based
    - Neural-based
- Cloth appearance reconstruction
  - Differentiable rendering
  - Generative AI
  - Data-driven reconstruction

Overview

- State-of-the-art surface-based cloth model
  - **Realistic appearance modeling**
  - Level of detail
    - Precomputation-based
    - Neural-based
- Cloth appearance reconstruction
  - Differentiable rendering
  - Generative AI
  - Data-driven reconstruction

Surface-based cloth model

- Advantages:
  - Strikes a balance in **realism, storage** and **performance**
  - Easier for **reconstruction**
- State-of-the-art example:
  - *A realistic surface-based cloth rendering model [Zhu et al. 2023]*

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Realism: Accounts for complex visual effects.
    - **Asymmetric reflection highlight**

Photo

Render

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Realism: Accounts for complex visual effects.
    - **Cross shaped transmission highlight**

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Realism: Accounts for complex visual effects.
    - **Transparency**

Render

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Yarn-level geometric maps + yarn-level shading

Photo reference

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Storage: Yarn-level **geometric maps**

Mesh

ID

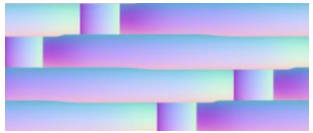

normal

tangent

position

Minimally tileable maps

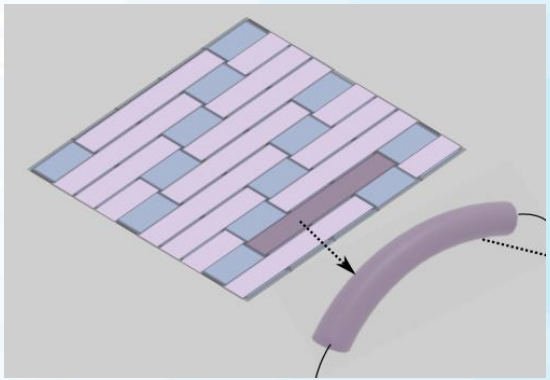

Yarn segment from maps

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Performance: Yarn-level **shading**

Yarn

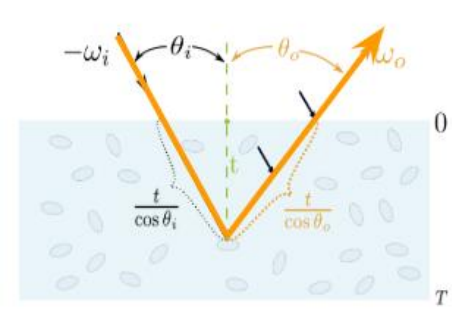

Micro-flake volume
Spongcake [Wang et al. 2020]

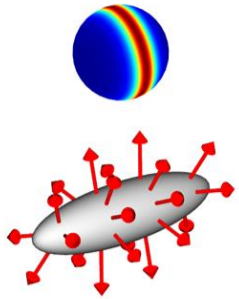

Fiber-like normal distribution
SGGX [Heitz et al. 2015]

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Yarn-level geometric maps + yarn-level shading

Photo

Render w/ [Zhu et al. 2023]

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Realism: **Procedural noise for irregularity.**

w/o noise w/ noise

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Realism: **Shadowing and masking**, considering per-pixel effective area.

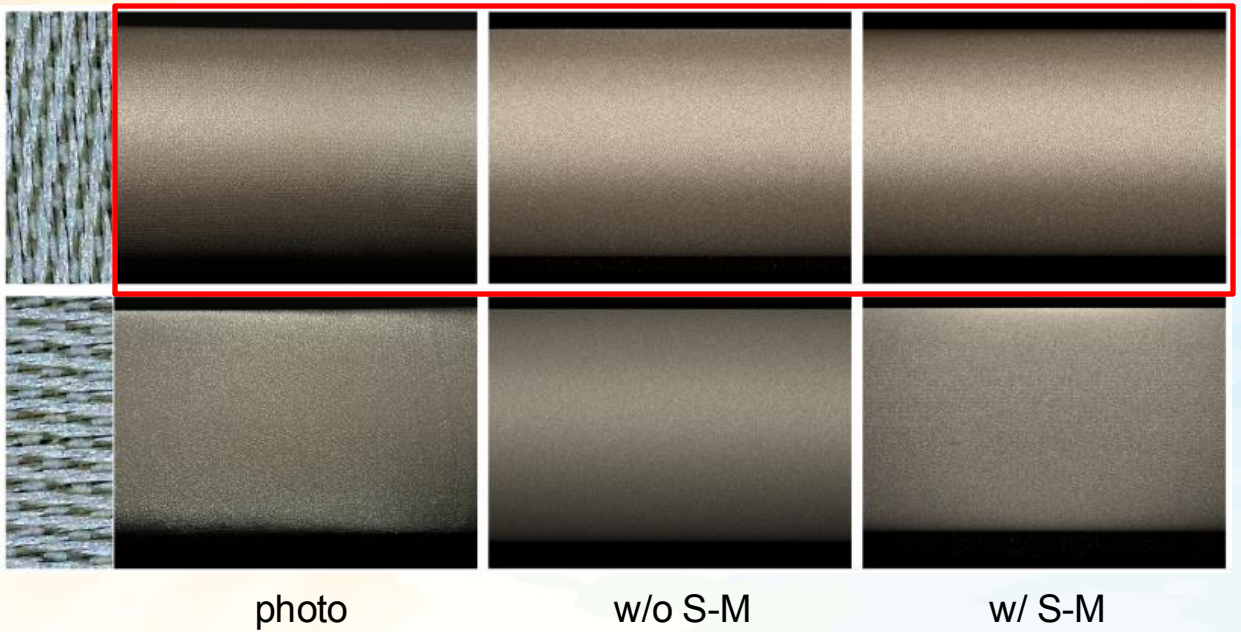

Surface-based cloth model

- *A realistic surface-based cloth rendering model [Zhu et al. 2023]*
  - Realism: **Shadowing and masking**, considering per-pixel effective area.

