thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

Accurate Appearance Preserving Prefiltering for Rendering Displacement-Mapped Surfaces

Lifan Wu[1] Shuang Zhao[2]
Ling-Qi Yan[3] Ravi Ramamoorthi[1]

[1]University of California, San Diego
[2]University of California, Irvine
[3]University of California, Santa Barbara

Realistic Appearance Models

Simple surface

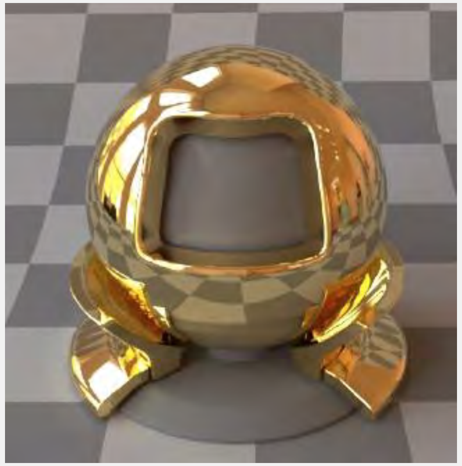

Complex surface

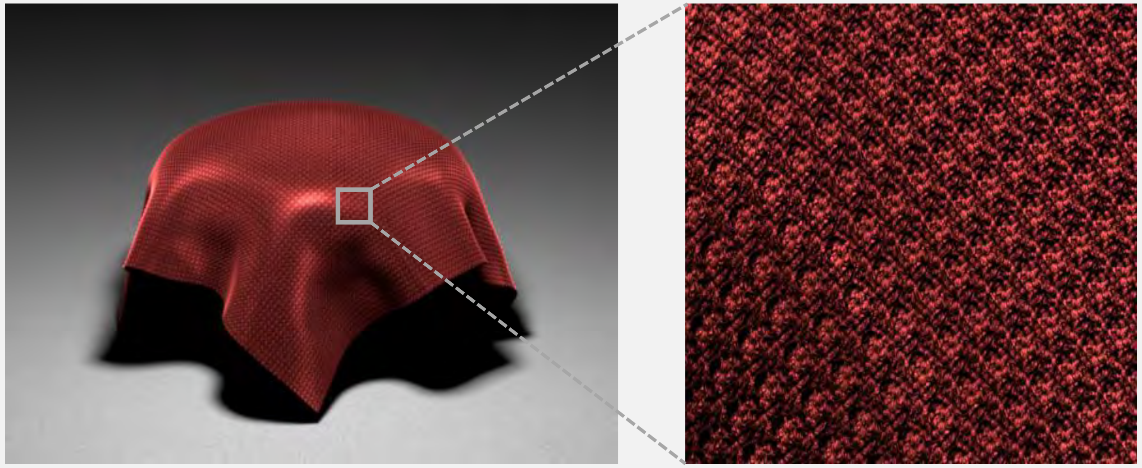

Image courtesy of Mitsuba [Jakob 2010]

Appearance Models with Rich Details


[Jakob et al. 2010]


[Zhao et al. 2011]


[Heitz et al. 2015]


[Khungurn et al. 2015]

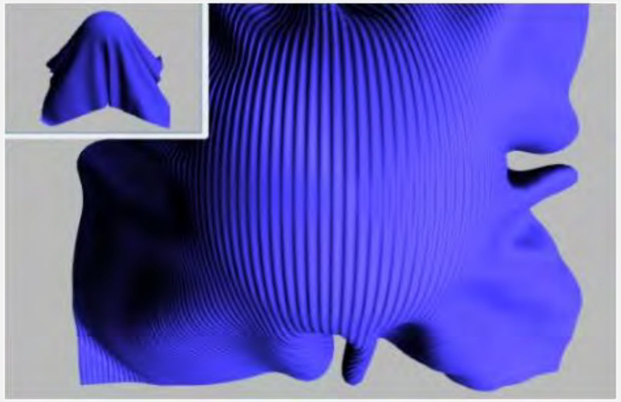
[Han et al. 2007]

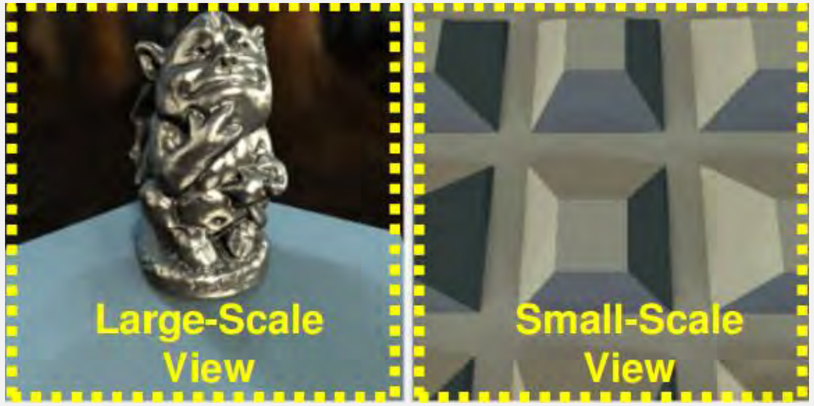

[Wu et al. 2011]


[Yan et al. 2014, 2016]

Modeling Details

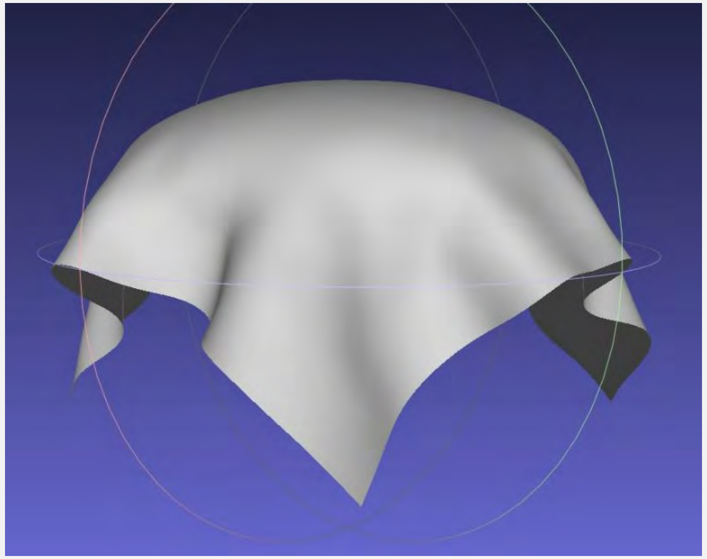

Base shape

+

Micro-geometry

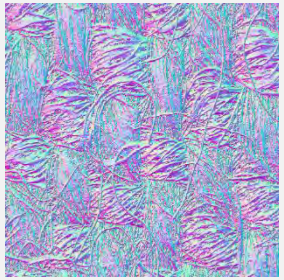

Normal map (2D)

Disp. map (2D)

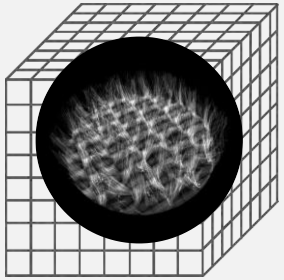

Volume (3D)

Micro-scattering

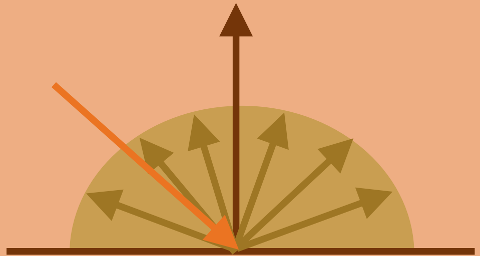

BRDF (Hemispherical)

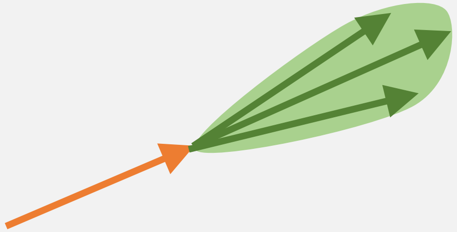

Phase function (Spherical)

Problems

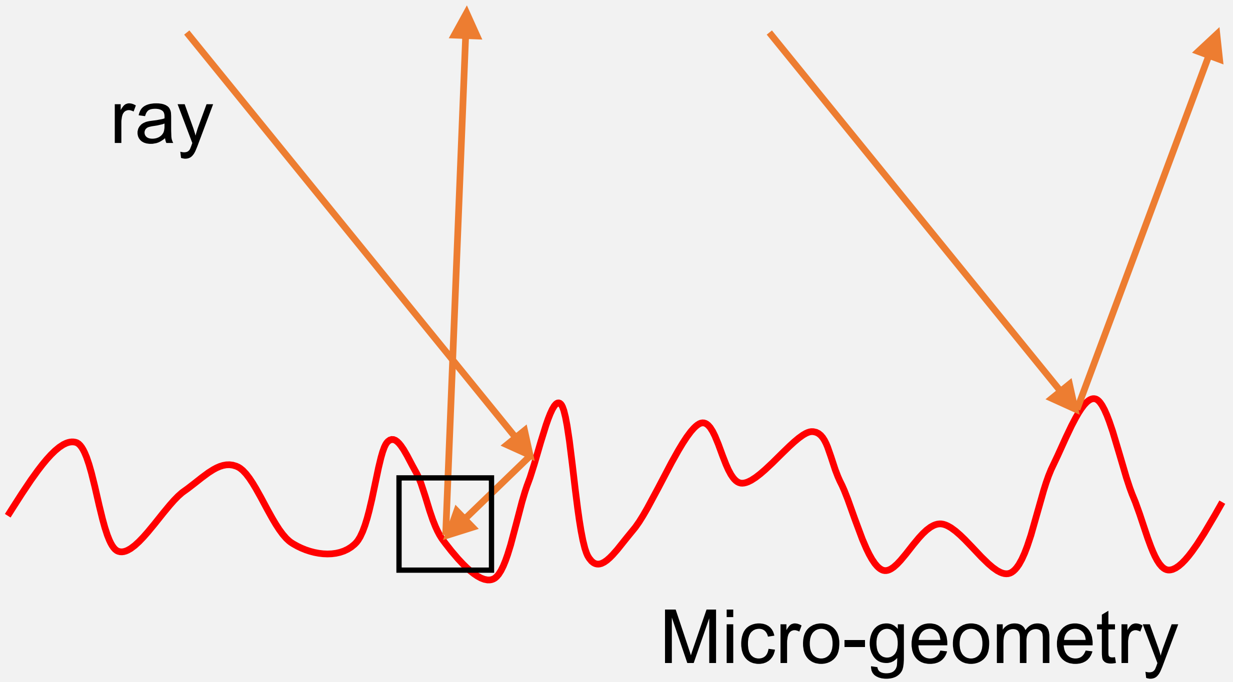

$$F = \int_{A_n} \cdots \int_{A_1} f(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)\, \mathrm{d}A(\boldsymbol{x}_1) \ldots \mathrm{d}A(\boldsymbol{x}_n)$$

Complex light-surface interaction

Difficult to compute and analyze

Motivation

- Camera zooming out ➔ less details are visible ➔ use coarser models [Zhao et al. 2016]