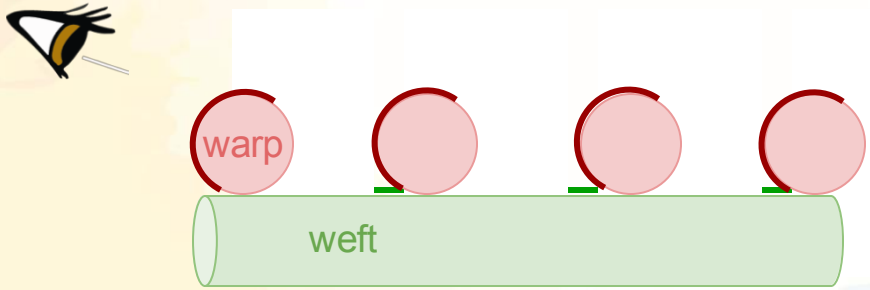

We need to consider an area instead of a point!

Effective area is varying at different viewing angle because of the masking effect between plies.

Surface-based cloth model

Photo

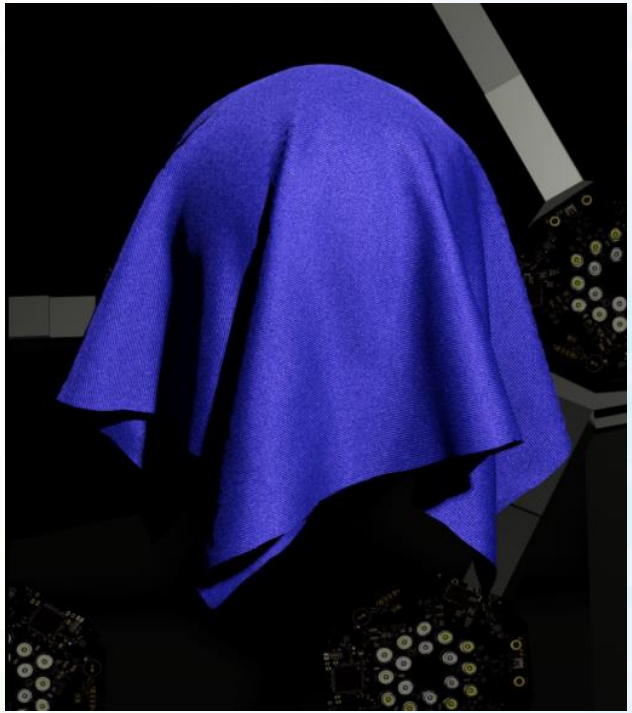

Render w/ [Zhu et al. 2023]

Overview

- State-of-the-art surface-based cloth model
  - Realistic appearance modeling
  - **Level of detail**
    - Precomputation-based
    - Neural-based
- Cloth appearance reconstruction
  - Differentiable rendering
  - Generative AI
  - Data-driven reconstruction

Cloth Level-of-Details

- Multi-scale problem
  - Diffuse from far away but very sparkly close up
  - Requires high SPP (samples-per-pixel) to resolve such sparkles

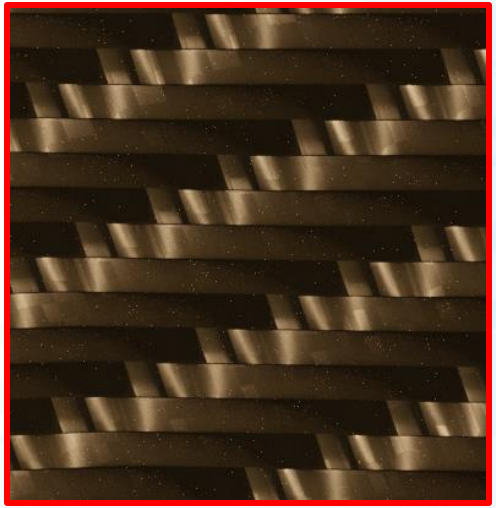

Cloth Level-of-Details

- Multi-scale problem
  - **1K** SPP. For real-time applications we need to reduce it to **1** SPP.

Photo

Render w/ [Zhu et al. 2023]

Cloth Level-of-Details

- Goal:
  - Integrate BRDF * visibility over a pixel footprint
  - As few samples as possible (1K SPP -> 1 SPP)

$$F_{\mathcal{P}} = \int_{\mathcal{P}} \hat{f}\left(x(p), \mathrm{i}, \mathrm{o}\right) V(x(p), \mathrm{i}) \mathrm{d}p$$

pixel footprint — BRDF — visibility

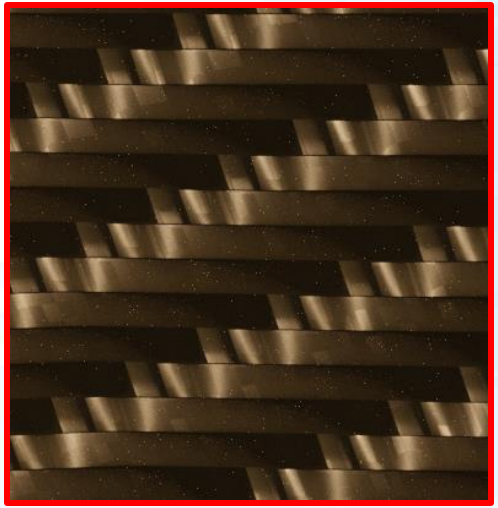

Cloth Level-of-Details

- *A multi-scale surface-based cloth appearance model [Zhu et a. 2024]*
  - Integration of BRDF only first, ignoring visibility.
  - **Precomputed Summed area table.** O(1) time complex.

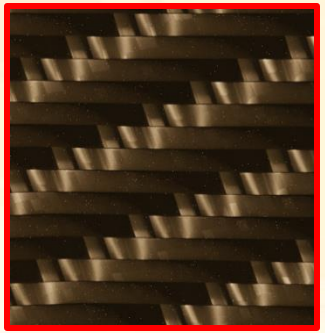

BRDF

31	2	4	33	5	36
12	26	9	10	29	25
13	17	21	22	20	18
24	23	15	16	14	19
30	8	28	27	11	7
1	35	34	3	32	6

Original data

31	33	37	70	75	111
43	71	84	127	161	222
56	101	135	200	254	333
80	148	197	278	346	444
110	186	263	371	450	555
111	222	333	444	555	666

SAT

15 + 16 + 14 + 28 + 27 + 11 =
101 + 450 − 254 − 186 = 111

4 queries

Cloth Level-of-Details

- *A multi-scale surface-based cloth appearance model [Zhu et a. 2024]*
  - Take into account visibility by using **control variates**

Base term

Residual term

$$\int_{\mathcal{P}} \hat{f}(p) \cdot v(p)\,\mathrm{d}p = \int_{\mathcal{P}} \alpha \cdot \hat{f}(p)\,\mathrm{d}p + \int_{\mathcal{P}} \hat{f}(p) \cdot (v(p) - \alpha)\,\mathrm{d}p$$

Precomputed

Monte Carlo sampling

base residual full

Mogwai

File View Help

120.2 FPS (8.3 ms/frame)

Graphs

		0.300	Diffuse Weight
0.025	0.190	0.120	Specular Reflecta
		0.400	Specular Roughne
		0.000	Ply Twist Angle
		0.000	Ply Flattening
0.000	0.000	0.000	Flyaway Sheen Co
		0.050	Flyaway Thicknes
		0.500	Flyaway Roughnes
		0.000	Texture Rotation
1000.000	1000.000		Texture Scale
		0	SAT Mode

Ply Normal Texture: C:\Users\mattchiang\

Texture info: 256x256 (RGBA32Float)

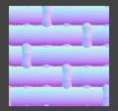

Remove texture

Ply Tangent Texture: C:\Users\mattchiang

Texture info: 256x256 (RGBA32Float)

Remove texture

Spline ID Texture missing!

Spline UV Texture missing!

▶ Grid volumes

▶ Statistics

Point light (Previous)

Cloth Level-of-Details

- *Real-time neural woven fabric rendering [Chen et al. 2024]*
  - **Encoder** takes in a woven material (geometric maps + shading parameters)
  - **Decoder** takes in a query (viewing + lighting + footprint) and returns BRDF value

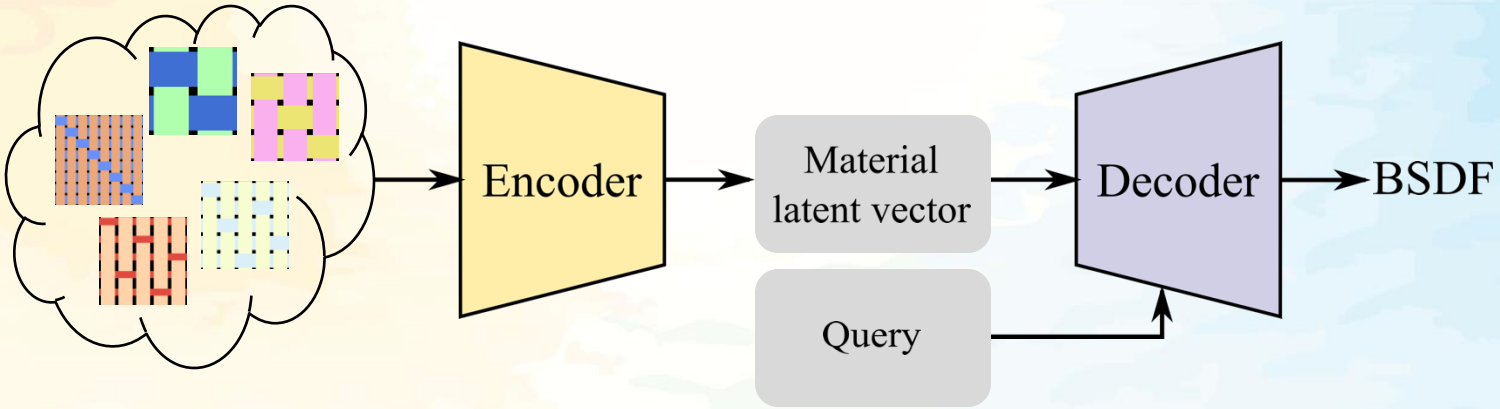

Cloth Level-of-Details

- *Real-time neural woven fabric rendering [Chen et al. 2024]*
  - **Encoder** is universal
  - Improvement over NeuMIP [] which is trained per-material
  - Data set: 7 patterns, 475 materials.

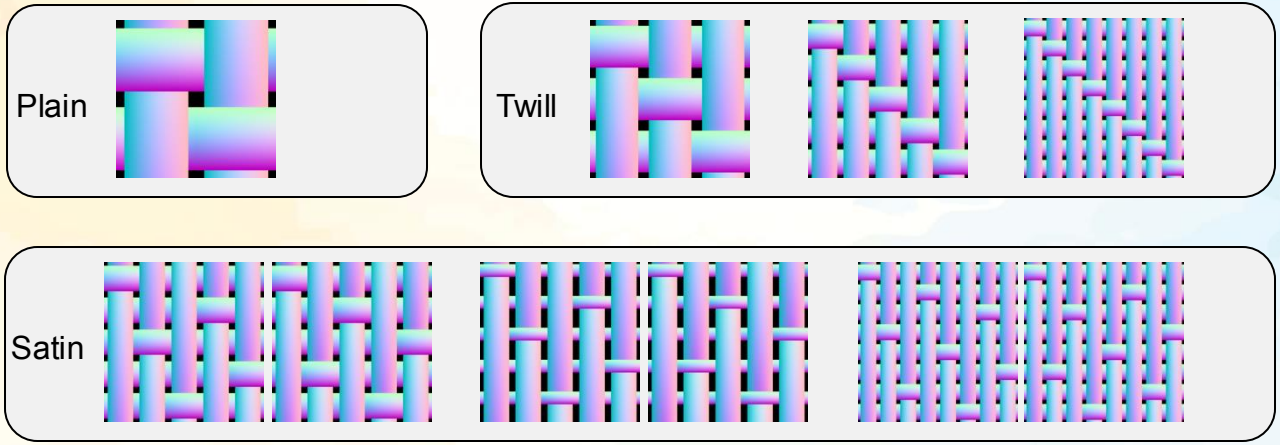

Cloth Level-of-Details

- *Real-time neural woven fabric rendering [Chen et al. 2024]*
  - **Decoder** is efficient
  - Real-time achieved by BRDF decomposition