Our Goal

- Prefilter high-resolution displacement maps + BRDFs

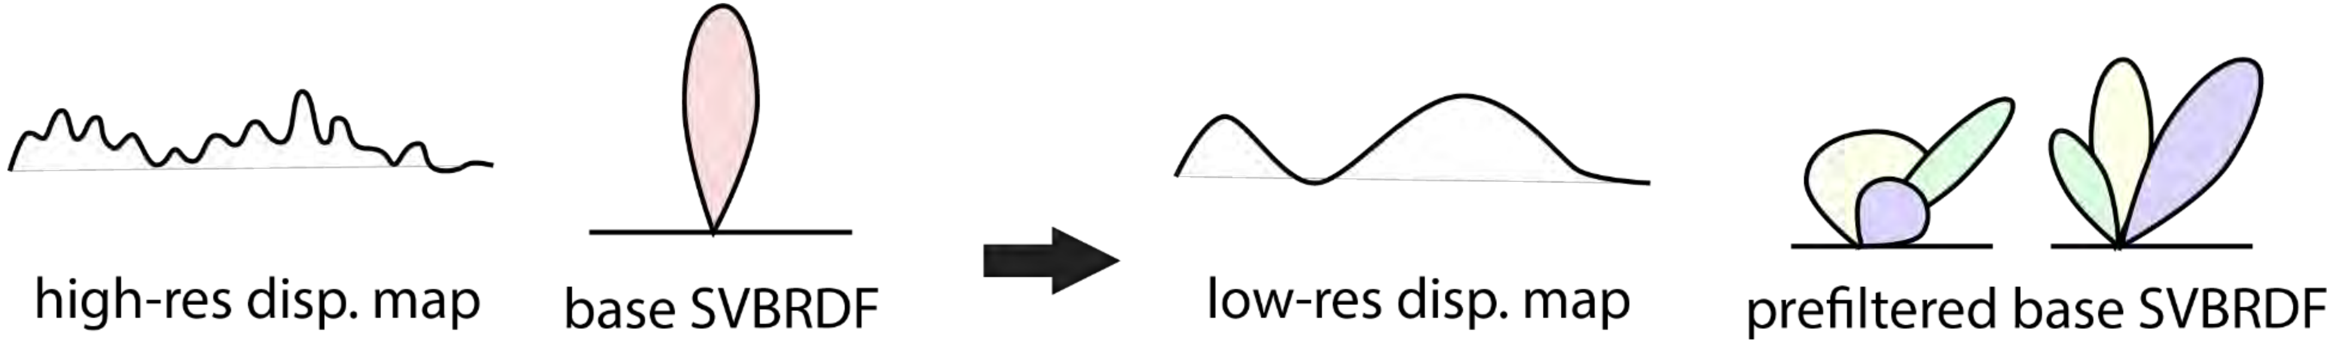

- Preserve appearance

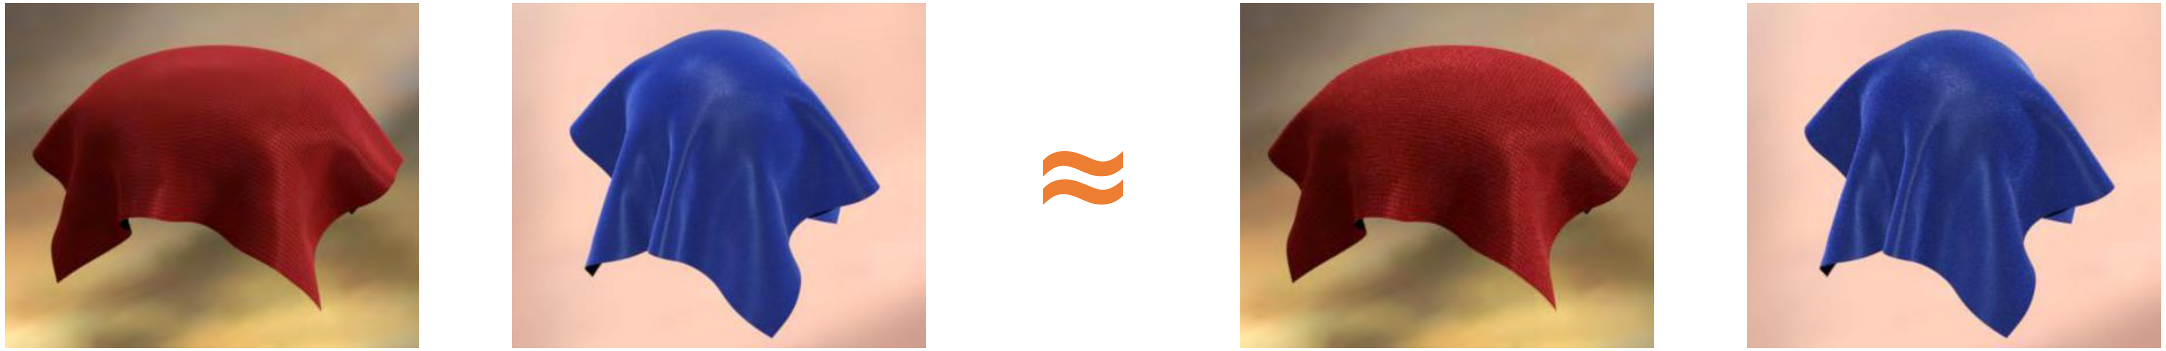

Original

Prefiltered

Prefiltering

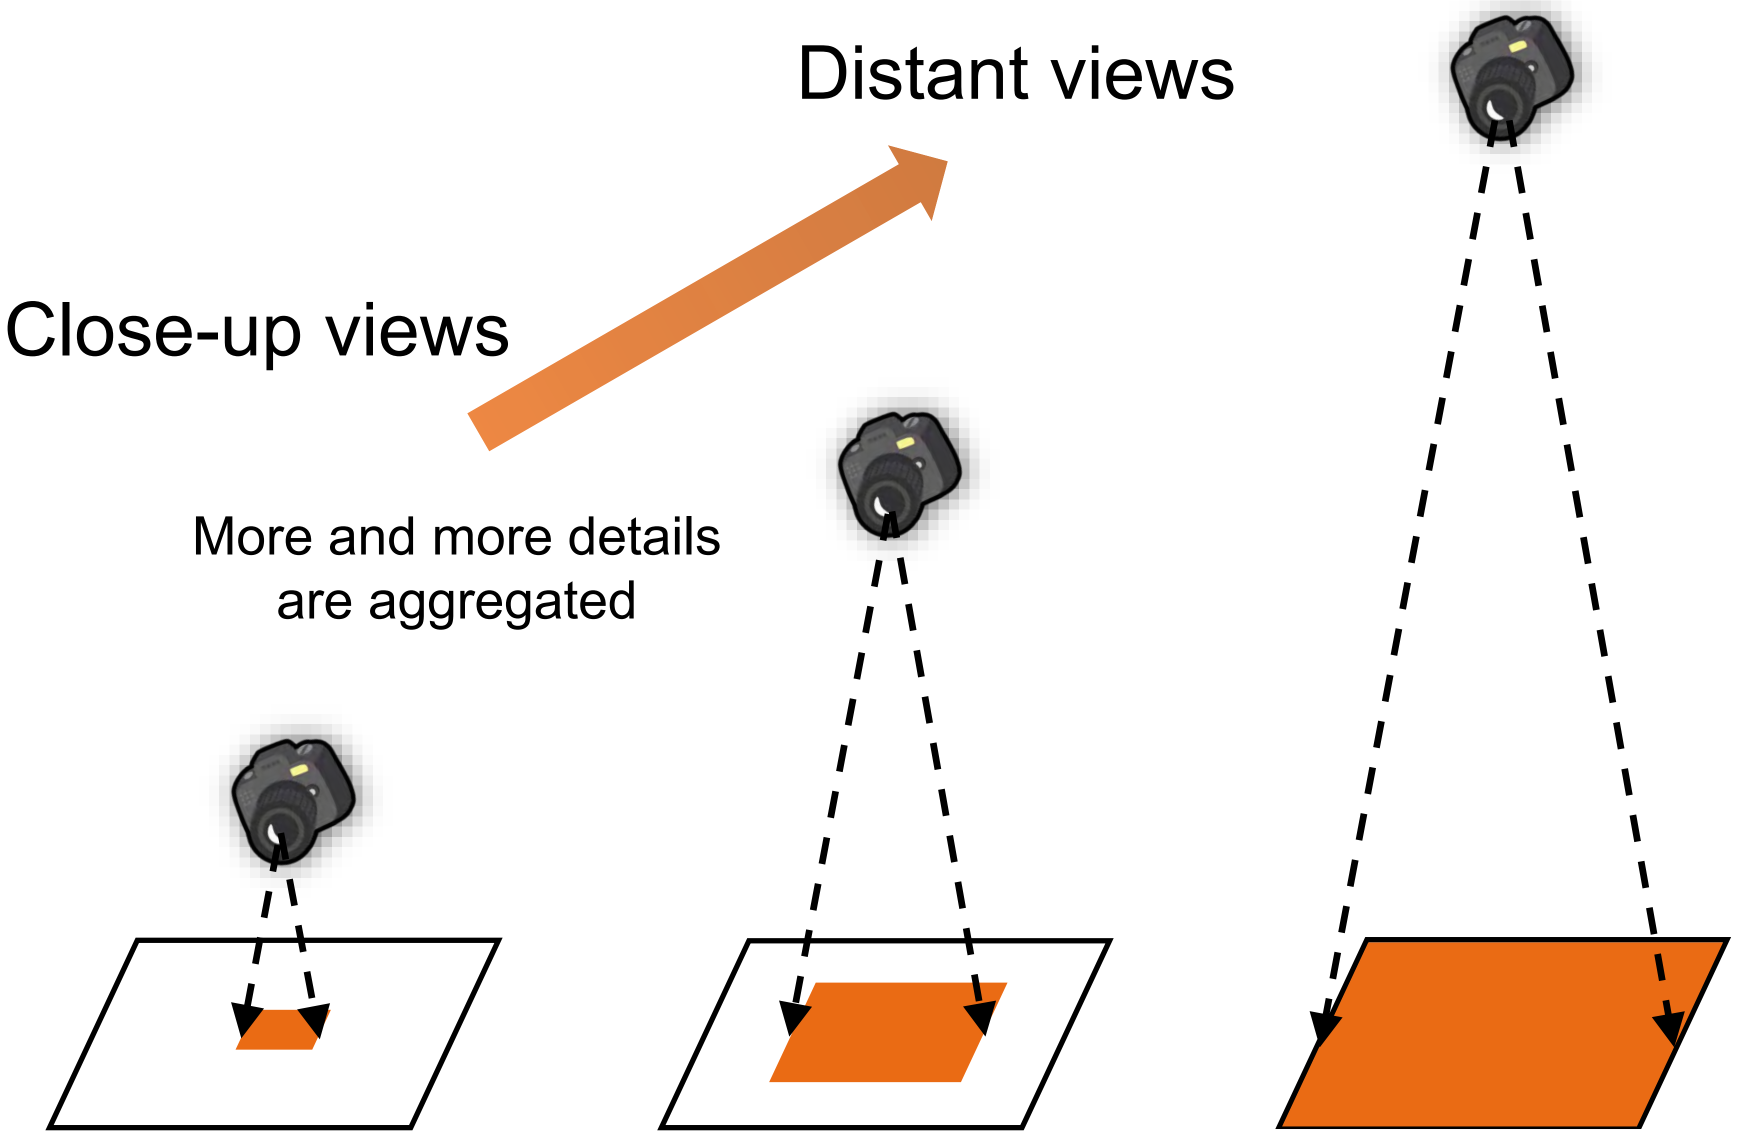

Coarser models representing aggregate micro-appearance

Match the desired appearance

Precomputed before rendering

Benefits

- Anti-aliasing, storage reduction

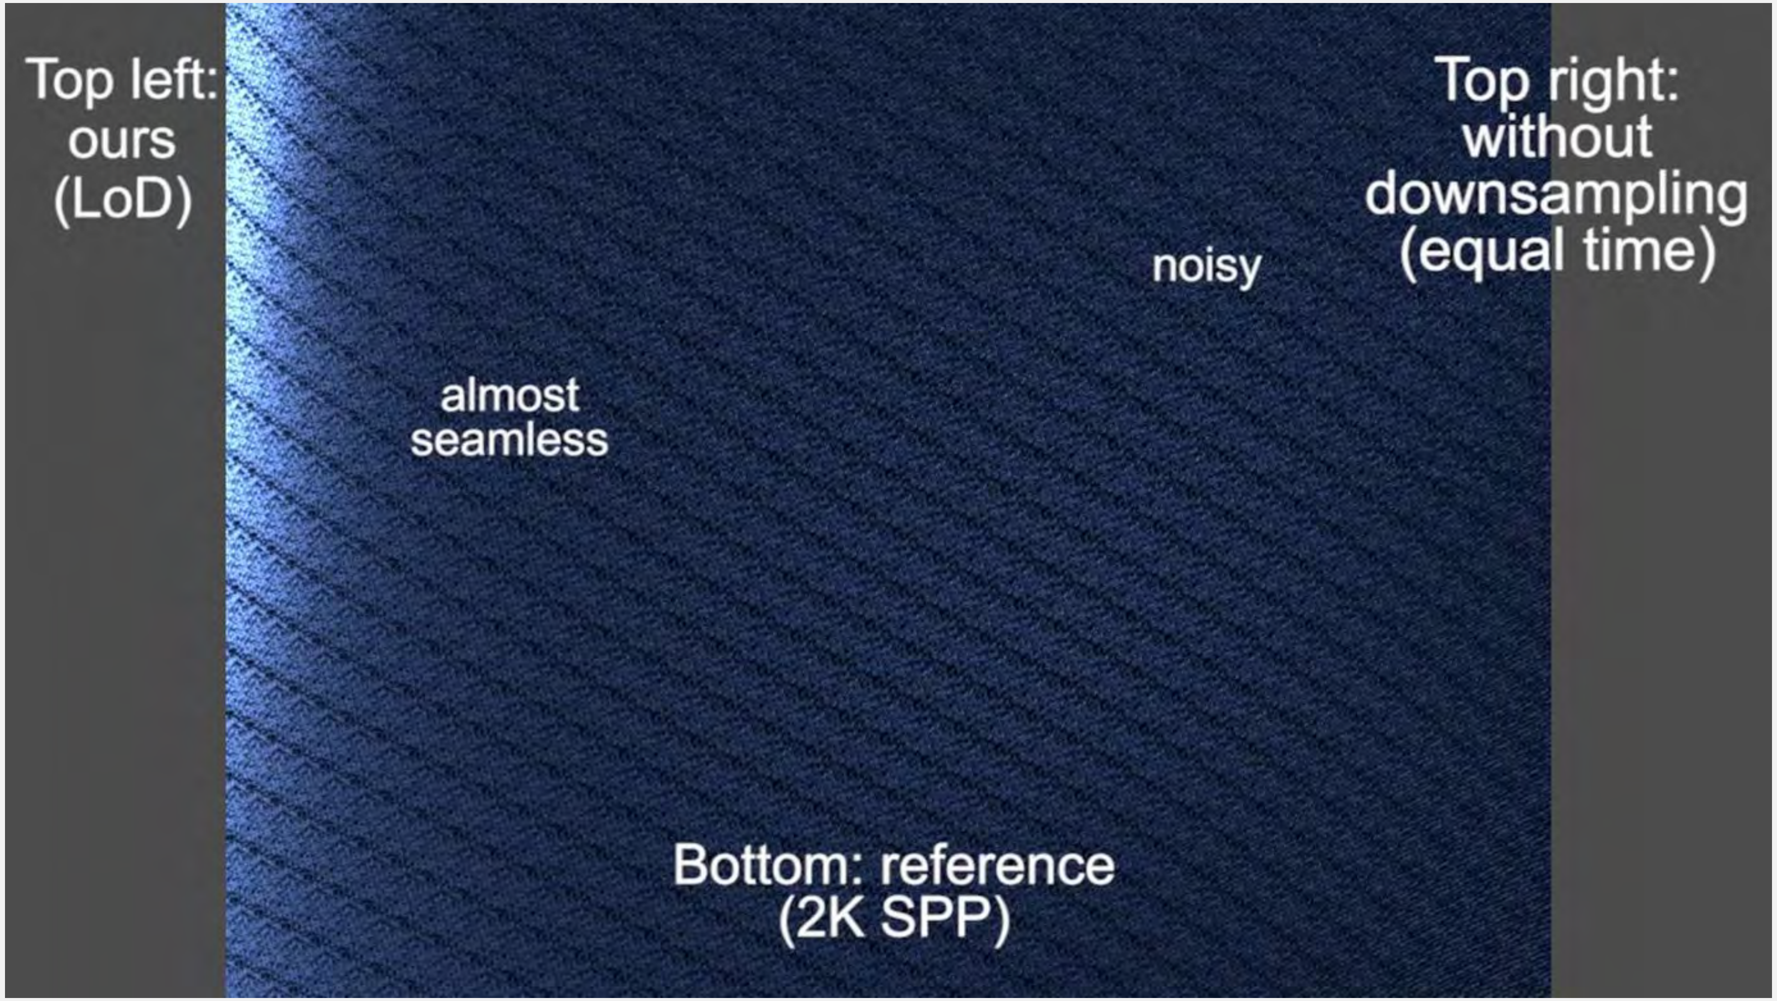

Challenges

- Difficult to accurately capture changes of illumination effects

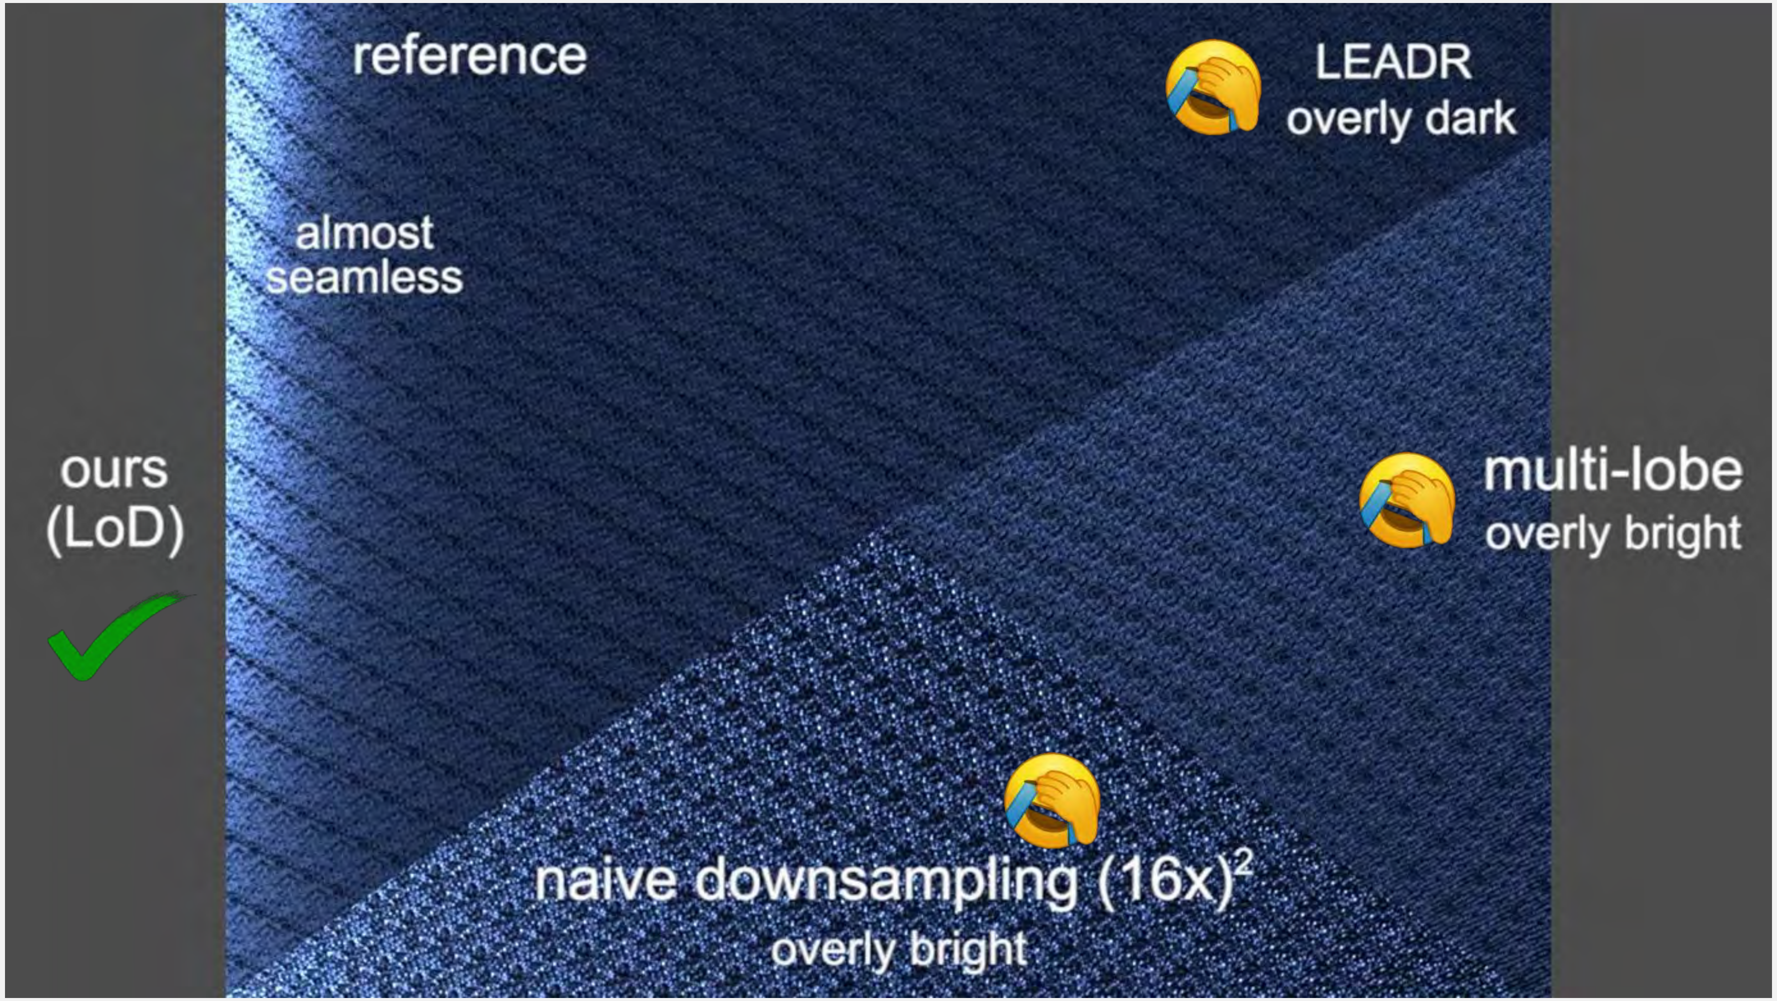

Our Contributions

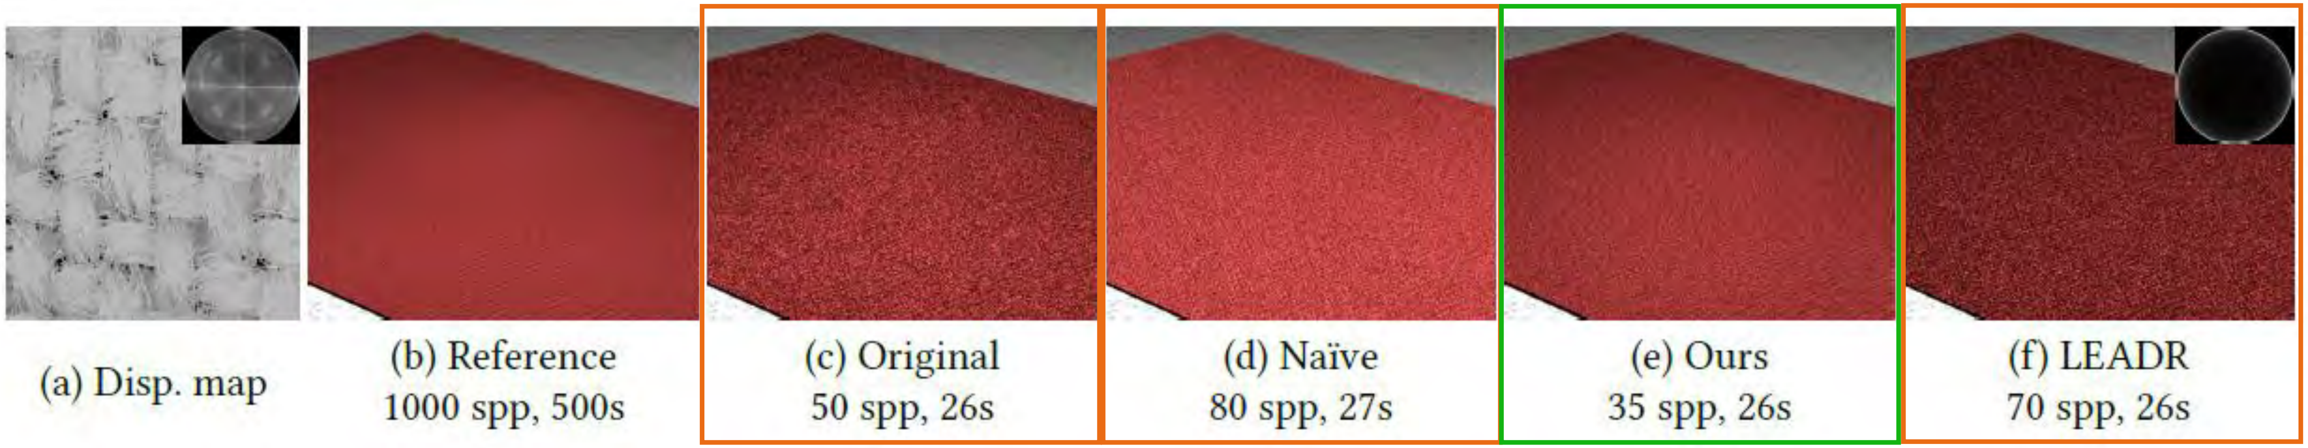

✓ Anti-aliased

✓ Accurate

✓ General surface

Background

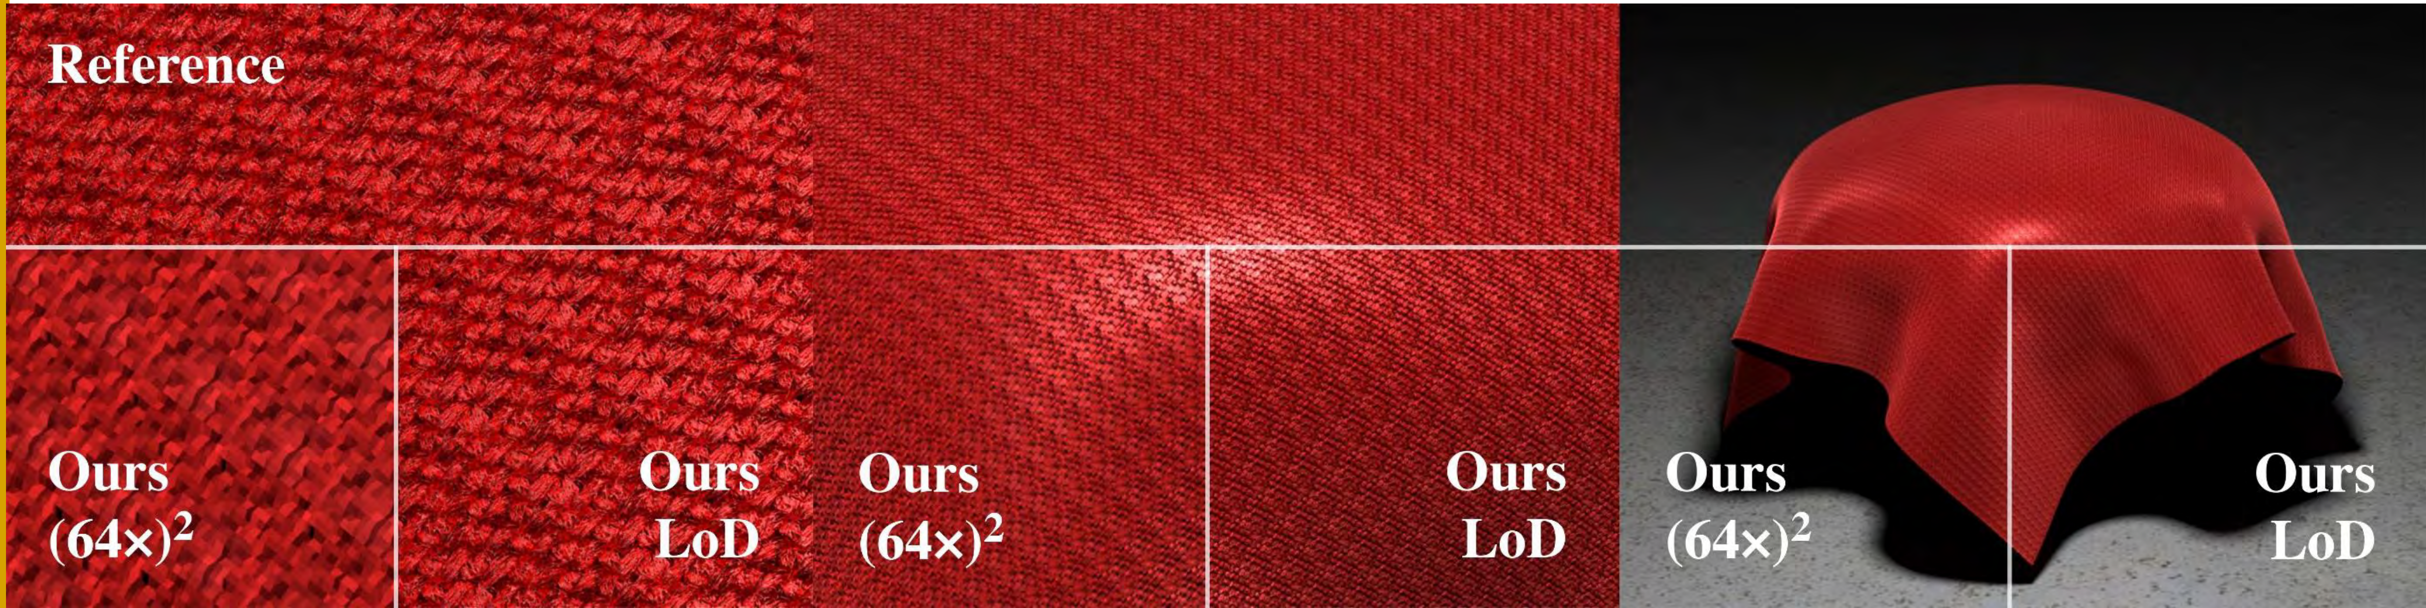
Reference
Ours
(64×)²
Ours
LoD
Ours
(64×)²
Ours
LoD
Ours
(64×)²
Ours
LoD