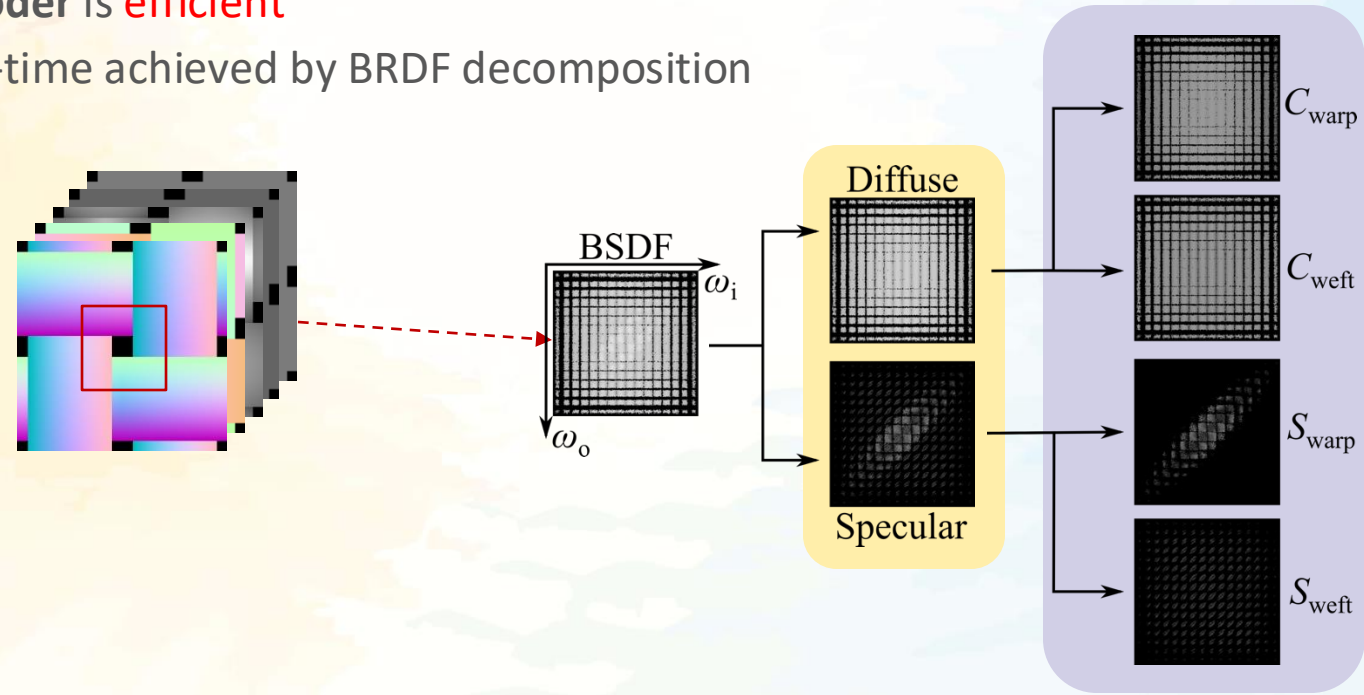

Cloth Level-of-Details

- *Real-time neural woven fabric rendering [Chen et al. 2024]*

GT	Ours	NeuMIP	Single-scale, 4 SPP
1177.6 ms	17.2 ms	9.6 ms	18.4 ms

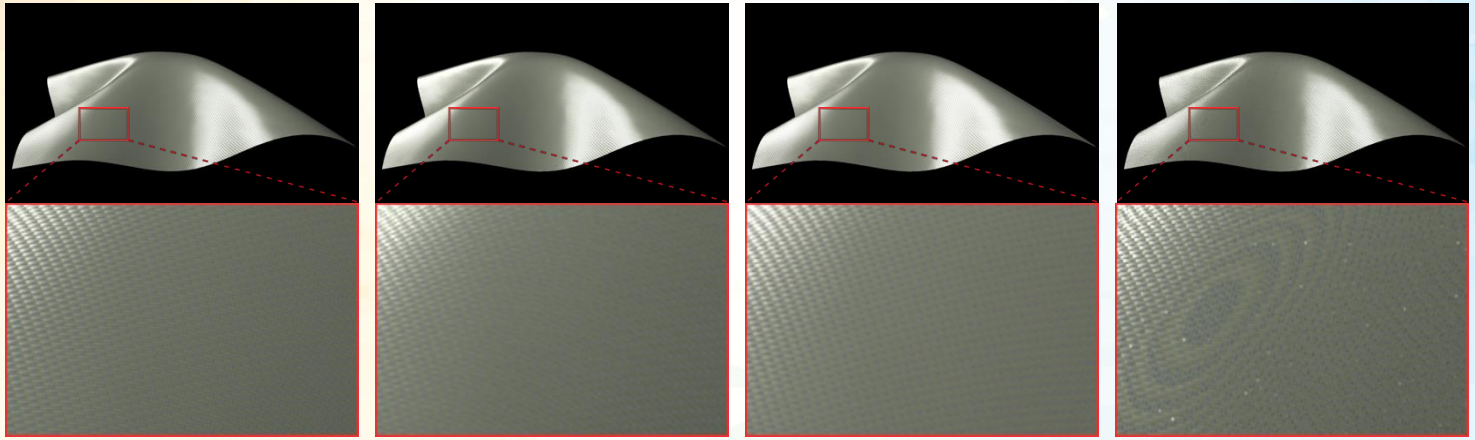

Limitations

- Flyaways, irregularity and non-repeating features.