2D Displacement Maps

- Describe surface details (micro-geometry)
- Need expensive super-sampling

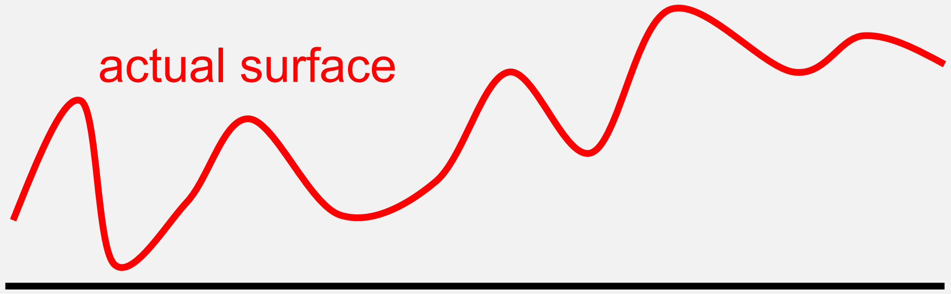

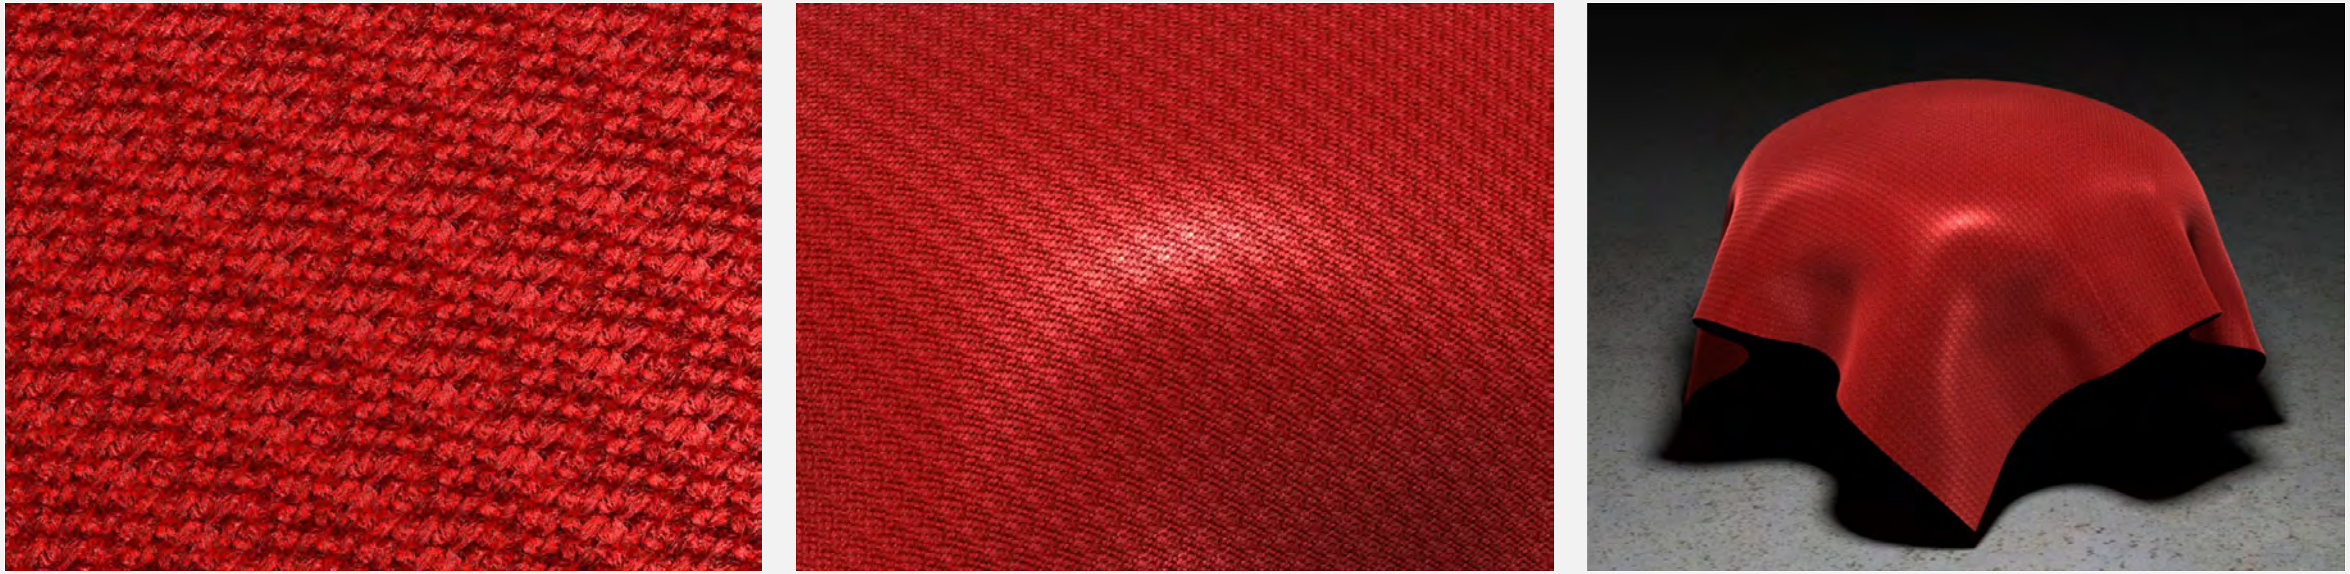

Close-up views → Distant views

Displacement mapping

- Surface patch

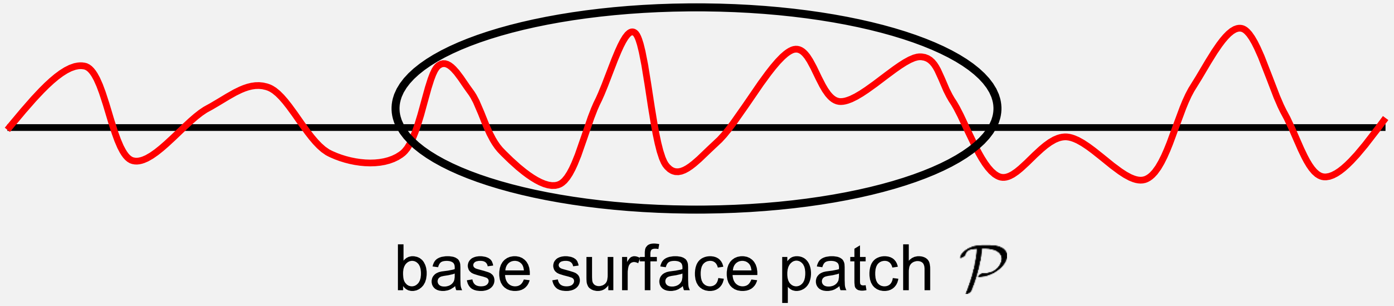

- Micro-geometry

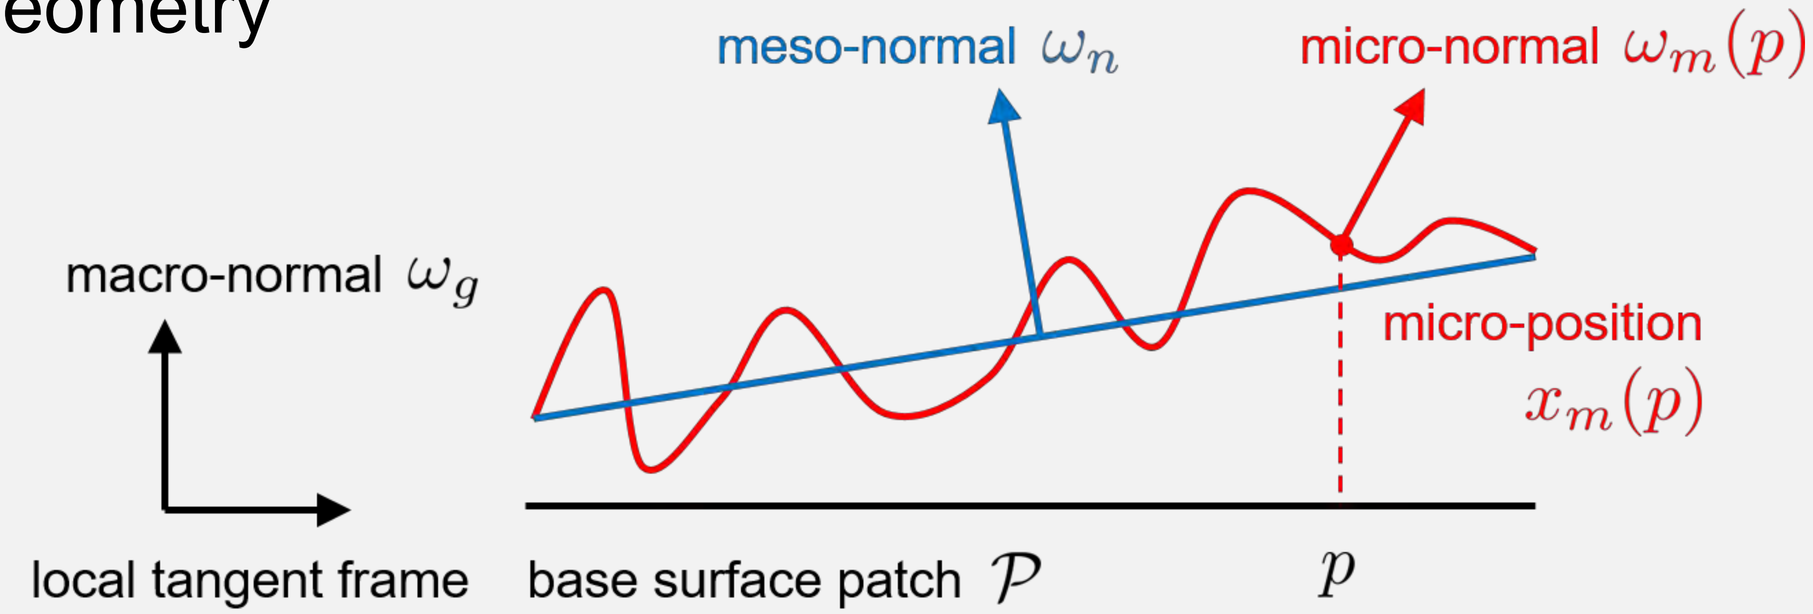

Prefiltering

- Jointly handle changes of illumination effects
- It is challenging due to non-linearity

Normal variation

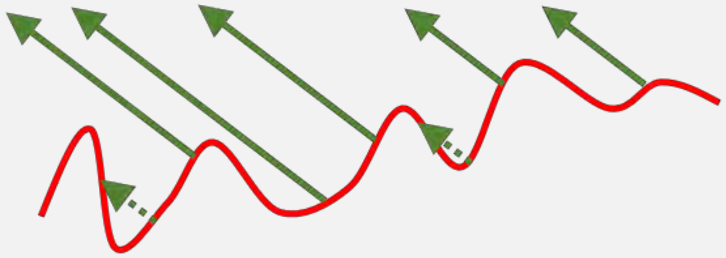

Shadowing-masking

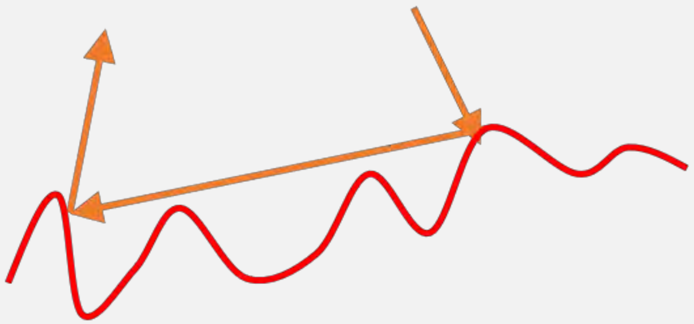

Interreflections

Previous Work

- Handle parts of illumination effects

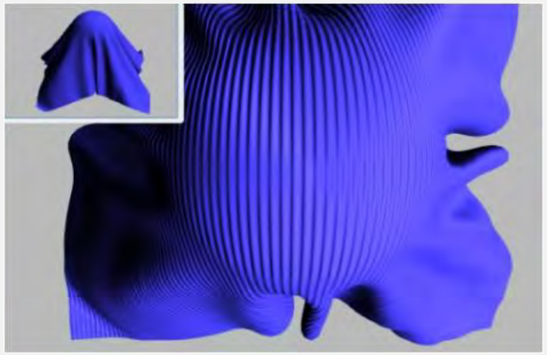

[Han et al. 2007]

Normal variation

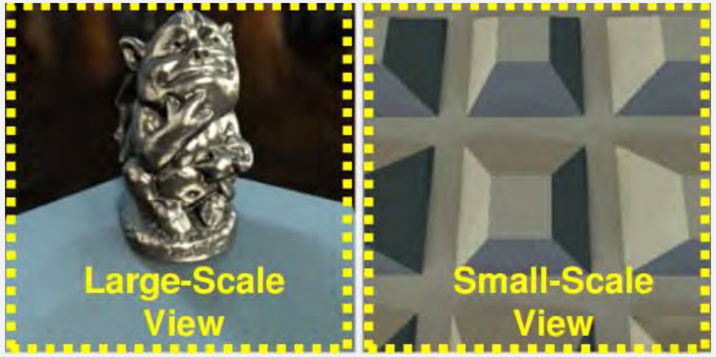

[Wu et al. 2011]

Normal variation + Shadowing-masking

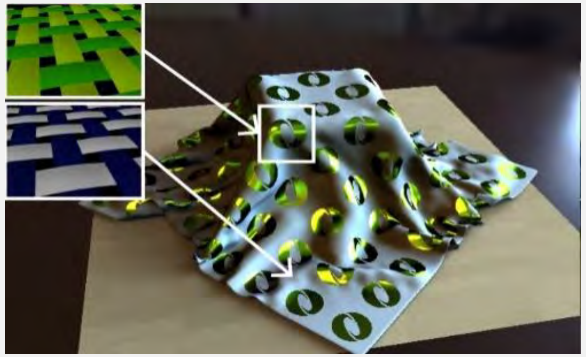

[Iwasaki et al. 2012]

Normal variation + Shadowing-masking

- Missing

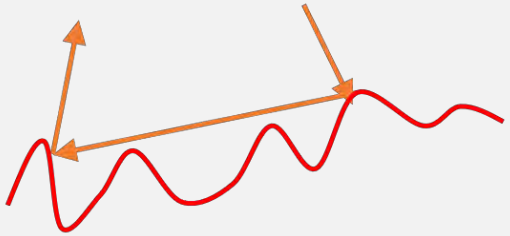

Interreflections

Previous Work

- Assuming certain types of surface (Gaussian/GGX/V-groove)

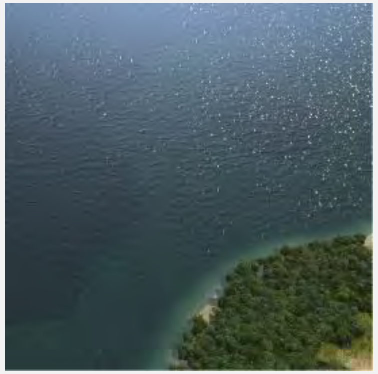

[Olano and Baker 2010]

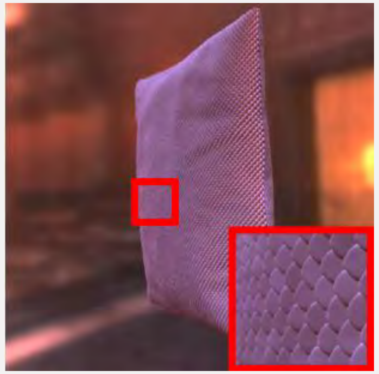

[Dupuy et al. 2013]

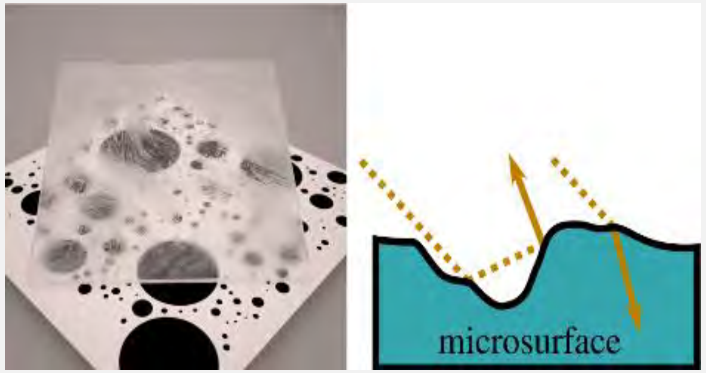

[Heitz et al. 2016]

[Lee et al. 2018]
[Xie and Hanrahan 2018]

- Fail to generalize

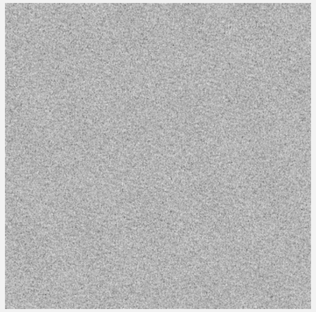

Gaussian surfaces

different

General surfaces

Previous Work

- Iterative inverse rendering (optimization) is expensive

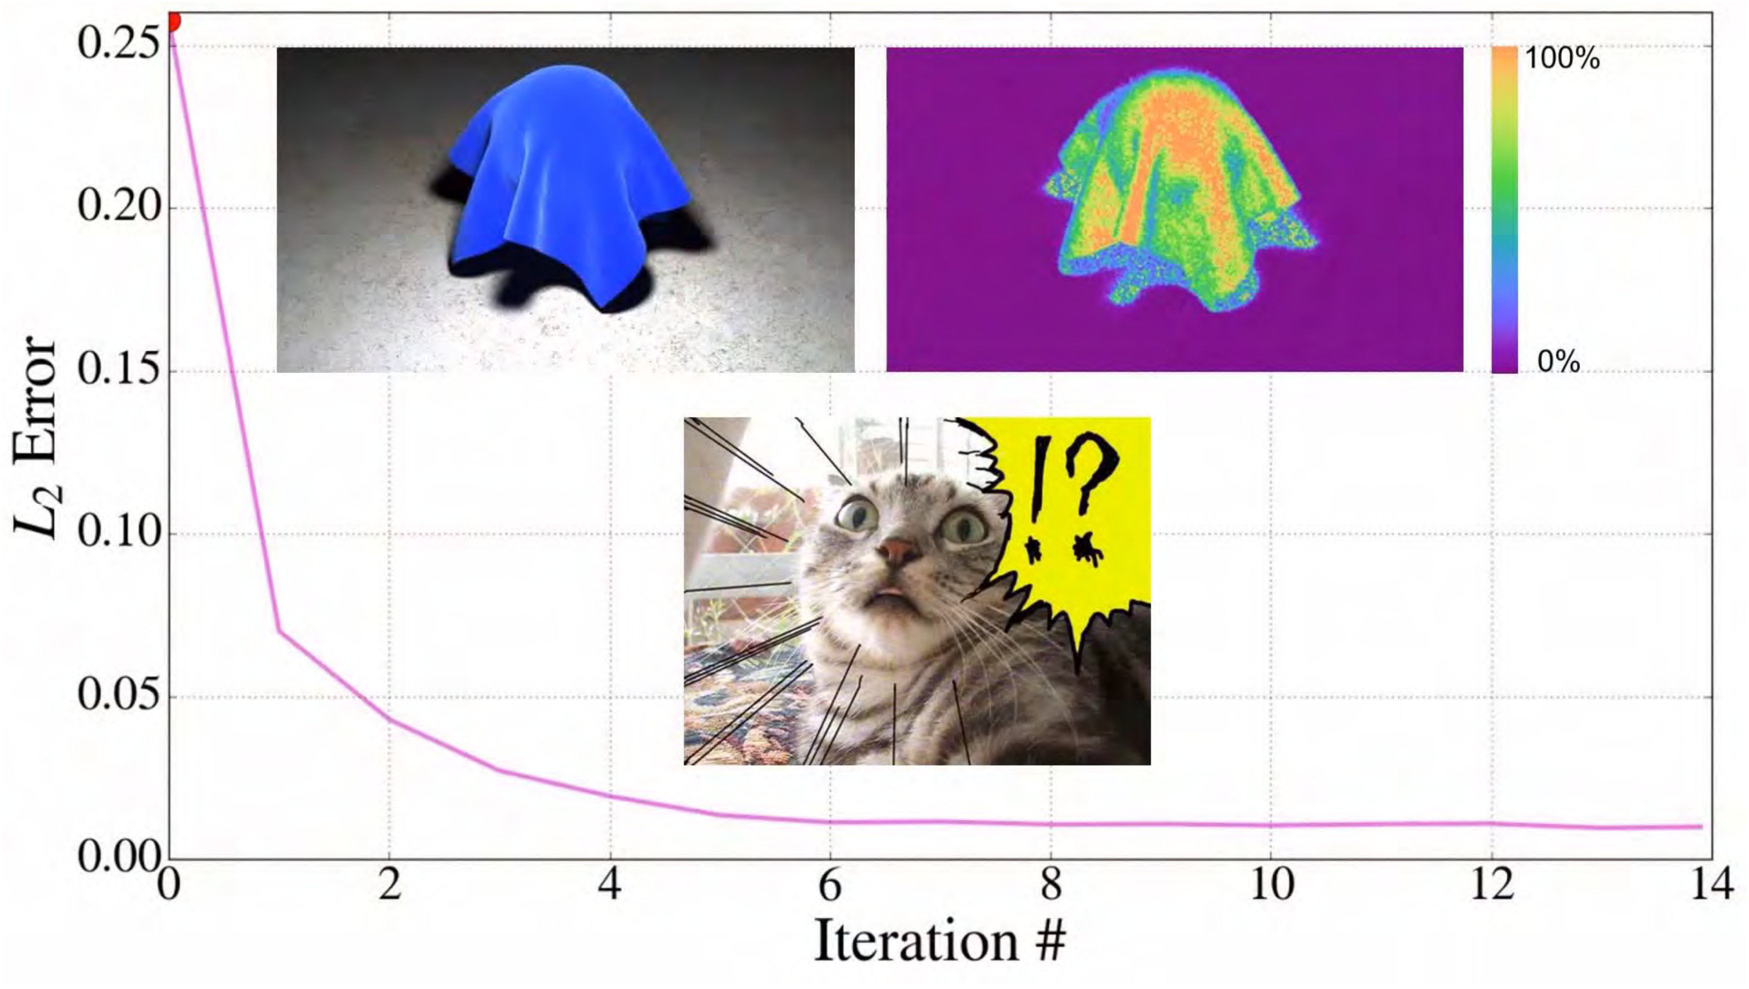

Our Approach vs. Previous Work

Method	Interreflections	General surfaces	Precomputation
Bi-Scale	No	Yes	Fast
Microfacet	Yes	No	Very fast
Inverse optimization	Yes	Yes	Slow
Ours	Yes	Yes	Fast

Effective BRDF

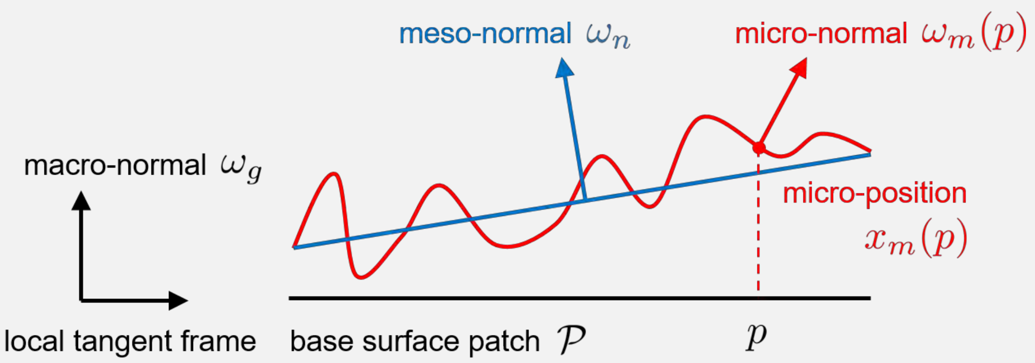

Micro-geometry

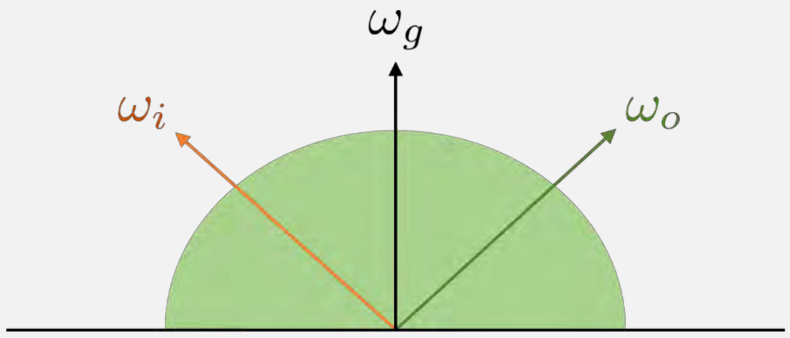

Micro-BRDF

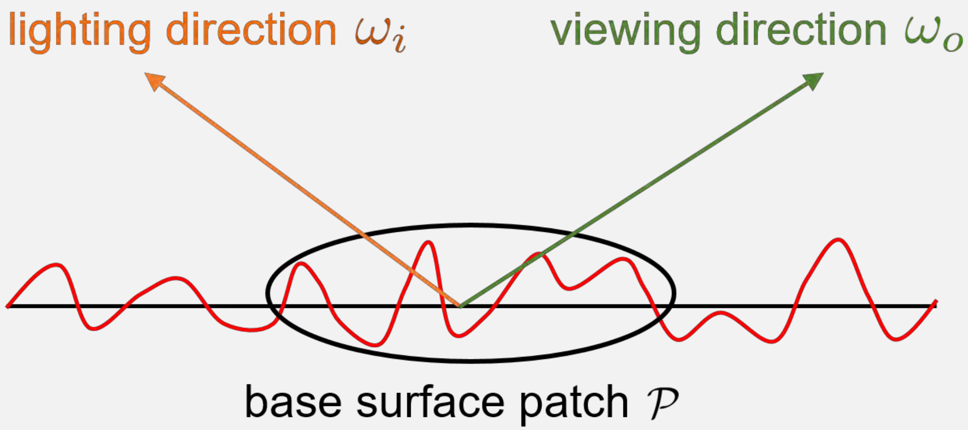

Effective BRDF

Effective BRDF

- Weighted average BRDF over $\mathcal{P}$ [Wu et al. 2011] [Dupuy et al. 2013]

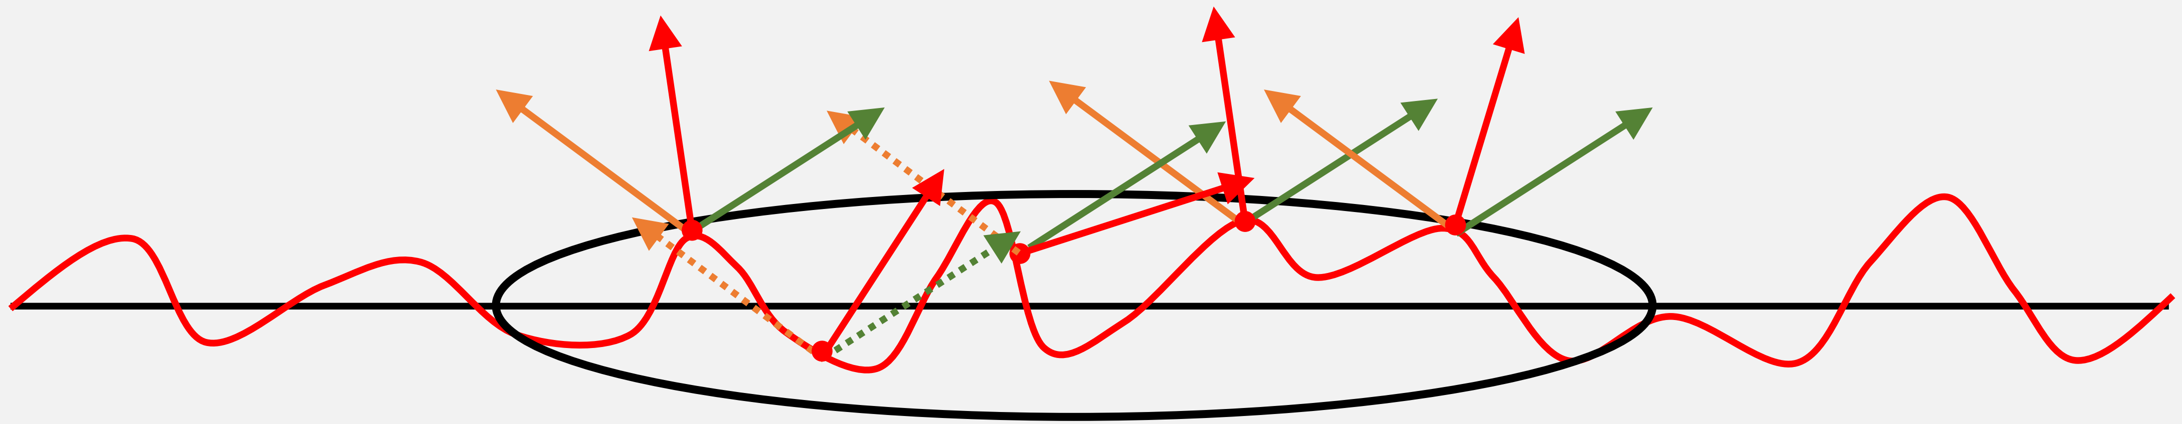

$$f^{\text{eff}}(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o; \mathcal{G}, f) =$$

Our Approach

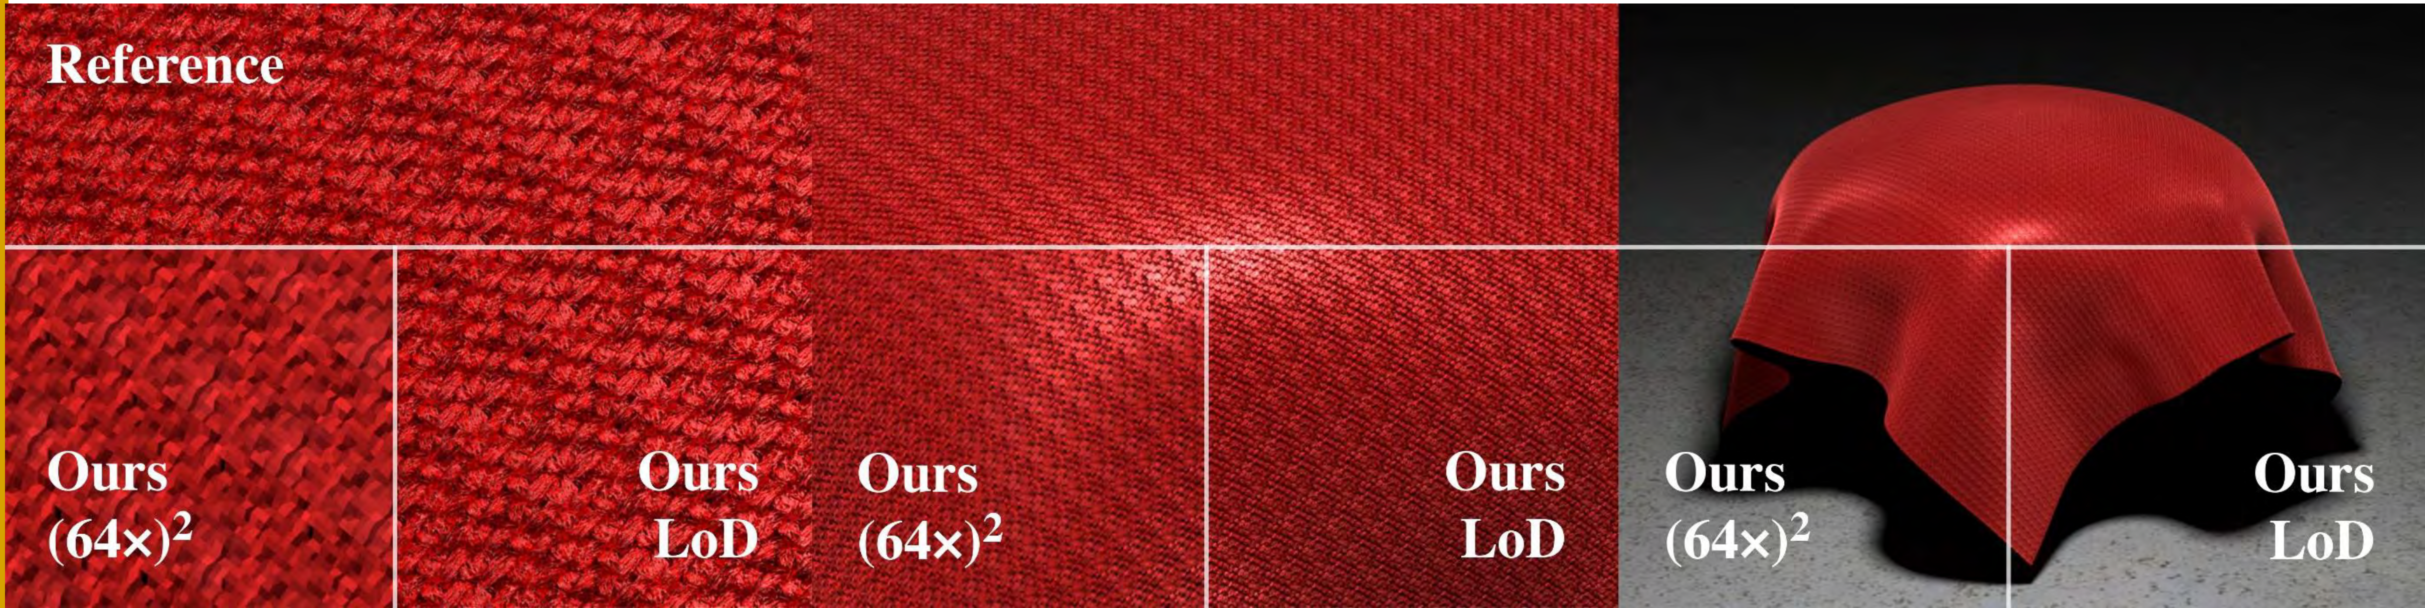

Overview

Joint prefiltering | Appearance matching | LoD rendering

- Joint prefiltering

high-res disp. map

G_{orig}

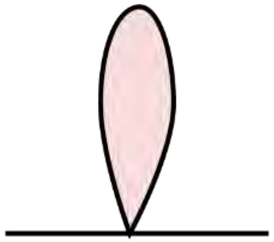

base SVBRDF

f_{orig}

low-res disp. map

G_{low}

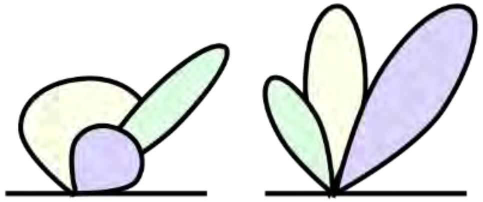

prefiltered base SVBRDF

f_{low}

- Appearance matching

ω_o

ω_i

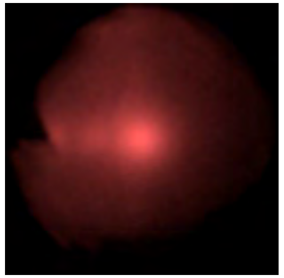

$f^{\mathrm{eff}}(\mathcal{G}_{\mathrm{orig}}, f_{\mathrm{orig}})$

$\approx$

ω_o

ω_i

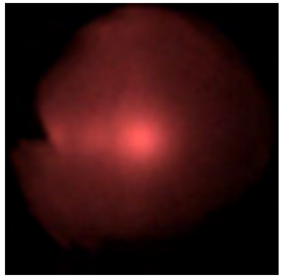

$f^{\mathrm{eff}}(\mathcal{G}_{\mathrm{low}}, f_{\mathrm{low}})$

Before prefiltering

$\approx$

After prefiltering

Effective BRDF with Interreflections

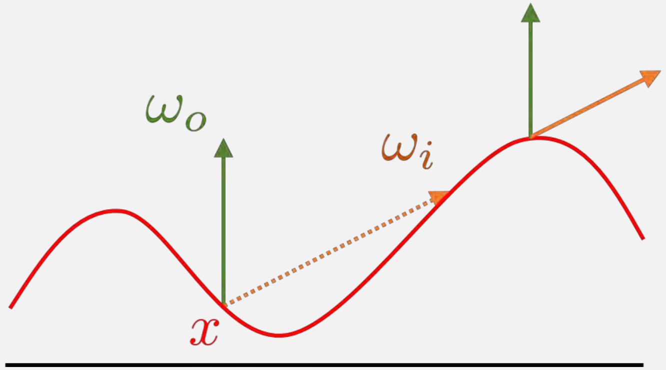

$$f^{\text{eff}}(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o; \mathcal{P}, f)$$

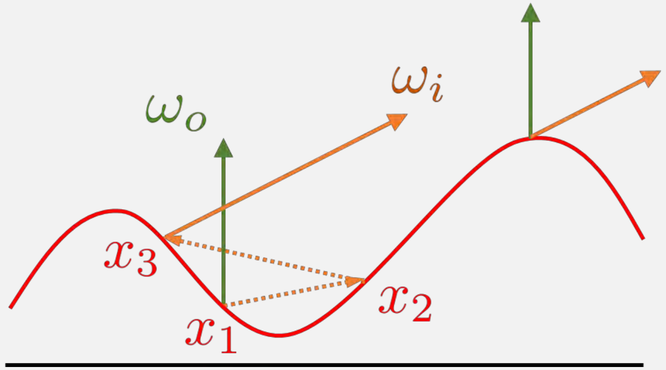

$$f^{\text{eff}}_{\text{gi}}(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o; \mathcal{P}, f)$$

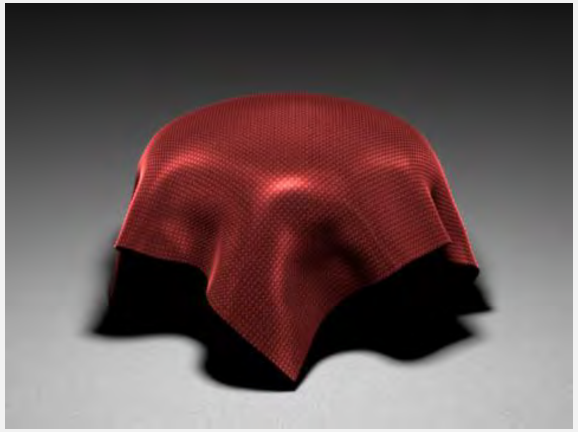

Without
interreflections

With
interreflections

Effective BRDF with Interreflections

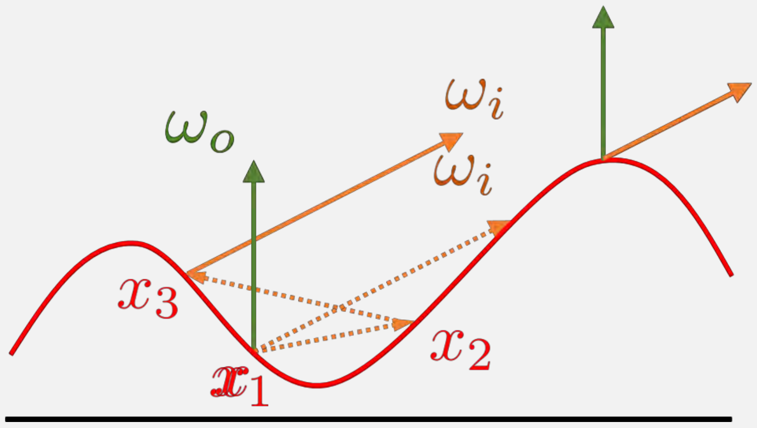

$$f^{\text{eff}}(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o; \mathcal{G}, f) = \frac{1}{A_{\mathcal{G}}(\boldsymbol{\omega}_o)} \int_{\mathcal{P}} \underbrace{f(\boldsymbol{x}_m(\boldsymbol{p}), \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \langle \boldsymbol{\omega}_m(p), \boldsymbol{\omega}_i \rangle V(\boldsymbol{x}_m(\boldsymbol{p}), \boldsymbol{\omega}_i)}_{\text{Single-bounce contribution}} A_{\mathcal{G}}(\boldsymbol{p}, \boldsymbol{\omega}_o) k_{\mathcal{P}}(\boldsymbol{p}) \, \mathrm{d}\boldsymbol{p}$$

$$f_{\text{ir}}^{\text{eff}}(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o; \mathcal{G}, f) = \frac{1}{A_{\mathcal{G}}(\boldsymbol{\omega}_o)} \int_{\mathcal{P}} \underbrace{\left[\int \cdots \int F(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) \, \mathrm{d}\mu(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) \right]}_{\text{Multi-bounce path integral}} A_{\mathcal{G}}(\boldsymbol{p}, \boldsymbol{\omega}_o) k_{\mathcal{P}}(\boldsymbol{p}) \, \mathrm{d}\boldsymbol{p}$$