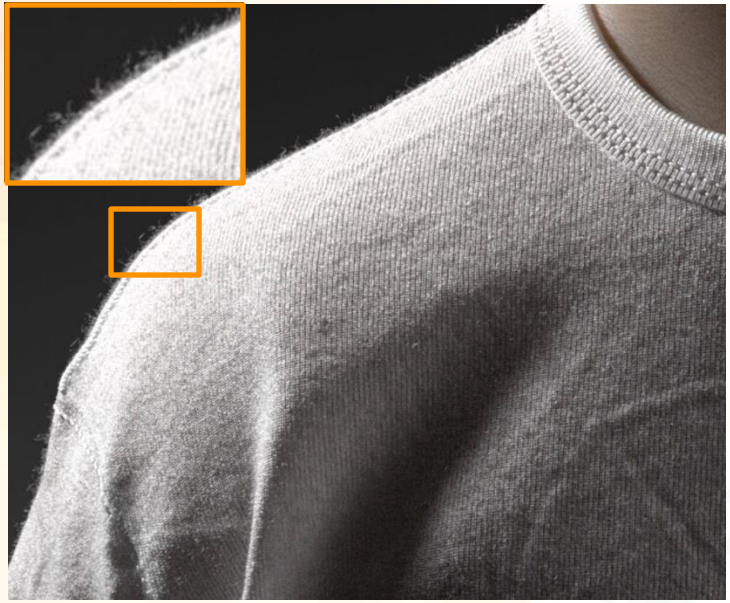

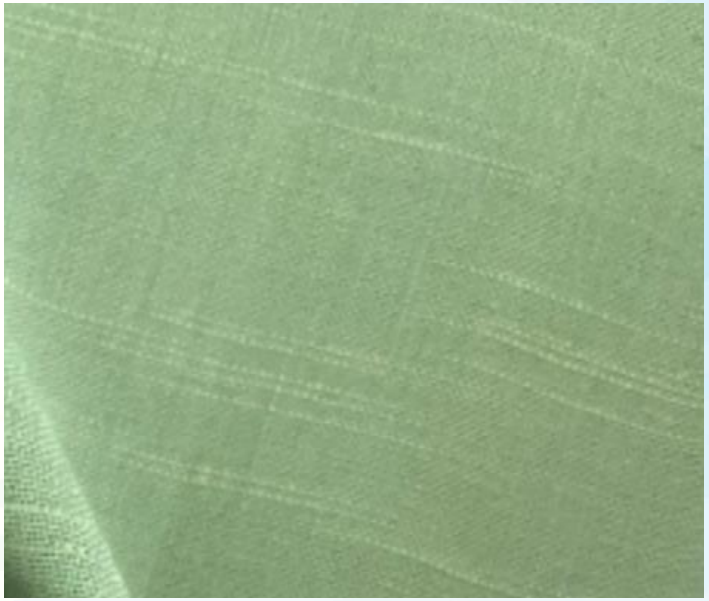

Limitations

- Precomputation-based approach [Zhu et al. 2024]
  - No interactive editing.
- Neural-based approach [Chen et al. 2024]
  - No importance sampling. Limited by network representation ability.

Overview

- State-of-the-art surface-based cloth model
  - Realistic appearance modeling
  - Level of detail
    - Precomputation-based
    - Neural-based
- **Cloth appearance reconstruction**
  - Differentiable rendering
  - Generative AI
  - Data-driven reconstruction

Cloth appearance reconstruction

- Hot topic in many fields that involve full-body digital avatar such as AR/VR.
- The ultimate goal:
  - *Rapid and automatic appearance reconstruction from sparse in-the-wild sensor data that can be rendered under novel geometric deformation, with a novel viewport and under novel lighting conditions, while support appearance editing, all in real-time framerate.*

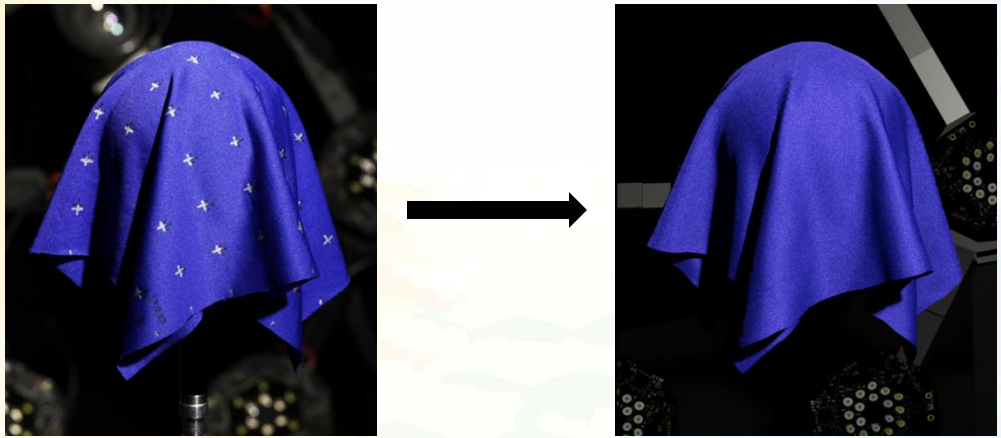

Photo

Render w/ [Zhu et al. 2023]

Differentiable rendering

- Auto material parameters estimation from a photograph using gradient decent.
- Most effective with scene parameters known
  - [Jin et al. 2022] Simple rig to measure light, camera and geometry

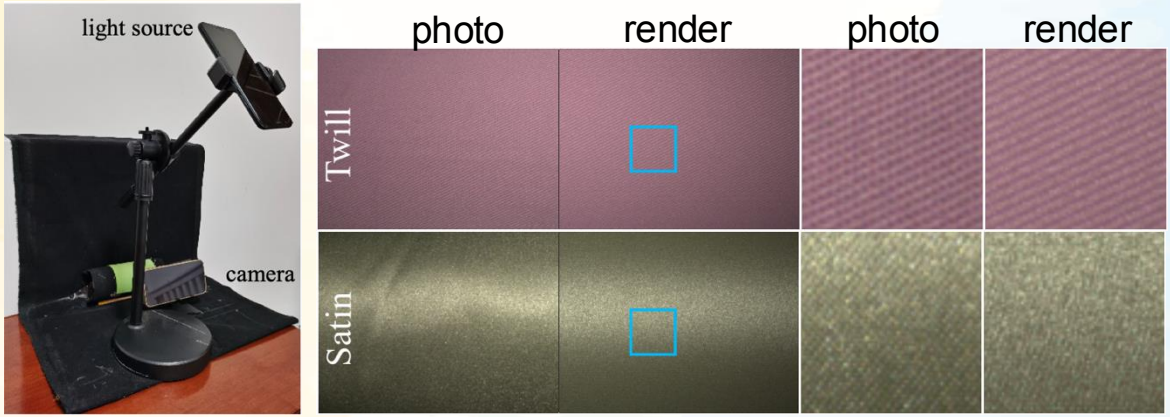

[Jin et al. 2022]

Differentiable rendering

- Subject to local minimums
  - [Jin et al. 2022] Two-stage estimation
    - Initialization w/ an MLP
    - Optimization w/ differentiable rendering

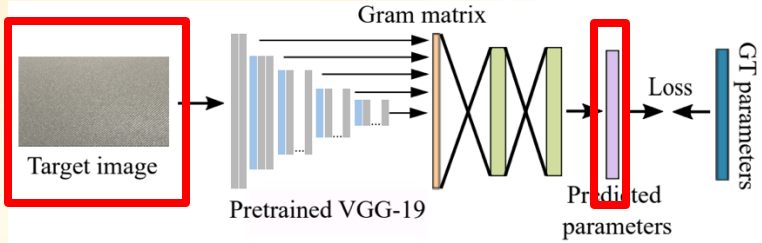

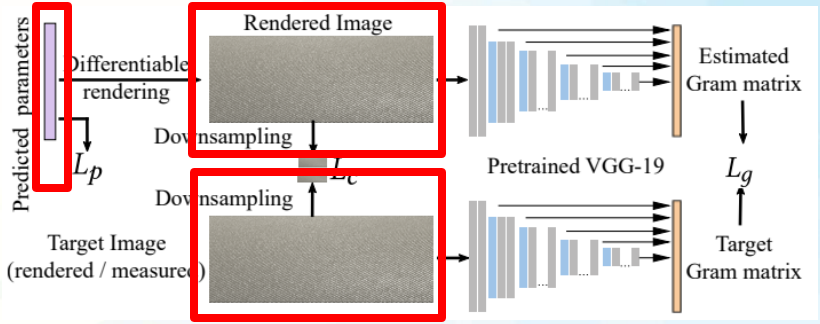

[Jin et al. 2022]

Generative AI

- Diffusion excels at image-to-image translation
- Material reconstruction as intrinsic image decomposition

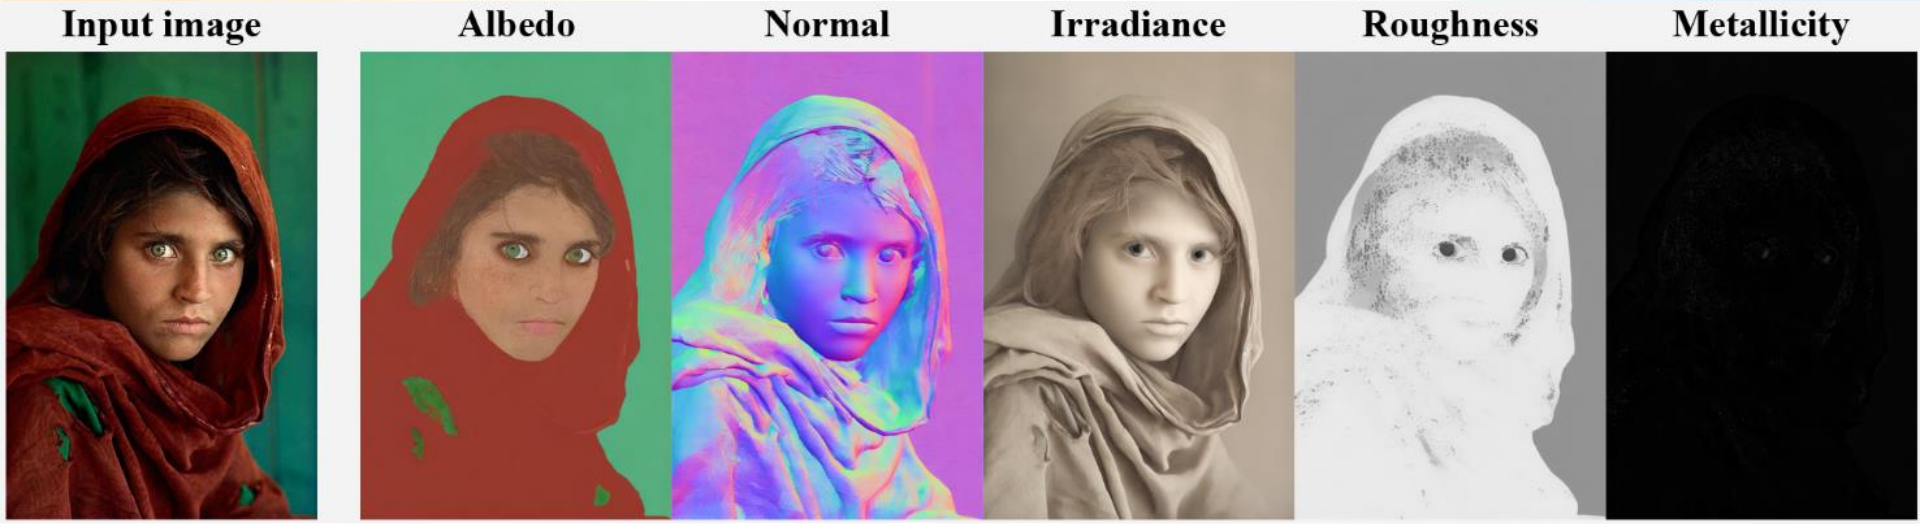

RGB↔X [Zheng et al. 2024]