Overview

Joint prefiltering

Appearance matching

LoD rendering

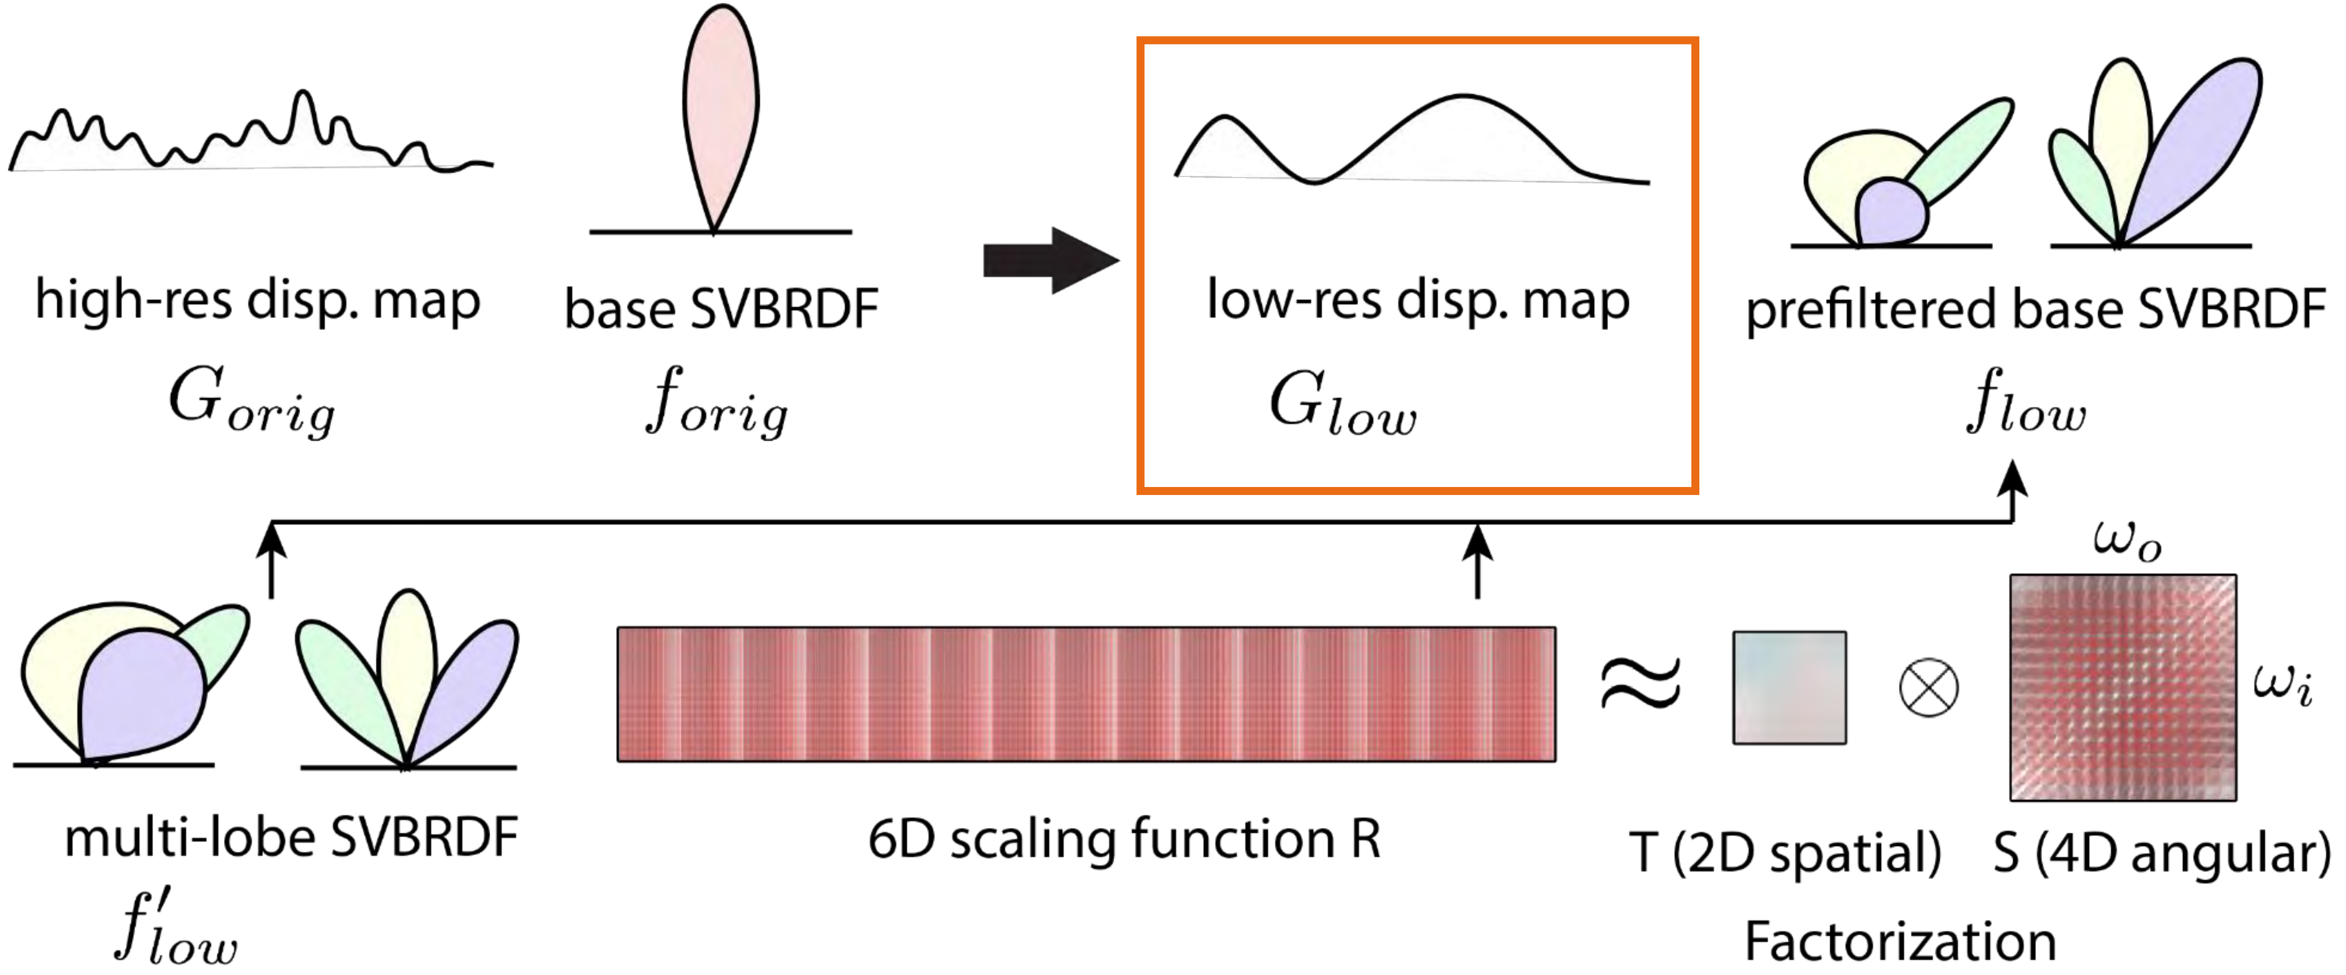

Downsampling Displacement Maps

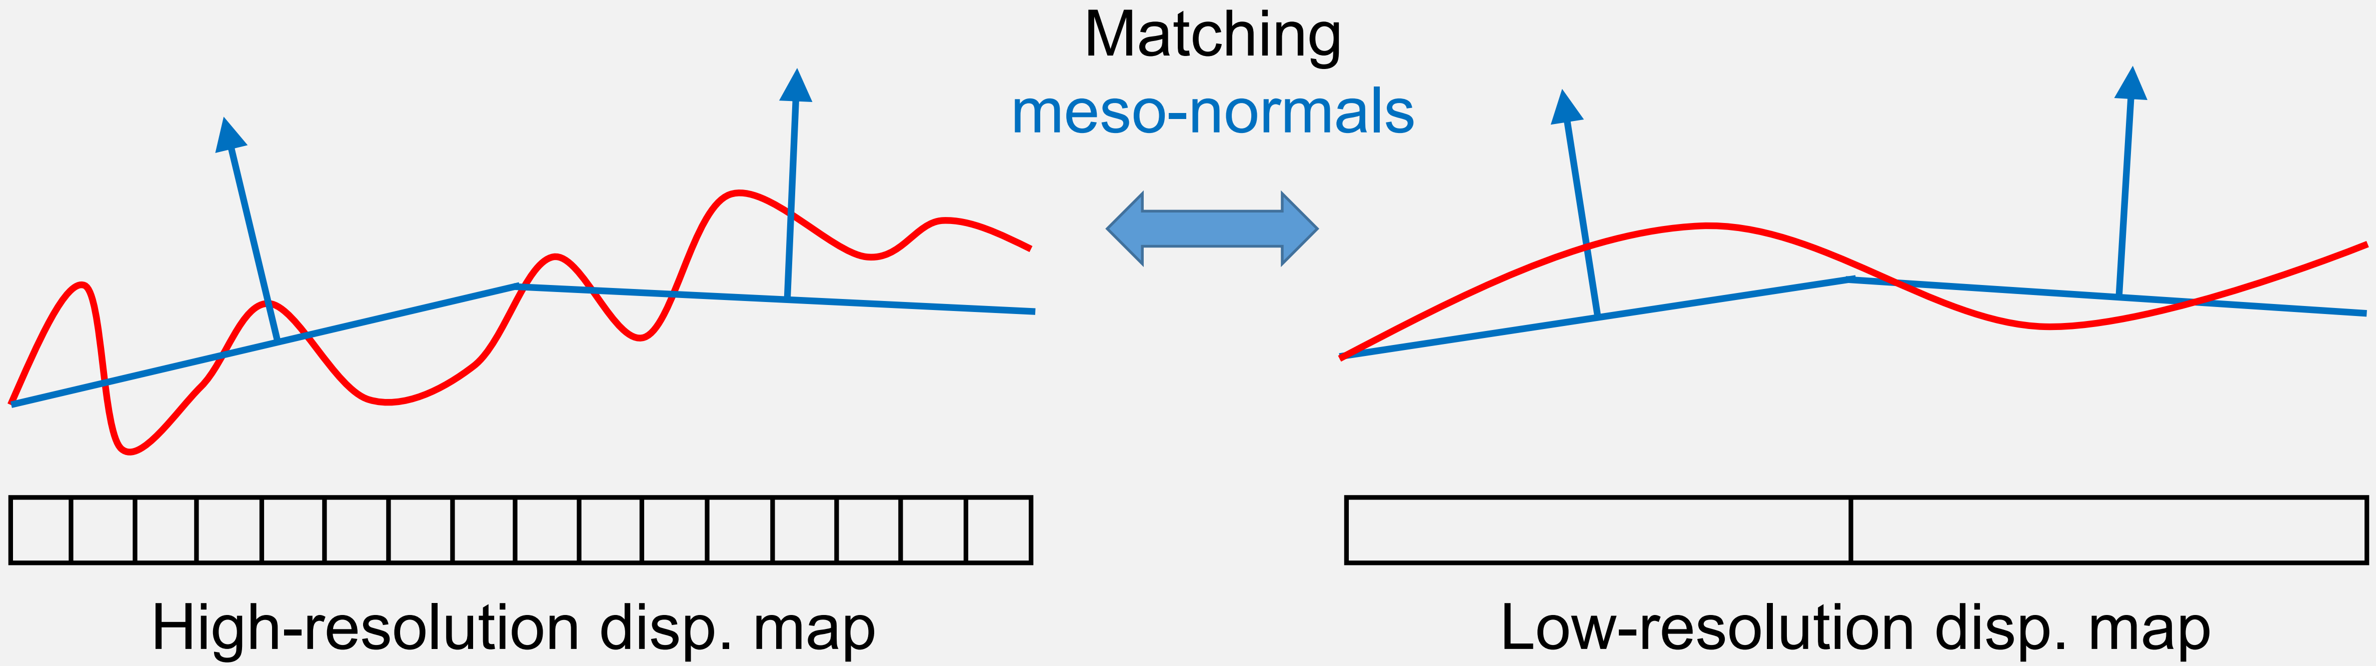

- Solved using least-squares

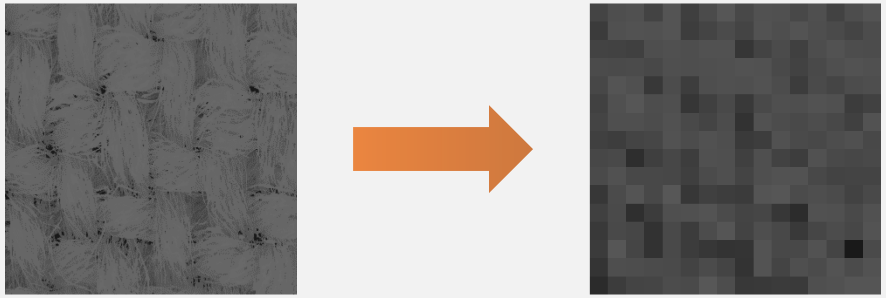

Overview

Joint prefiltering

Appearance matching

LoD rendering

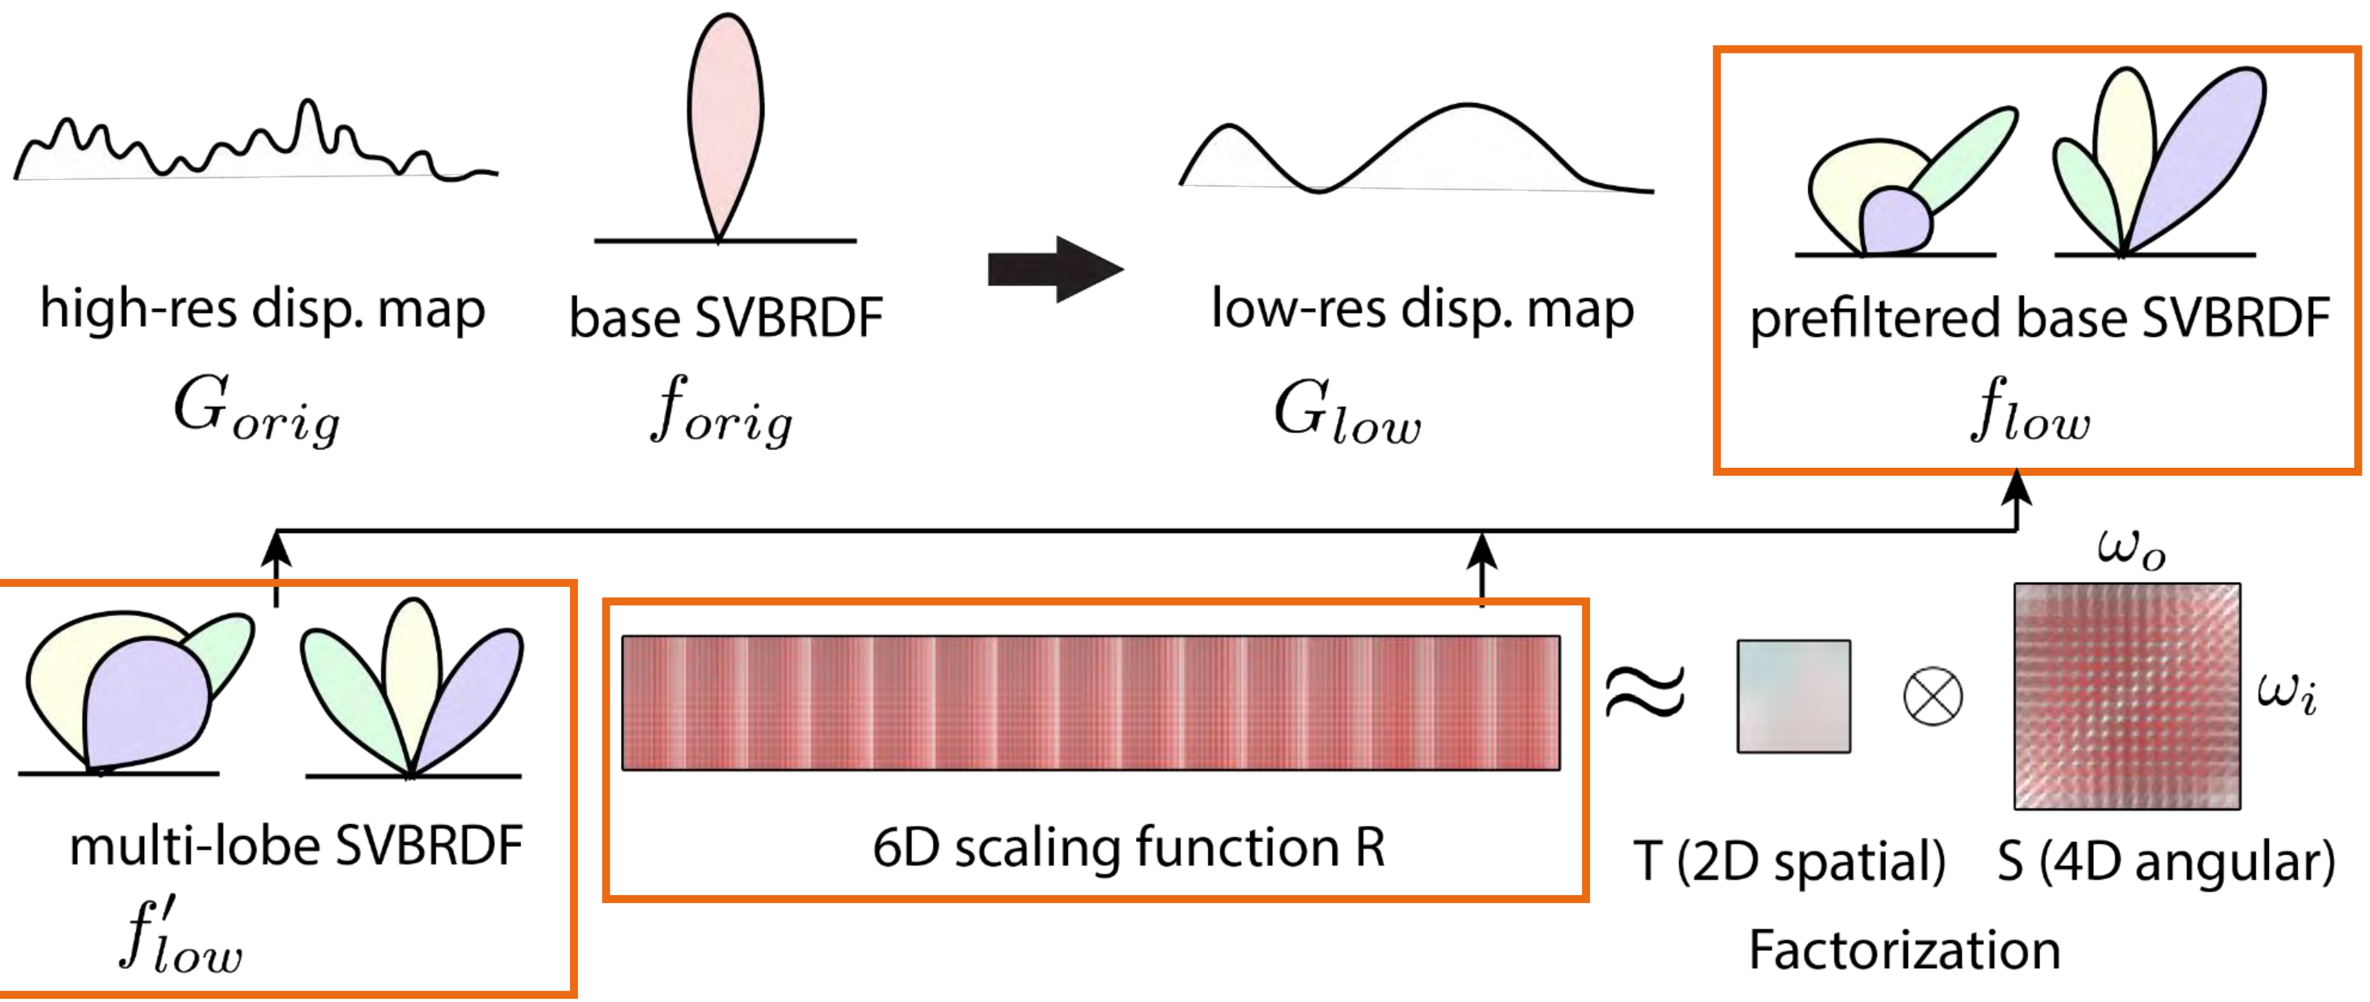

Step 1: Multi-Lobe SVBRDF

- NDF: A (hemi-)spherical distribution of normal directions

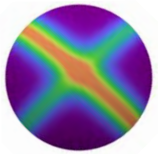

- Statistical representation: decorrelating positions and normals

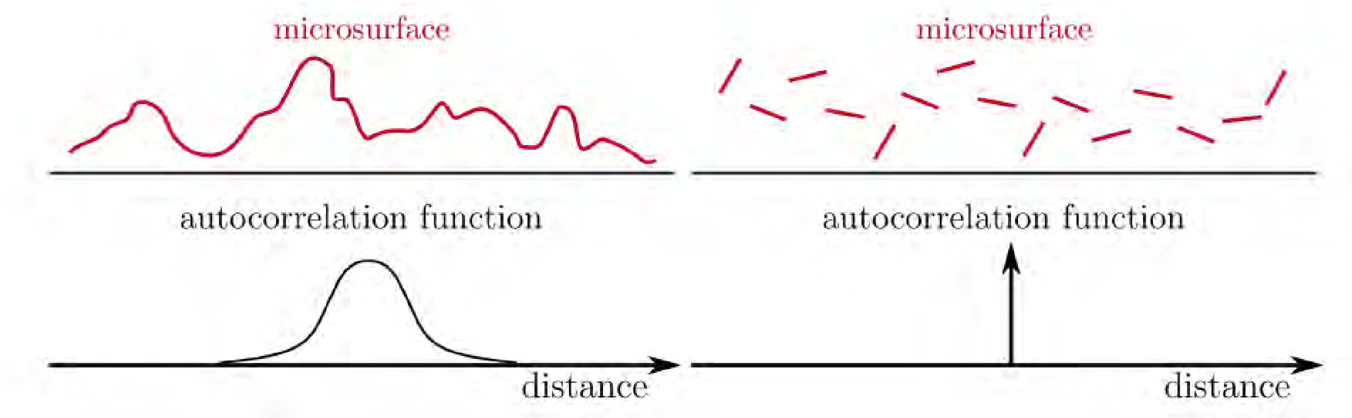

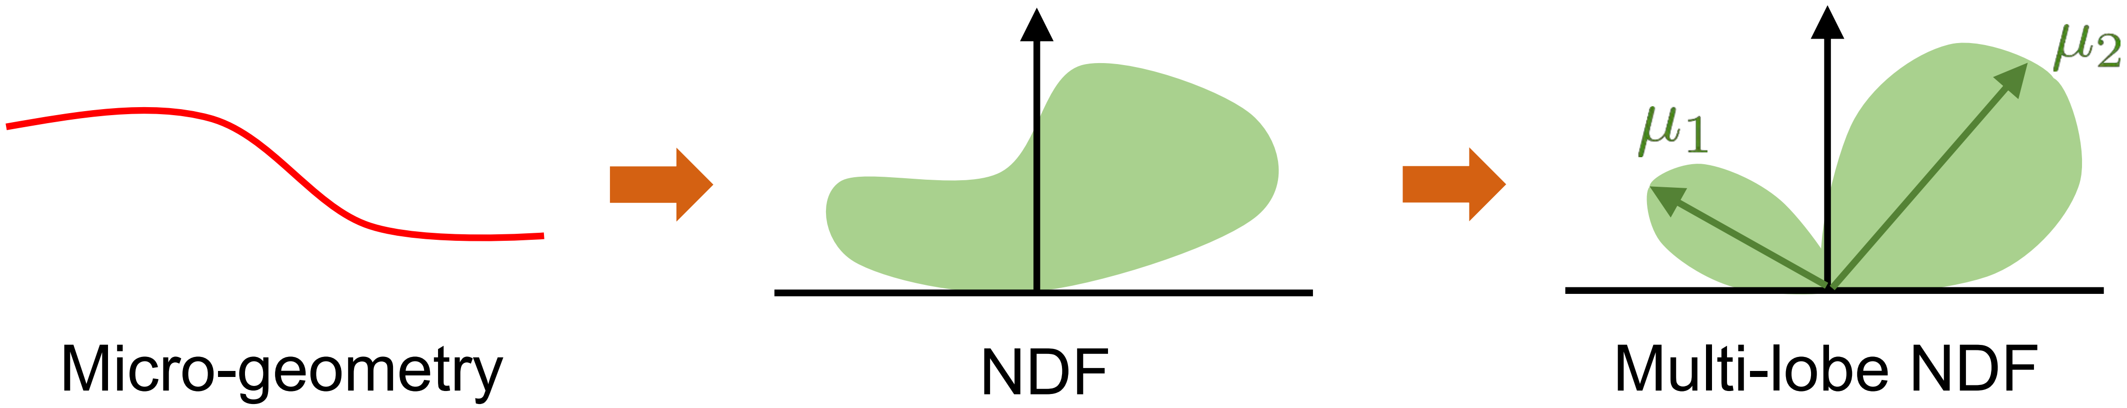

Image courtesy of [Heitz 2014]

Step 1: Multi-Lobe SVBRDF

- Normal mapping [Han et al. 2007]

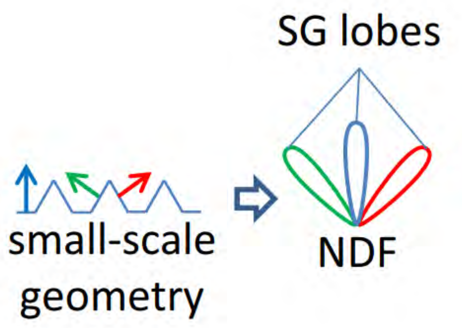

- Multi-lobe BRDF = Multi-lobe NDF $\otimes$ Micro-BRDF

Overview

Joint prefiltering

Appearance matching

LoD rendering

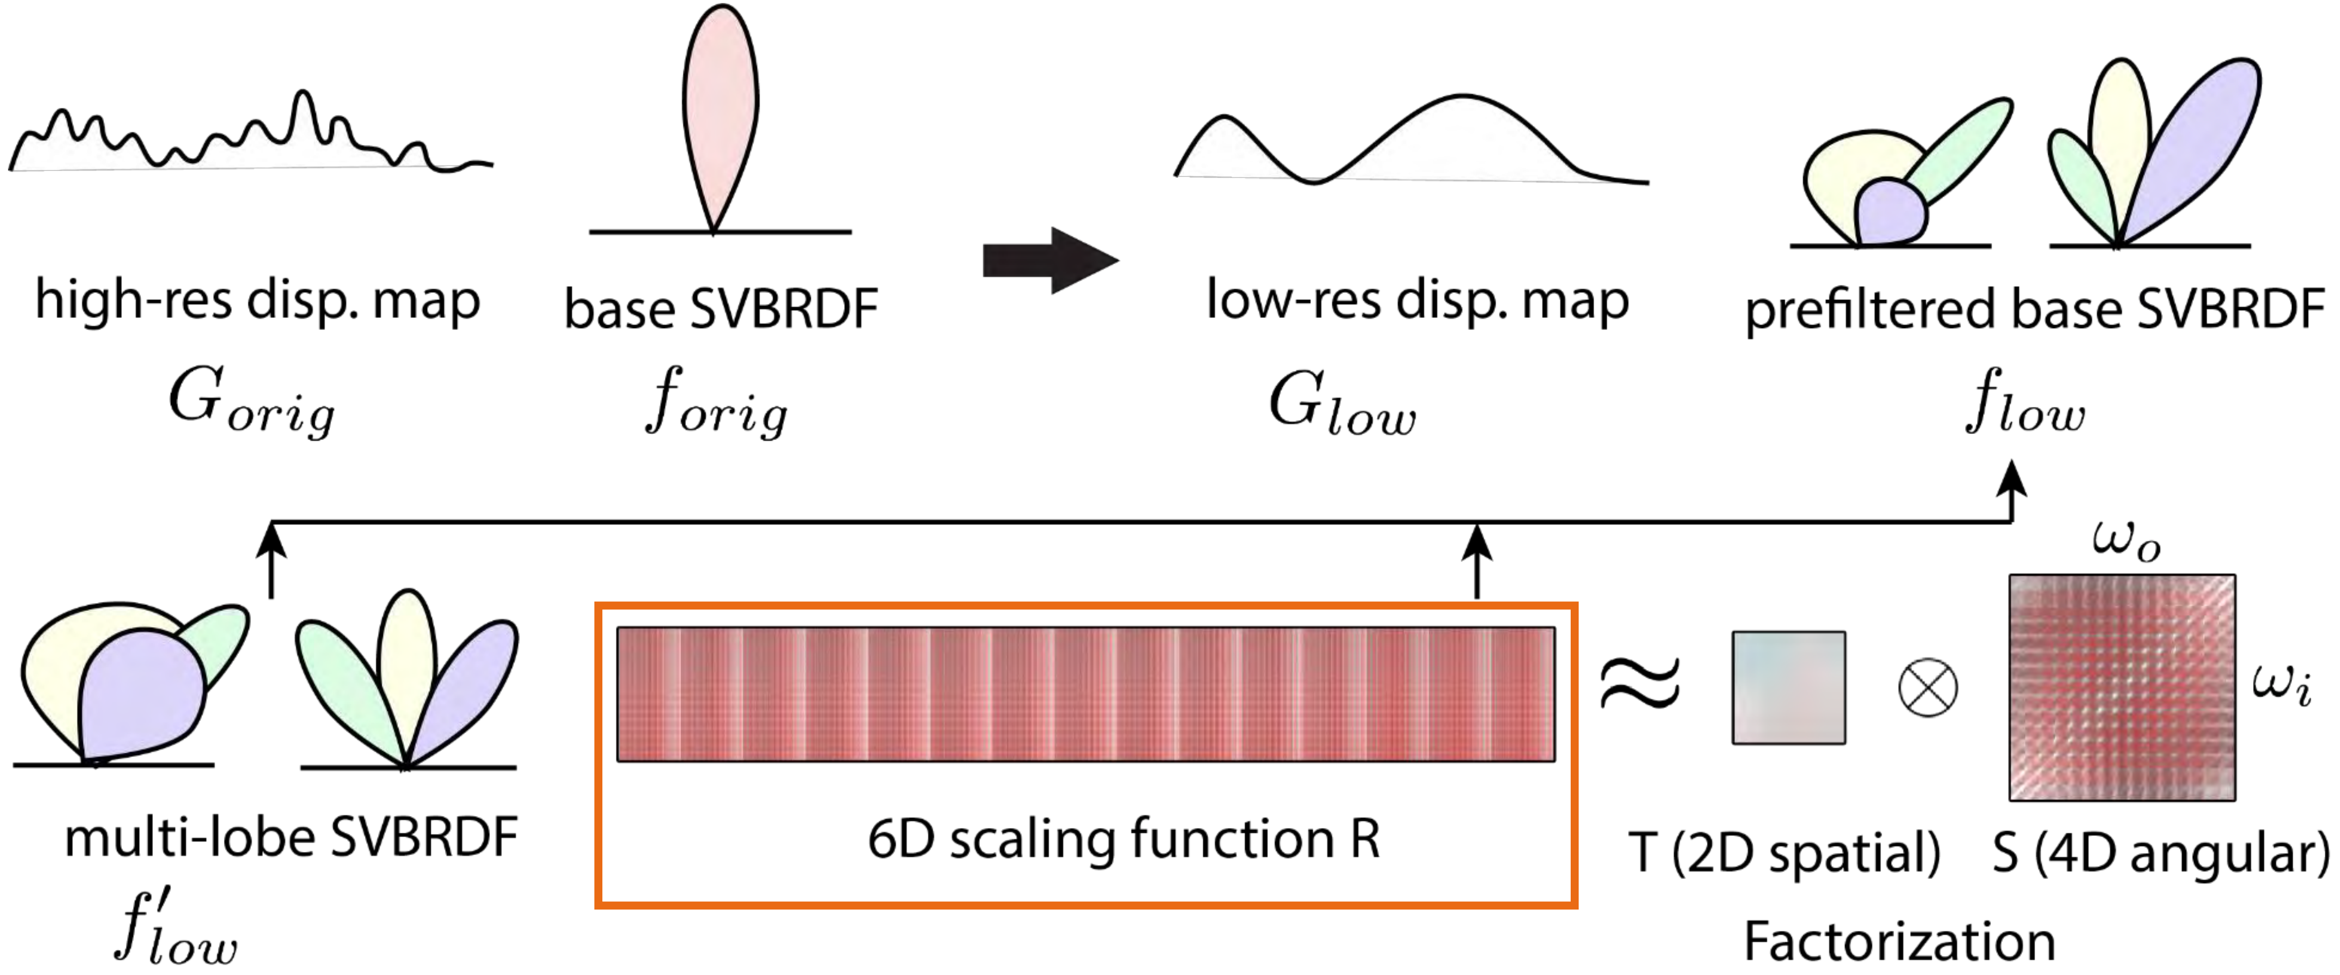

Step 2: Scaling Function

- Matching effective BRDFs

$$f_{\mathrm{ir}}^{\mathrm{eff}}(\mathcal{G}_{\mathrm{orig}}, f_{\mathrm{orig}}) \approx f_{\mathrm{ir}}^{\mathrm{eff}}(\mathcal{G}_{\mathrm{low}}, \boxed{R_{\mathrm{ir}}} \cdot f'_{\mathrm{low}})$$

- Computing the scaling function directly: $R_{\mathrm{ir}} = \dfrac{f_{\mathrm{ir}}^{\mathrm{eff}}(\mathcal{G}_{\mathrm{orig}}, f_{\mathrm{orig}})}{f_{\mathrm{ir}}^{\mathrm{eff}}(\mathcal{G}_{\mathrm{low}}, f'_{\mathrm{low}})}$
  - No need for iterative optimization

- Not a practical algorithm

Overview

Joint prefiltering

Appearance matching

LoD rendering

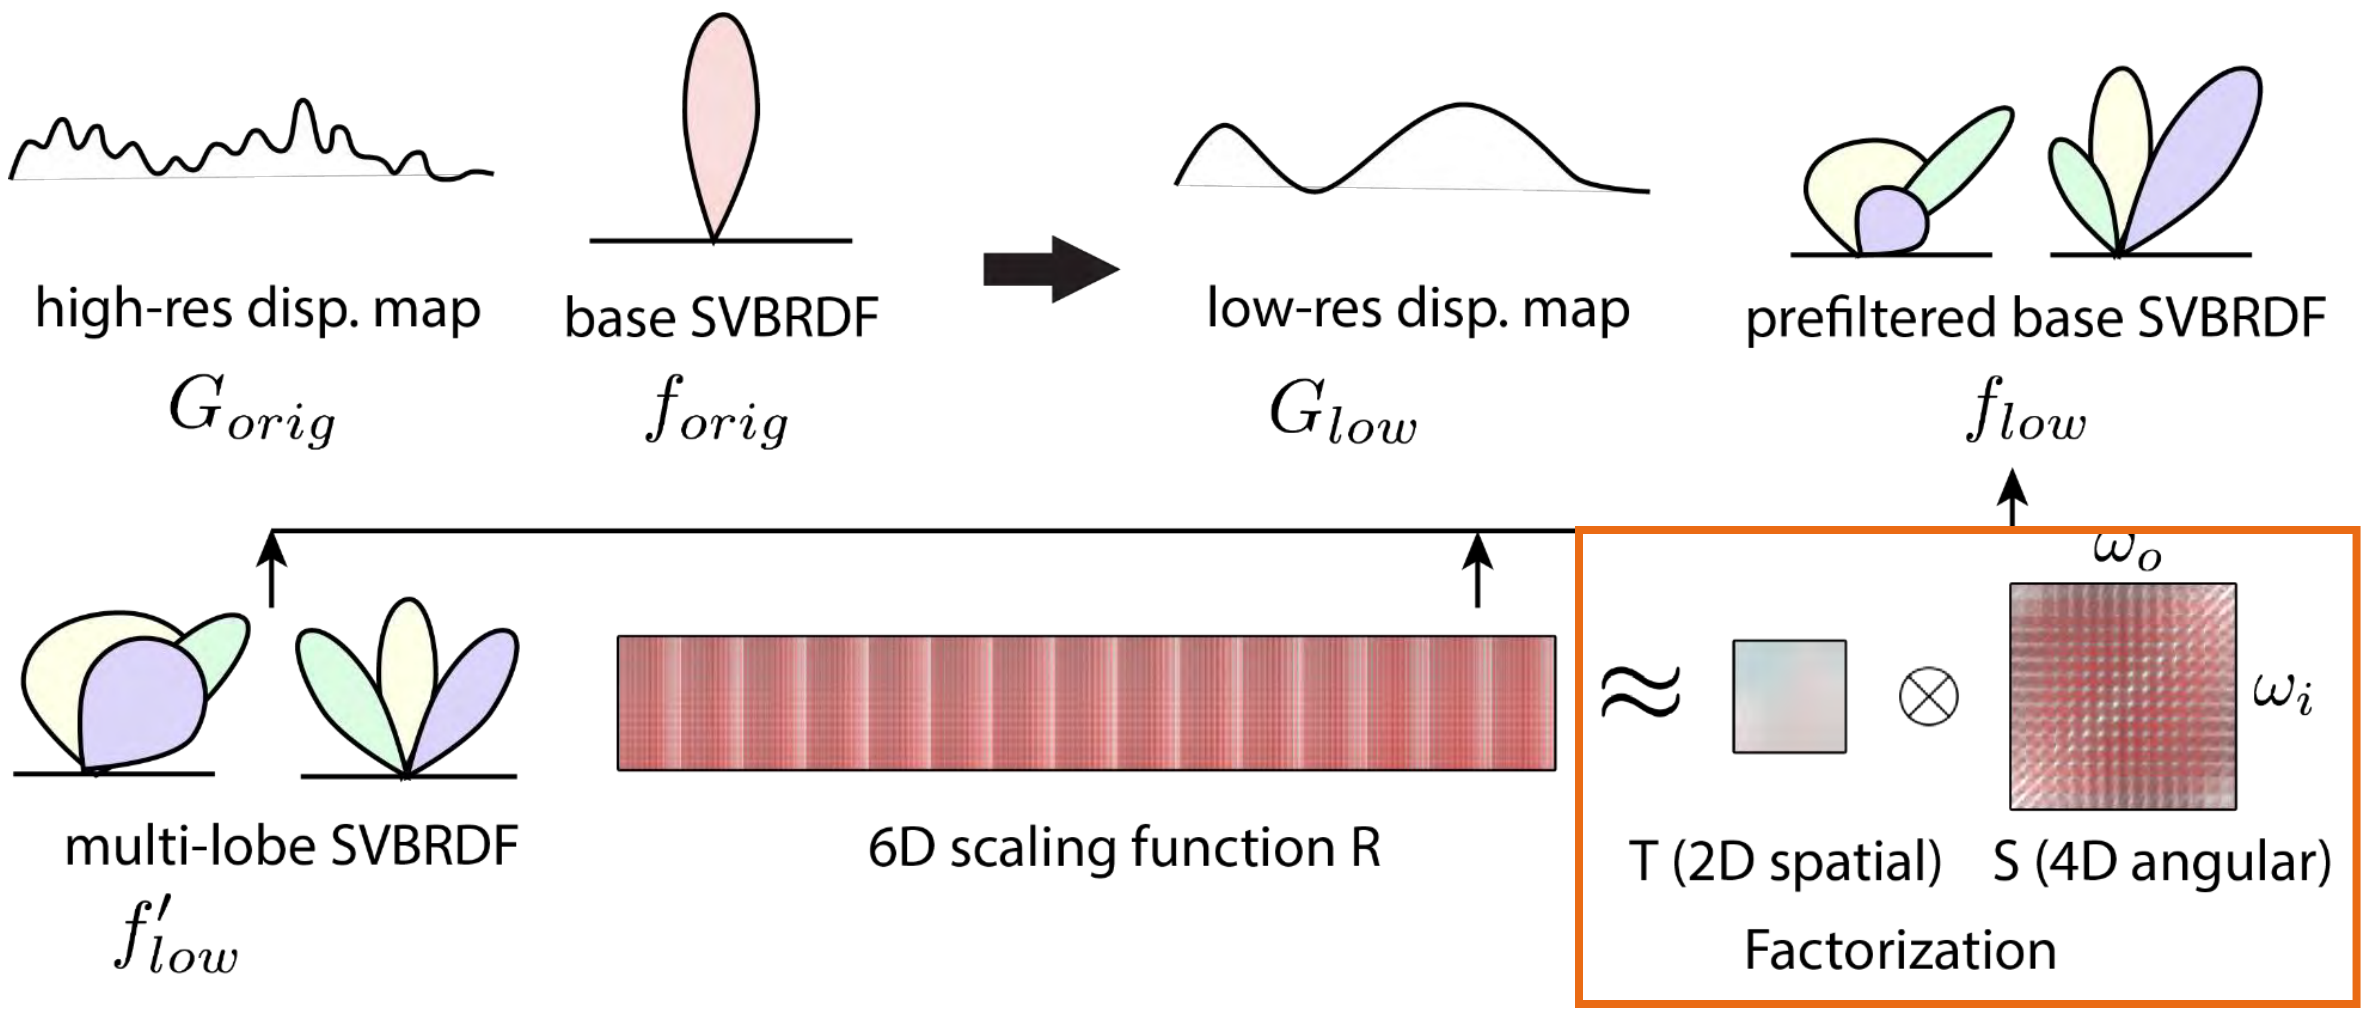

Efficient Factorization

- Impractical to compute and store the full 6D scaling function
- Rank-1 factorization $R_{\mathrm{ir}}(\boldsymbol{x}, \boldsymbol{\omega}_i, \boldsymbol{\omega}_o) \approx T_{\mathrm{ir}}(\boldsymbol{x}) \cdot S_{\mathrm{ir}}(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$

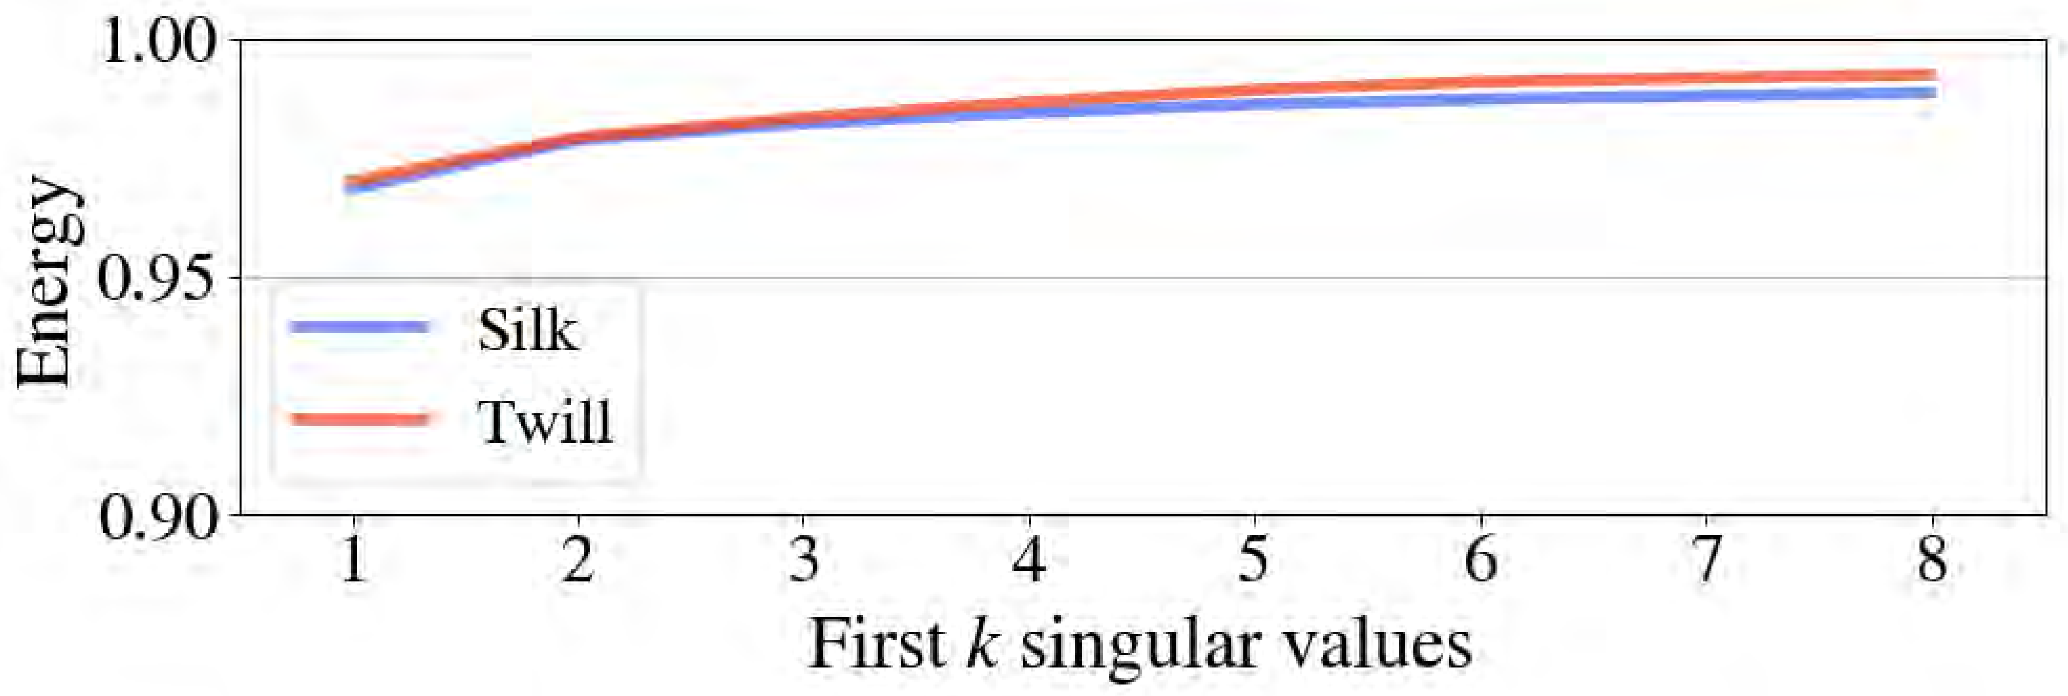

Efficient Factorization

- $T_{\mathrm{ir}}(\boldsymbol{x})$ and $S_{\mathrm{ir}}(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ can be tabulated coarsely (4^2 and 15^4)
- They can be reconstructed from sparse 6D samples