Generative AI

- RGB->X: reconstruction
- X->RGB: rendering

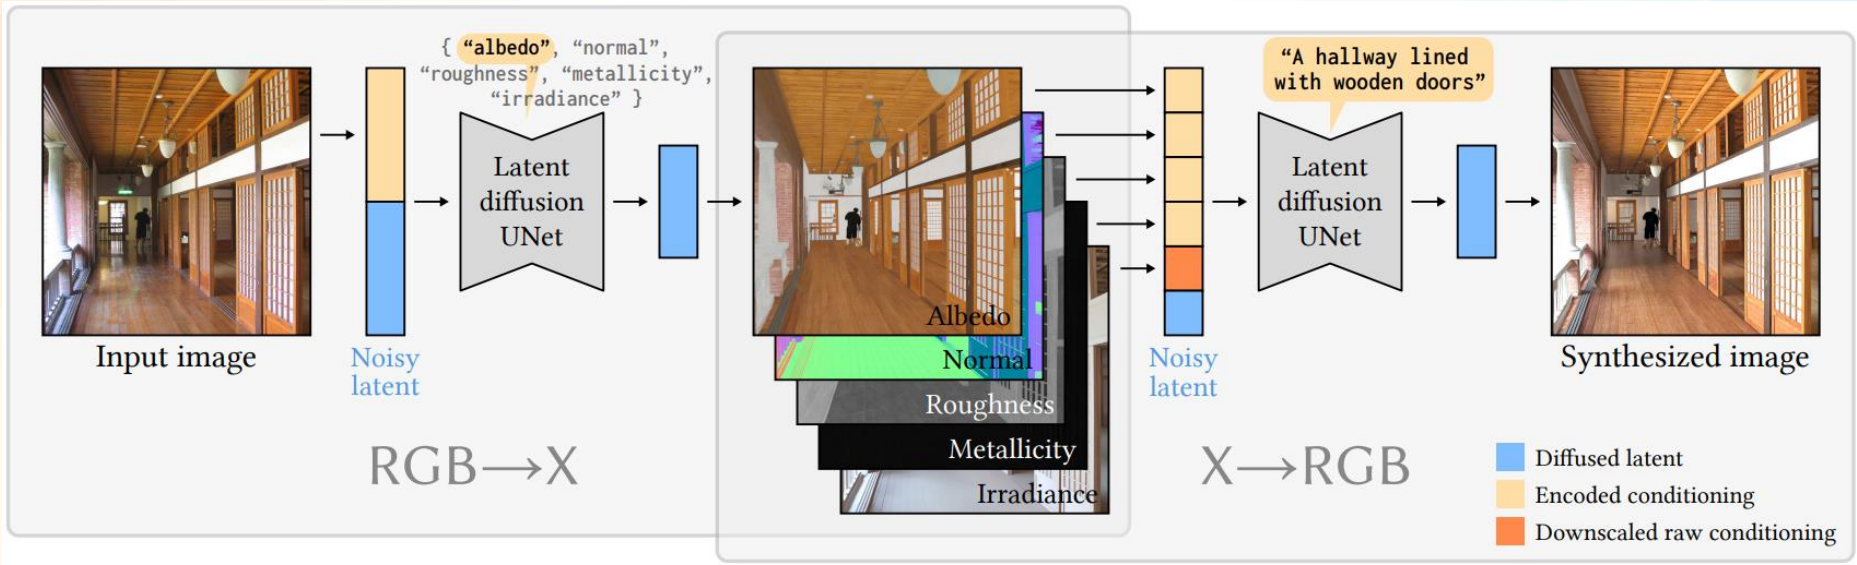

RGB↔X [Zheng et al. 2024]

Generative AI

- Conditional editing via a text prompt
- Limitations: primarily 2D. No novel pose or novel view.

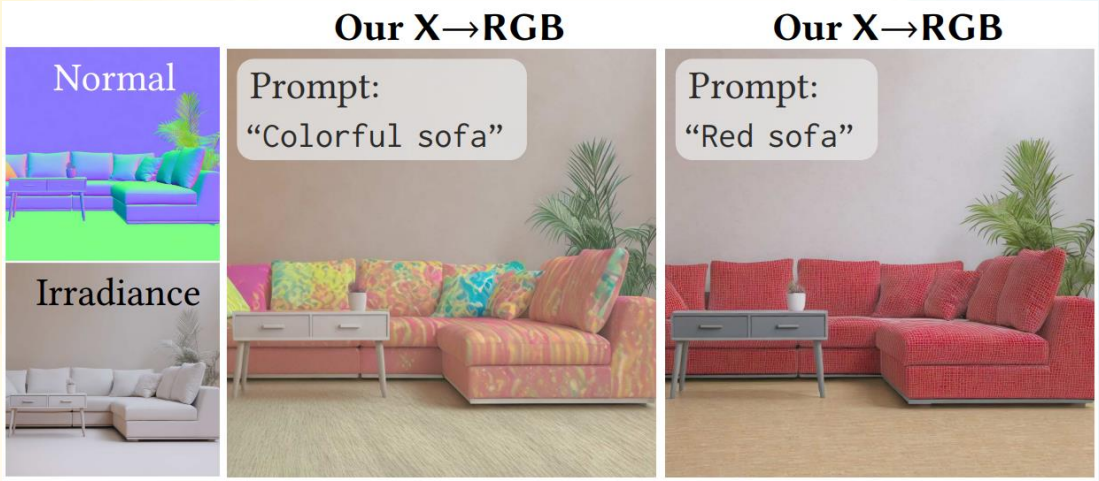

RGB↔X [Zheng et al. 2024]

Data-driven reconstruction

- 3D Gaussian splatting.
- Data-driven. Novel-view synthesis. Good at volumetric appearance for fabrics.