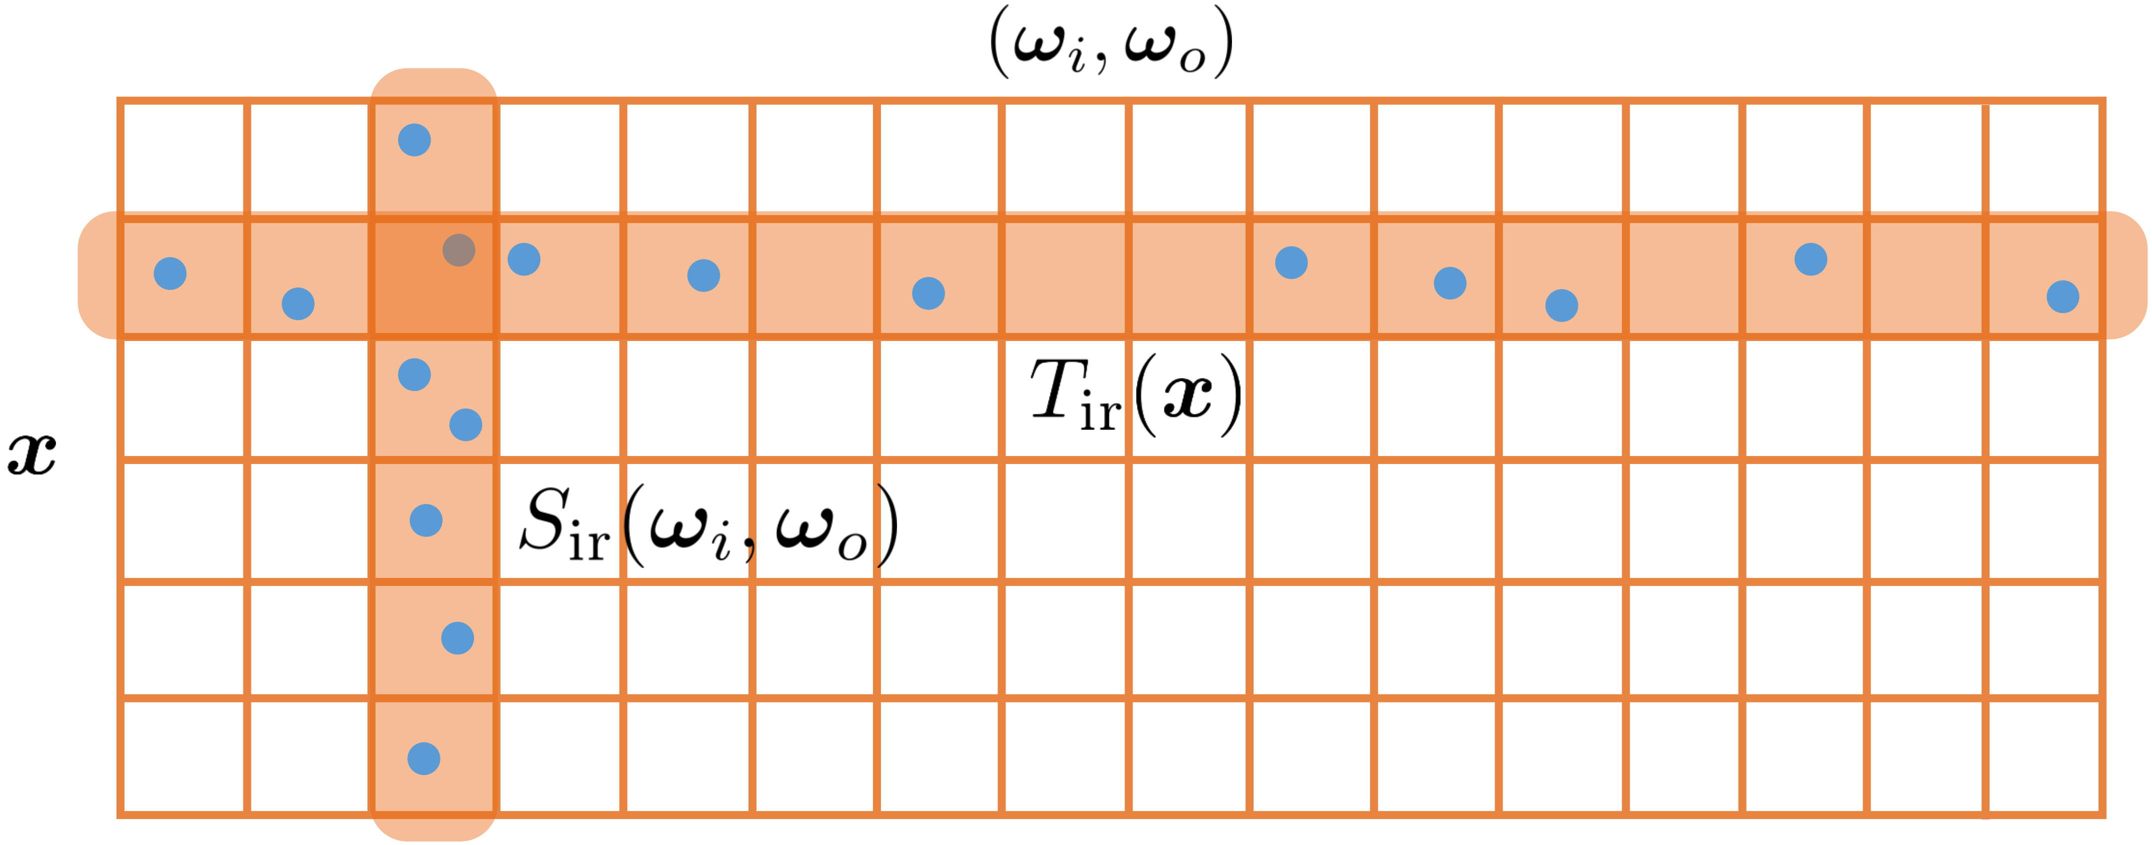

Efficient Factorization

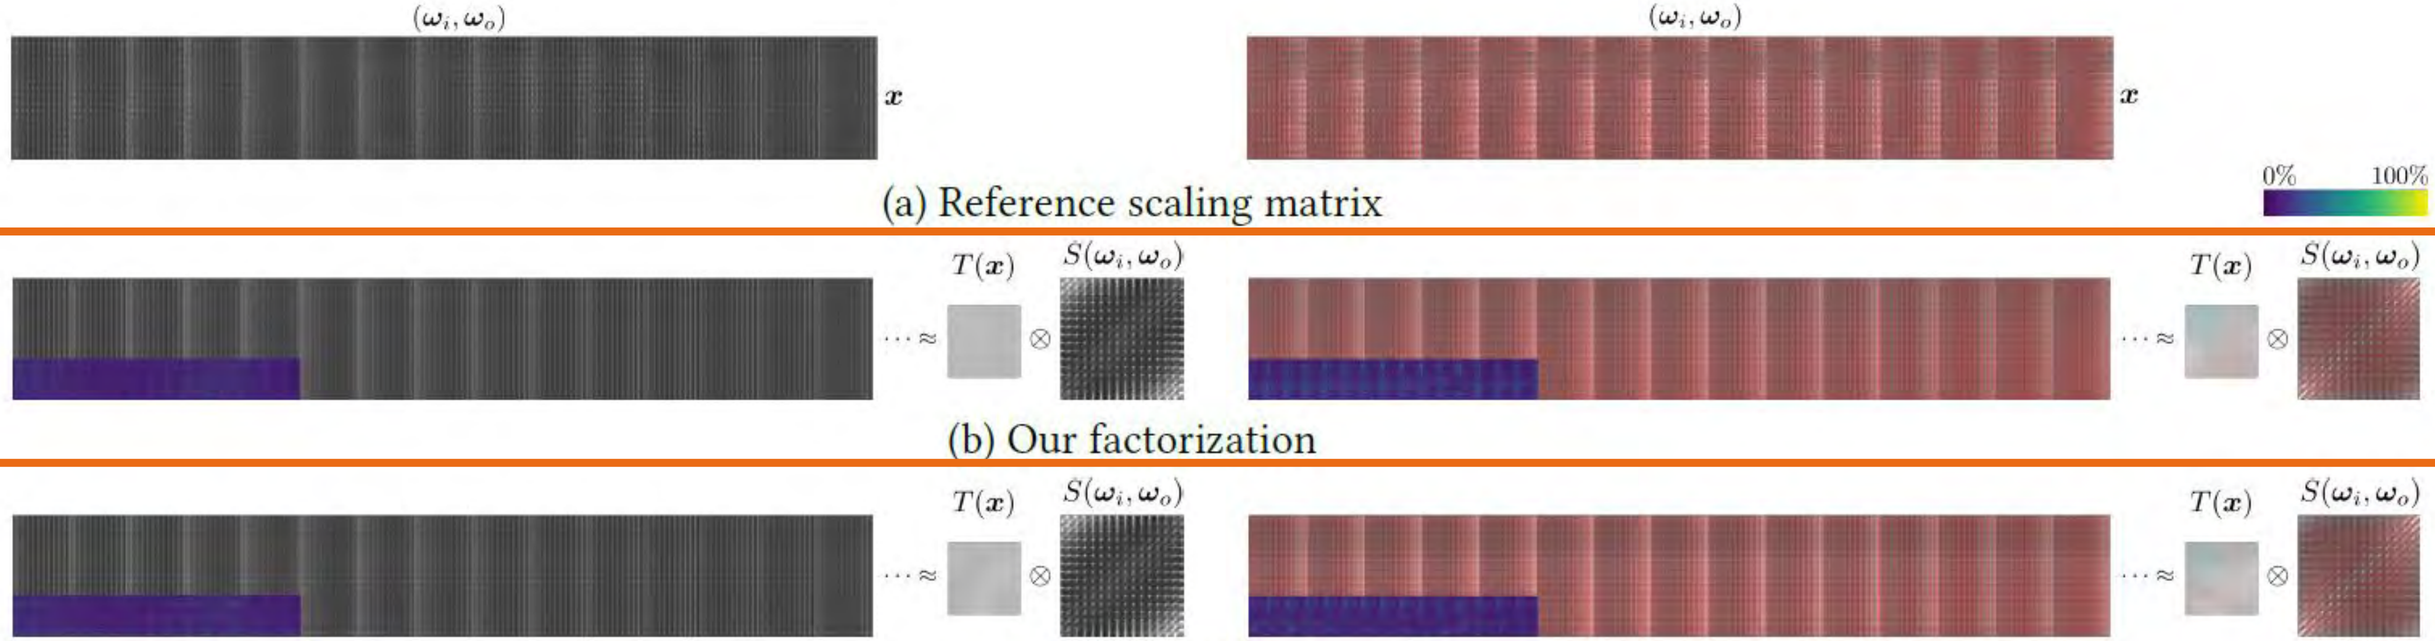

(a) Reference scaling matrix

(b) Our factorization

(c) Rank-1 SVD

Single Scale

✓ Joint prefiltering

✓ Appearance matching

? LoD rendering

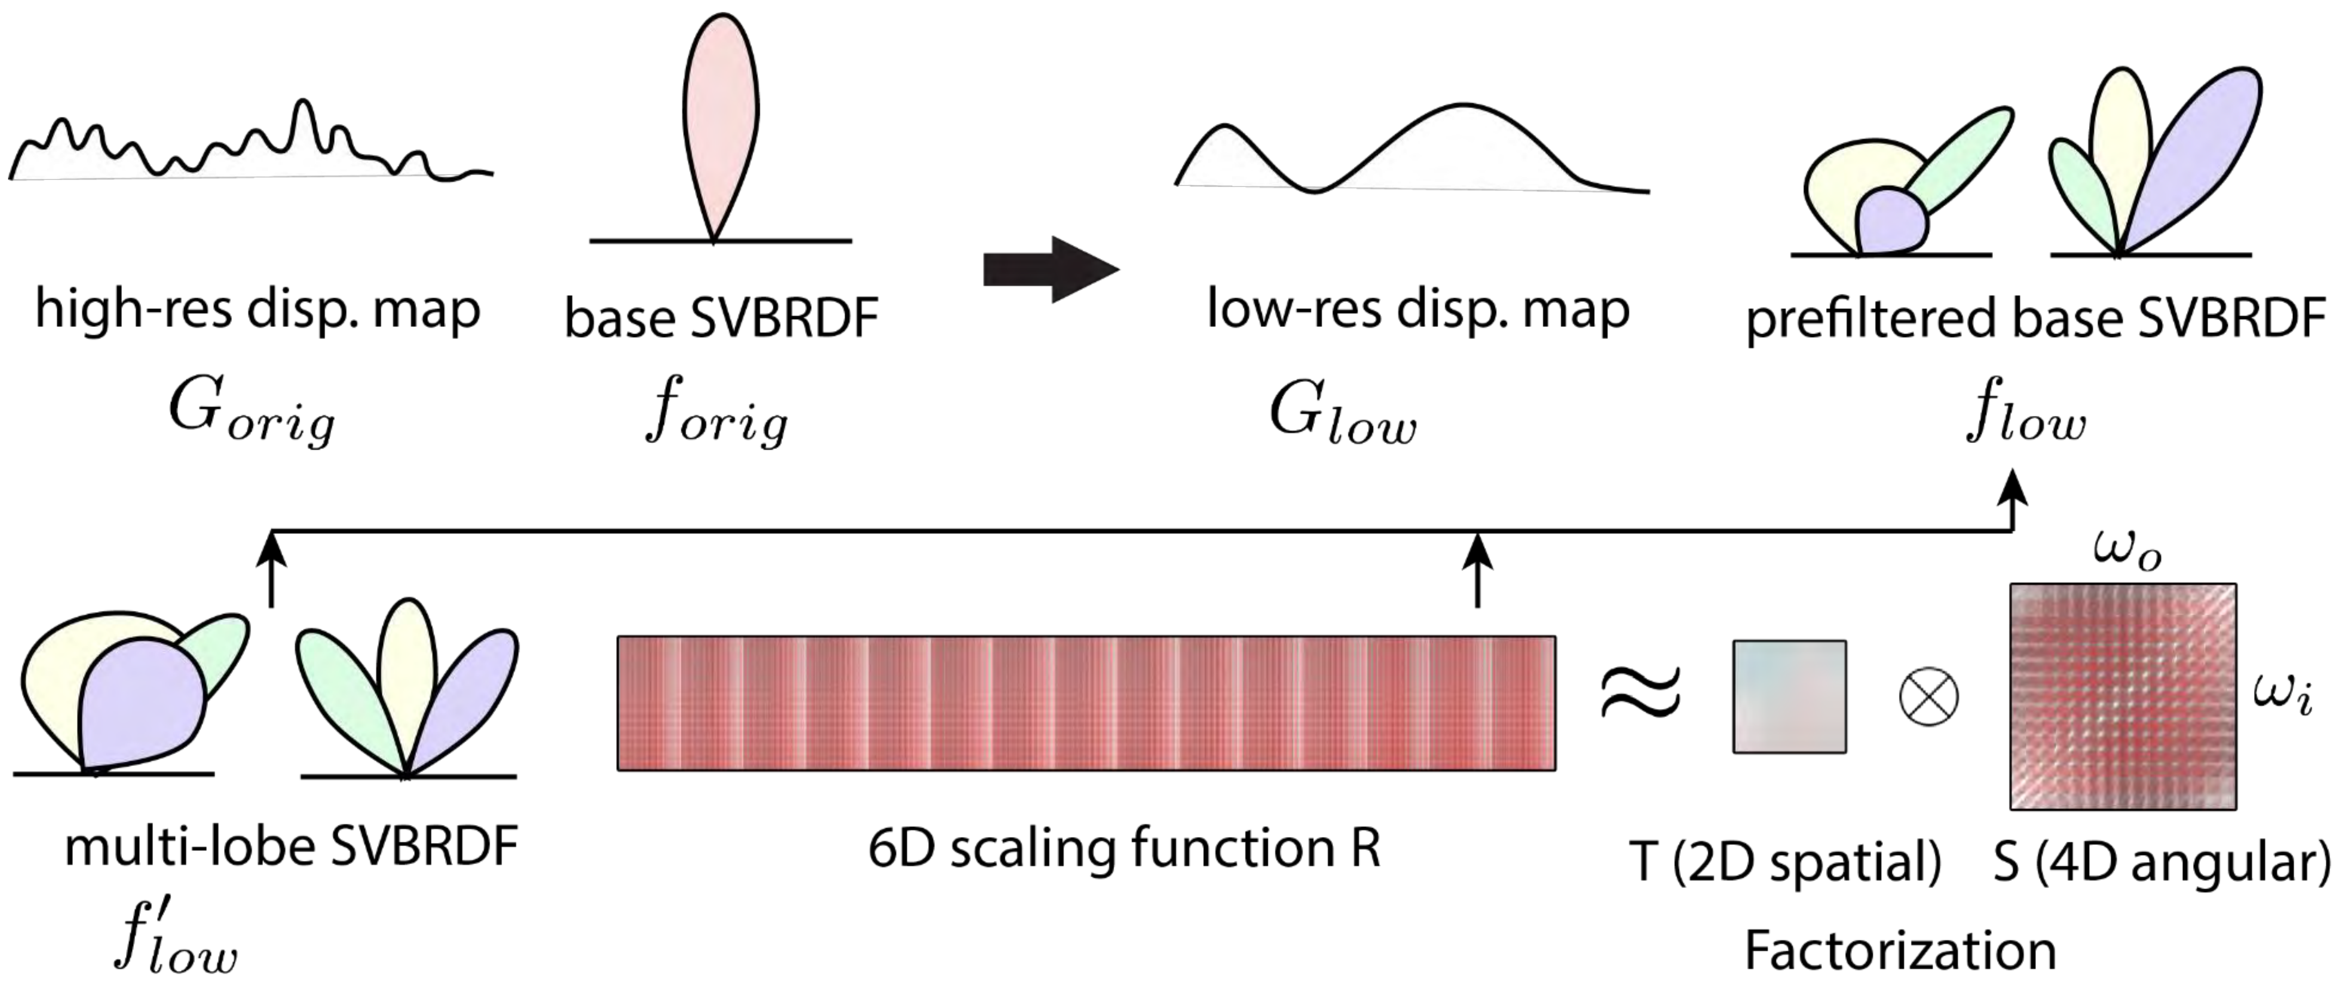

Multi-Scale LoD

✓ Joint prefiltering

✓ Appearance matching

✓ LoD rendering

- Prefilter at each mipmap level
- Interpolate path contributions traced on different levels