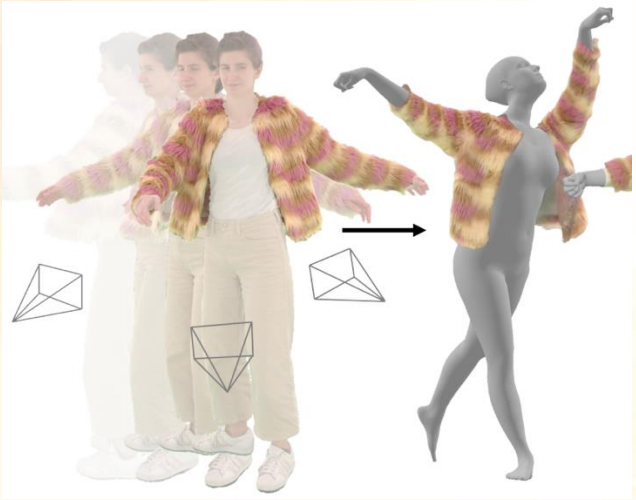

Gaussian Garments [Rong et al. 2024]

PhysAvatar [Zheng et al. 2024]

Data-driven reconstruction

- Generalize to novel pose
  - Learning **pose-dependent *3DGS parameter offsets*** using local geometric info from multi-view **video** sequences.

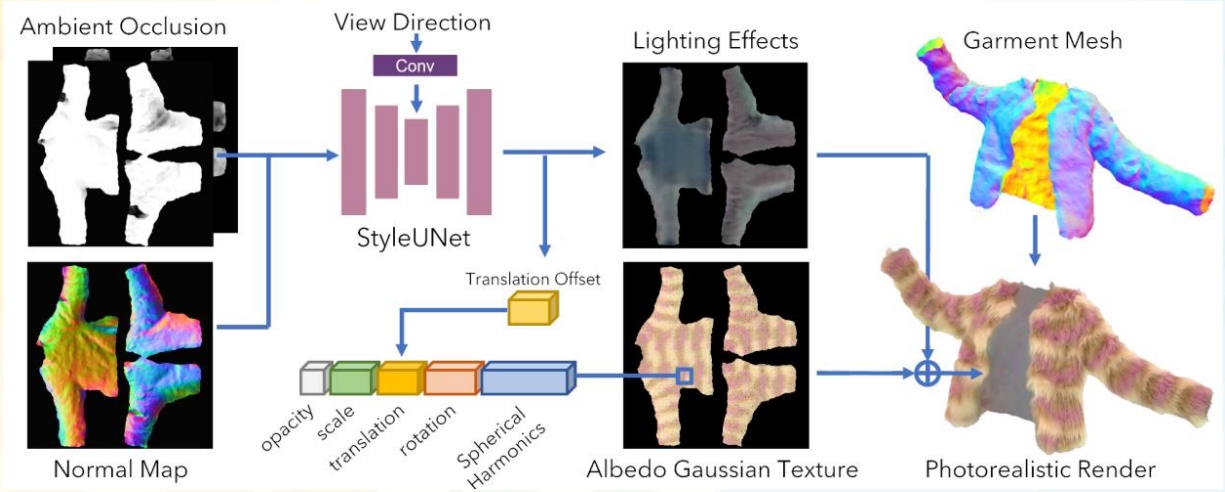

Gaussian Garments [Rong et al. 2024]

Data-driven reconstruction

- Generalize to novel pose and novel lighting.
  - Learning **pose-dependent *intrinsic maps*** using local geometric info from multi-view **video** sequences.

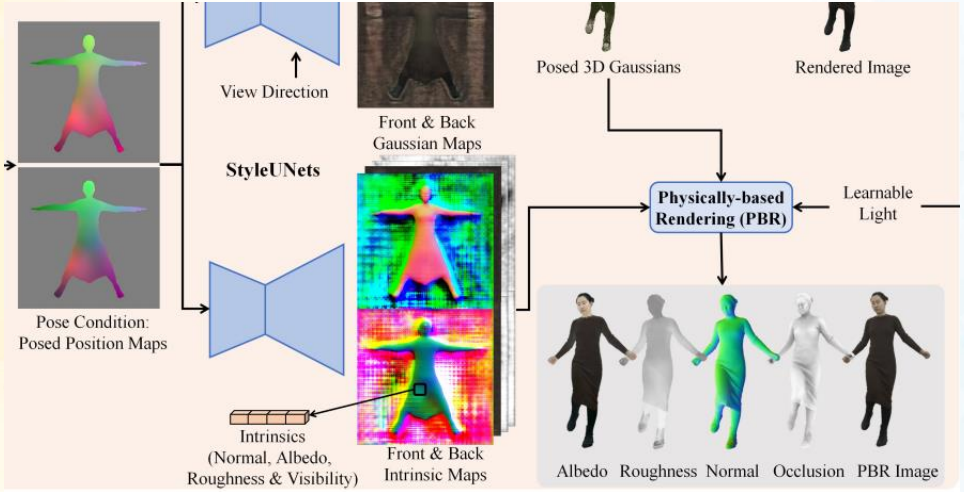

[Li et al. 2024]

Data-driven reconstruction

- Limitations:
  - Lots of data. Slow reconstruction.
  - Not easy for editing.
  - Visual quality is not thoroughly analyzed.

Overview

- State-of-the-art surface-based cloth model
  - Realistic appearance modeling
  - Level of detail
    - Precomputation-based
    - Neural-based
- Cloth appearance reconstruction
  - Differentiable rendering
  - Generative AI
  - Data-driven reconstruction