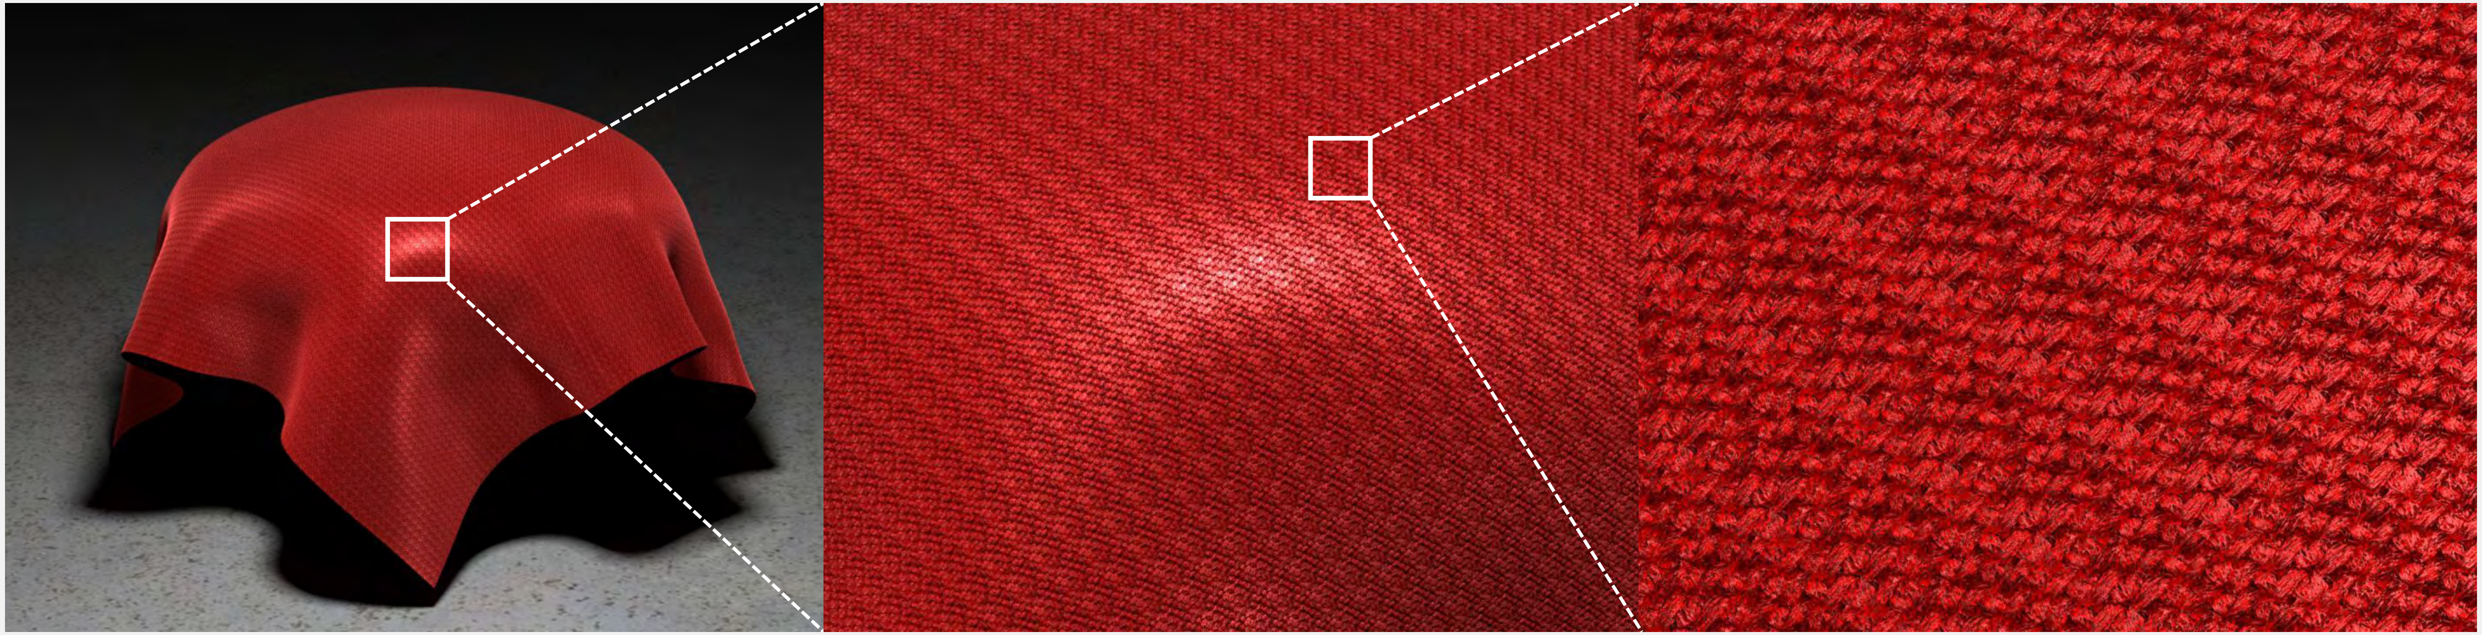

Results

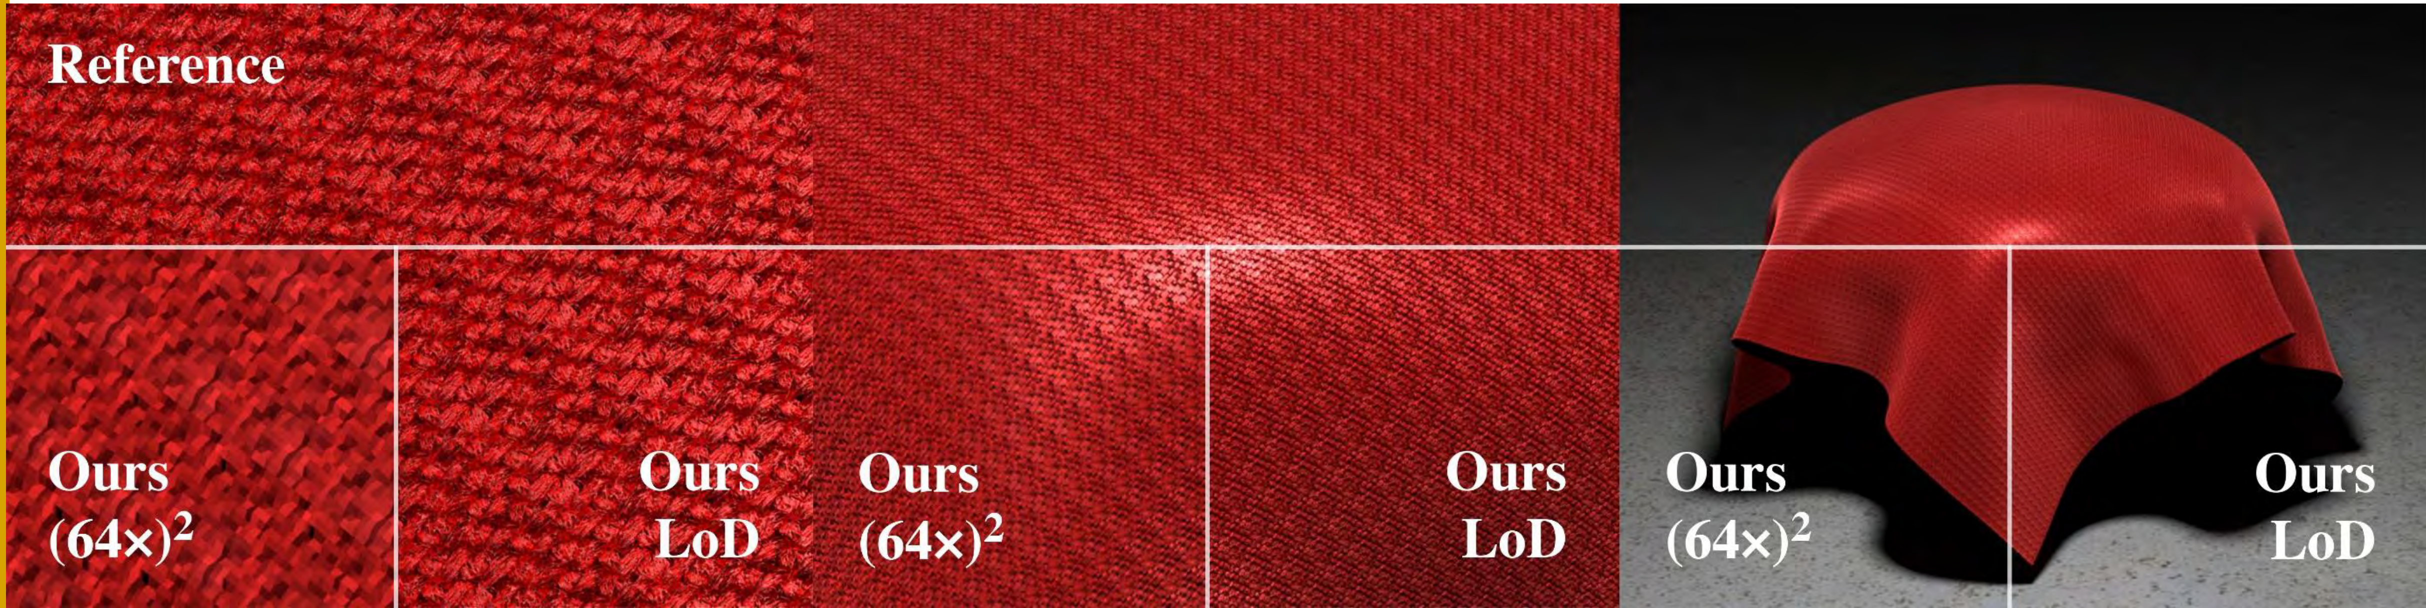
Reference
Ours
(64×)²
Ours
LoD
Ours
(64×)²
Ours
LoD
Ours
(64×)²
Ours
LoD

Scaling Function Resolution

- Determine angular resolutions (ω_i: 15^2, ω_o: 15^2)

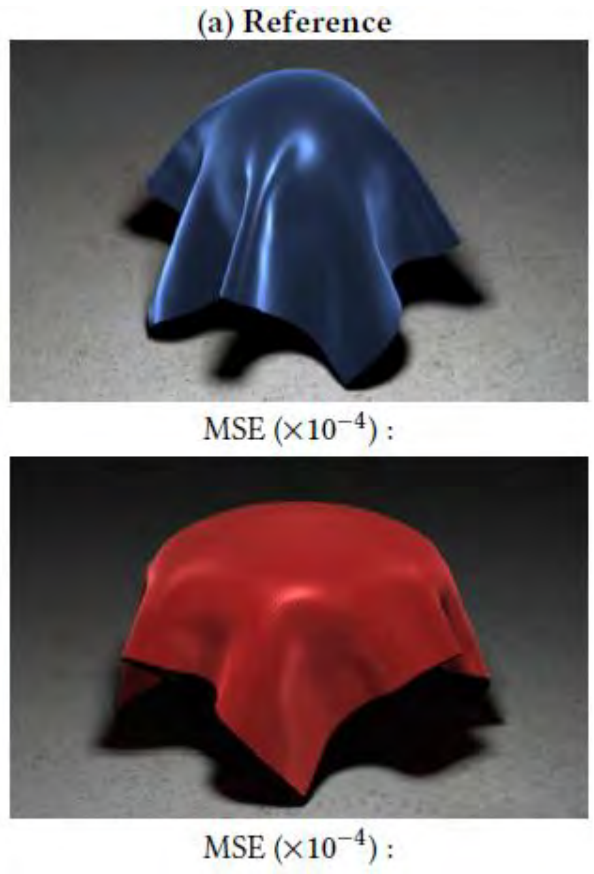

Scaling Function Resolution

- Determine spatial resolutions (uv: 4^2)

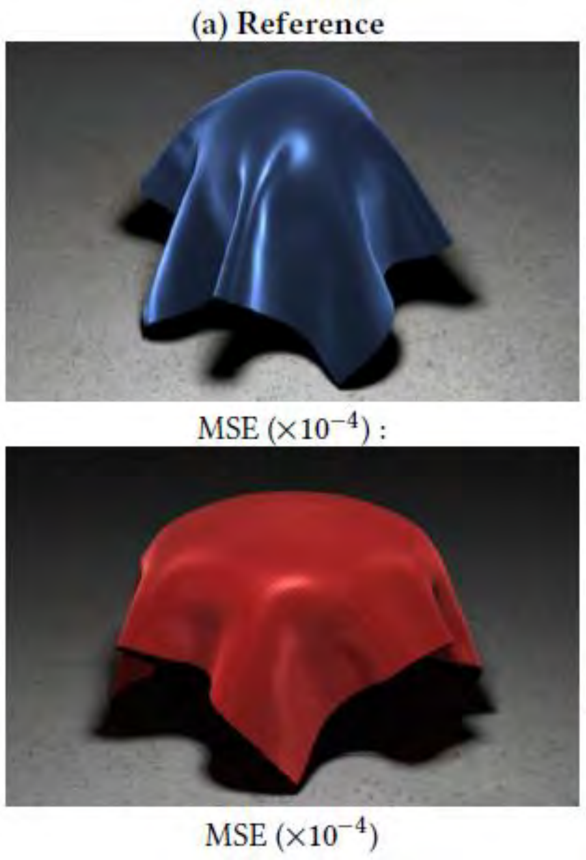

Validations

- Energy conservation

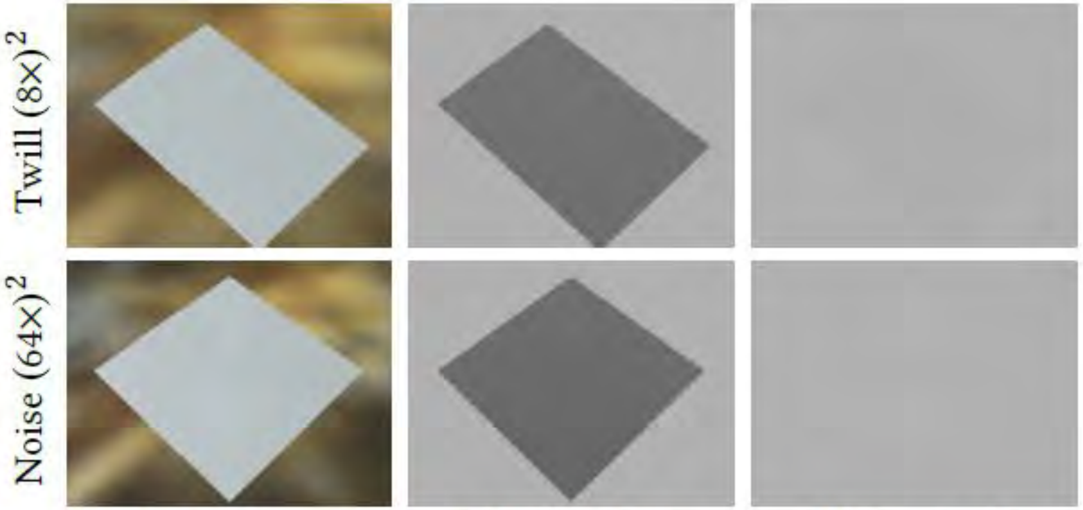

(a) Environmental lighting (b) Constant lighting with R (c) Constant lighting with R_{ir}

- Synthetic two-color V-grooves

(a) Reference (b) Ours ($R_{\text{ir}} \cdot f'_{\text{low}}$) (c) Multi-Lobe ($f'_{\text{low}}$)

Accuracy Comparison

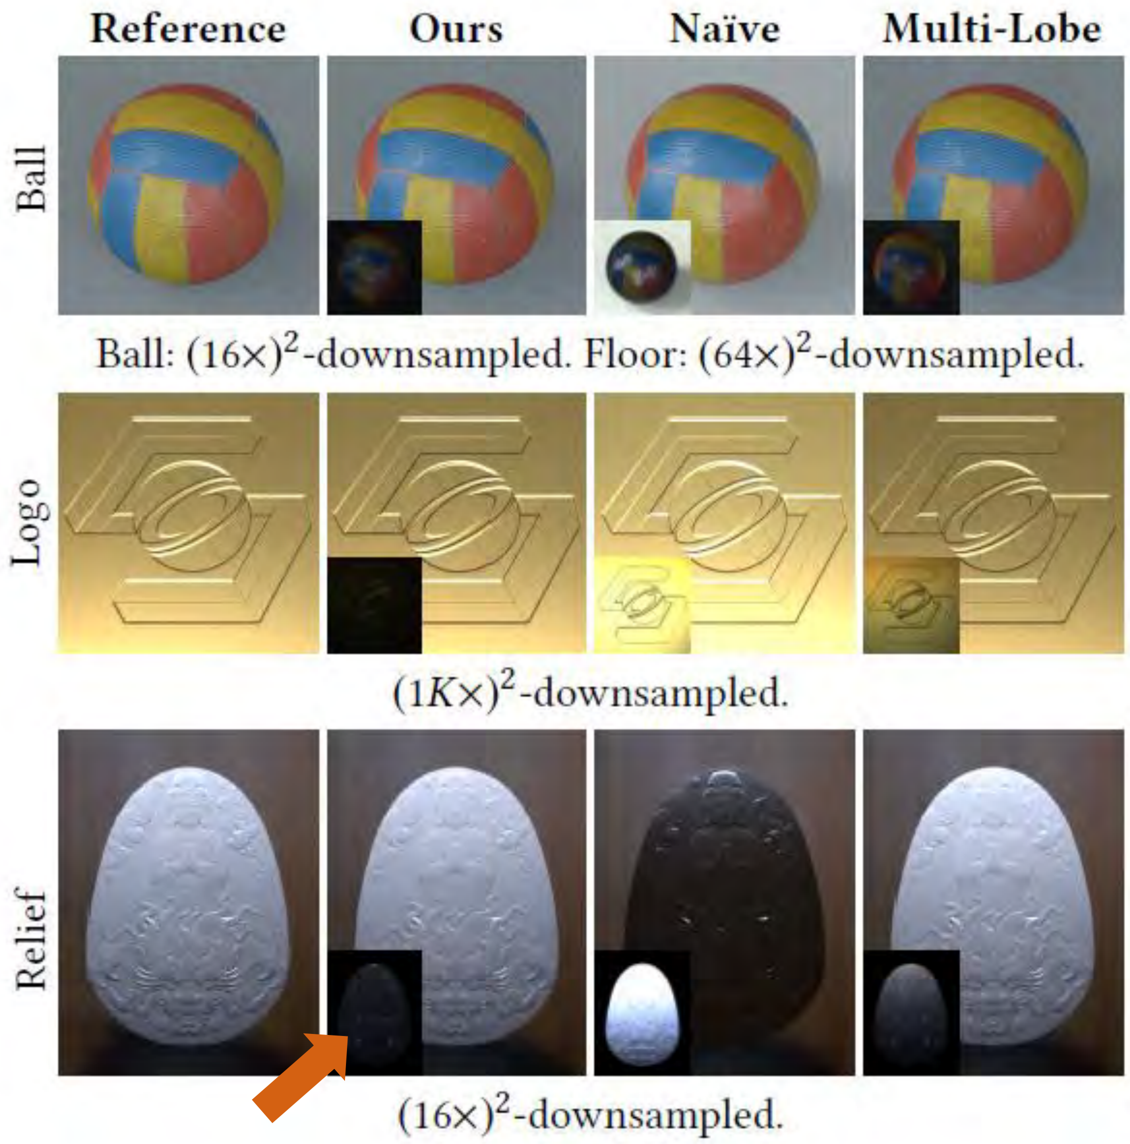

LoD Rendering

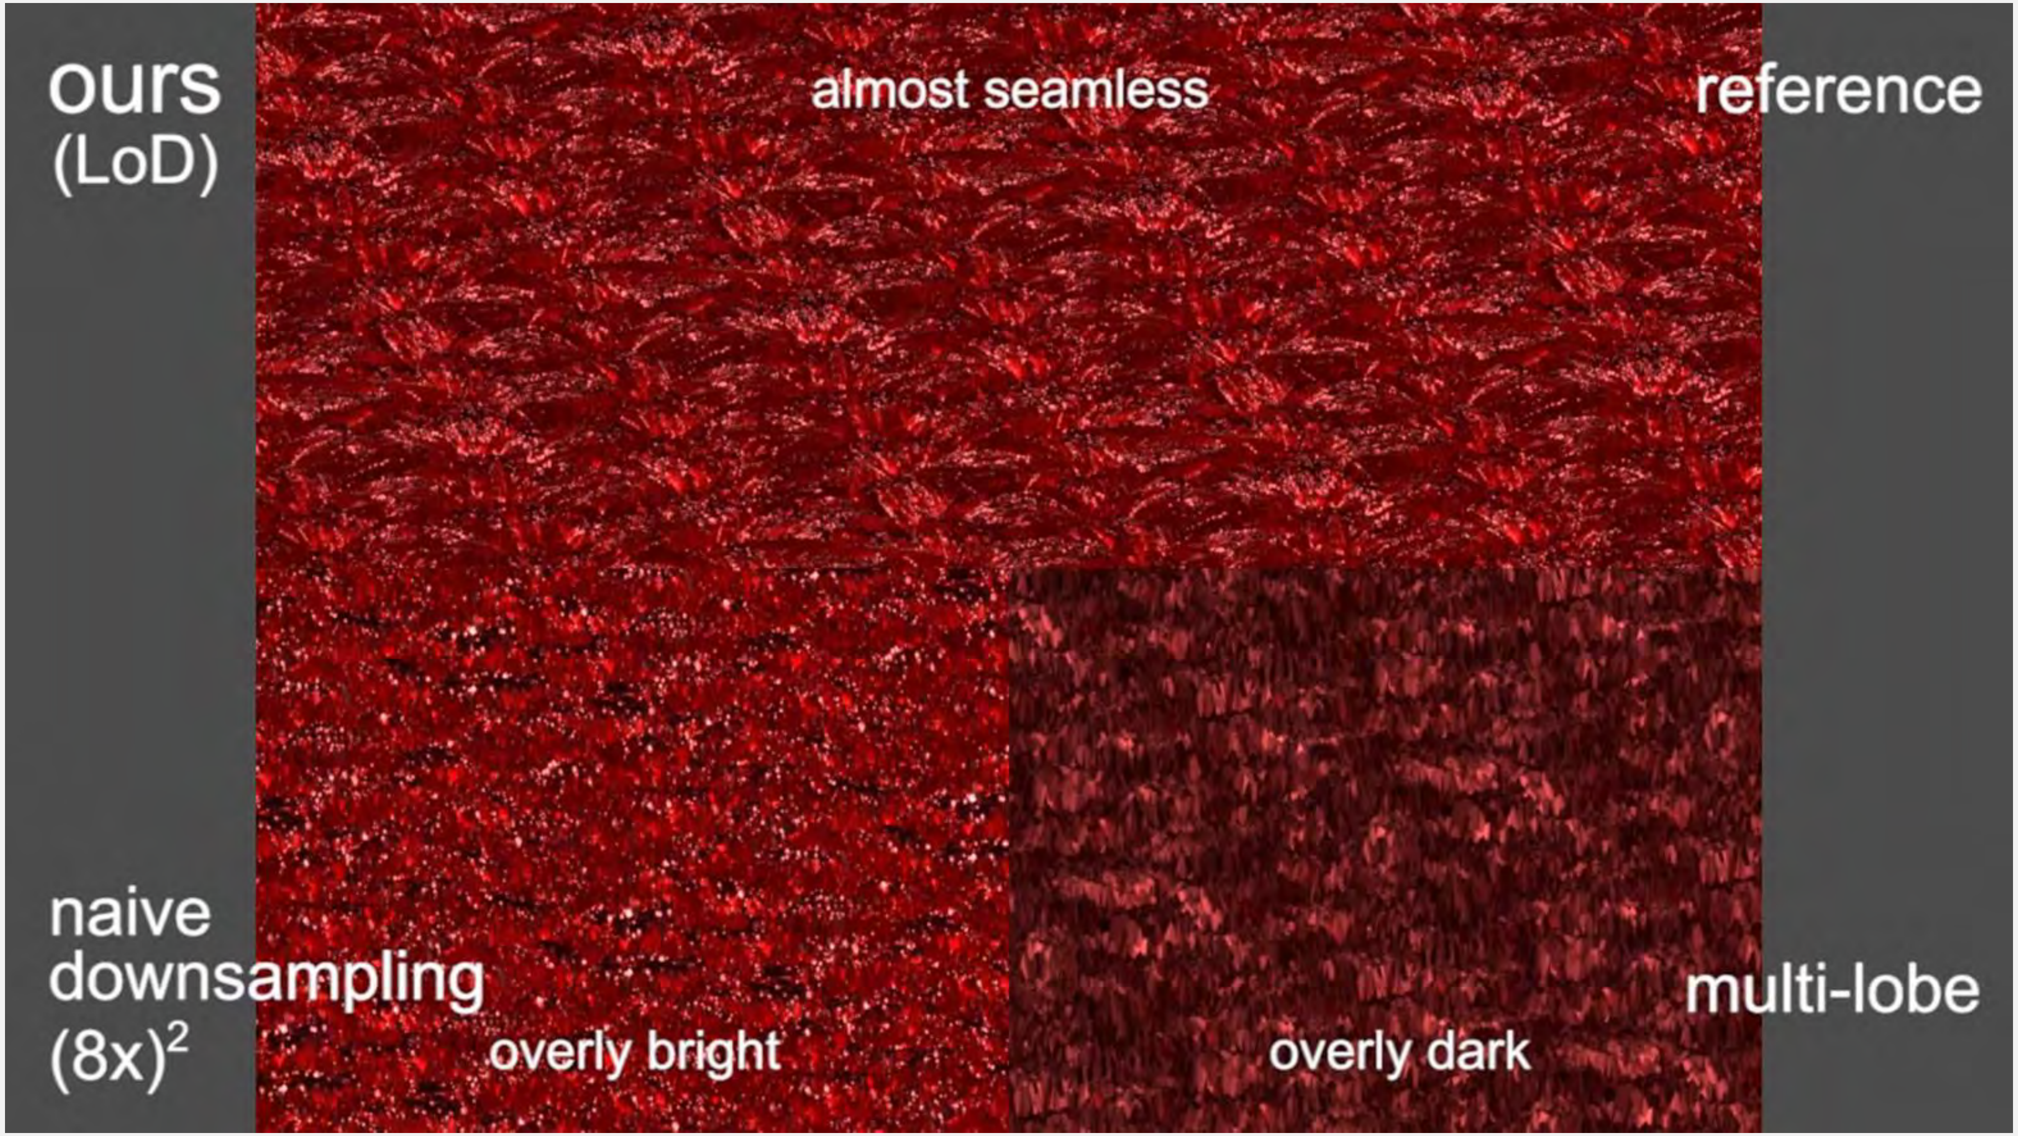

Changing Lighting/Viewing

Limitation / Future Work

- Fail when the vertical displacements are large
- Rely on model-dependent precomputation

- Theoretical analysis of appearance prefiltering
- Material editing

Future Work

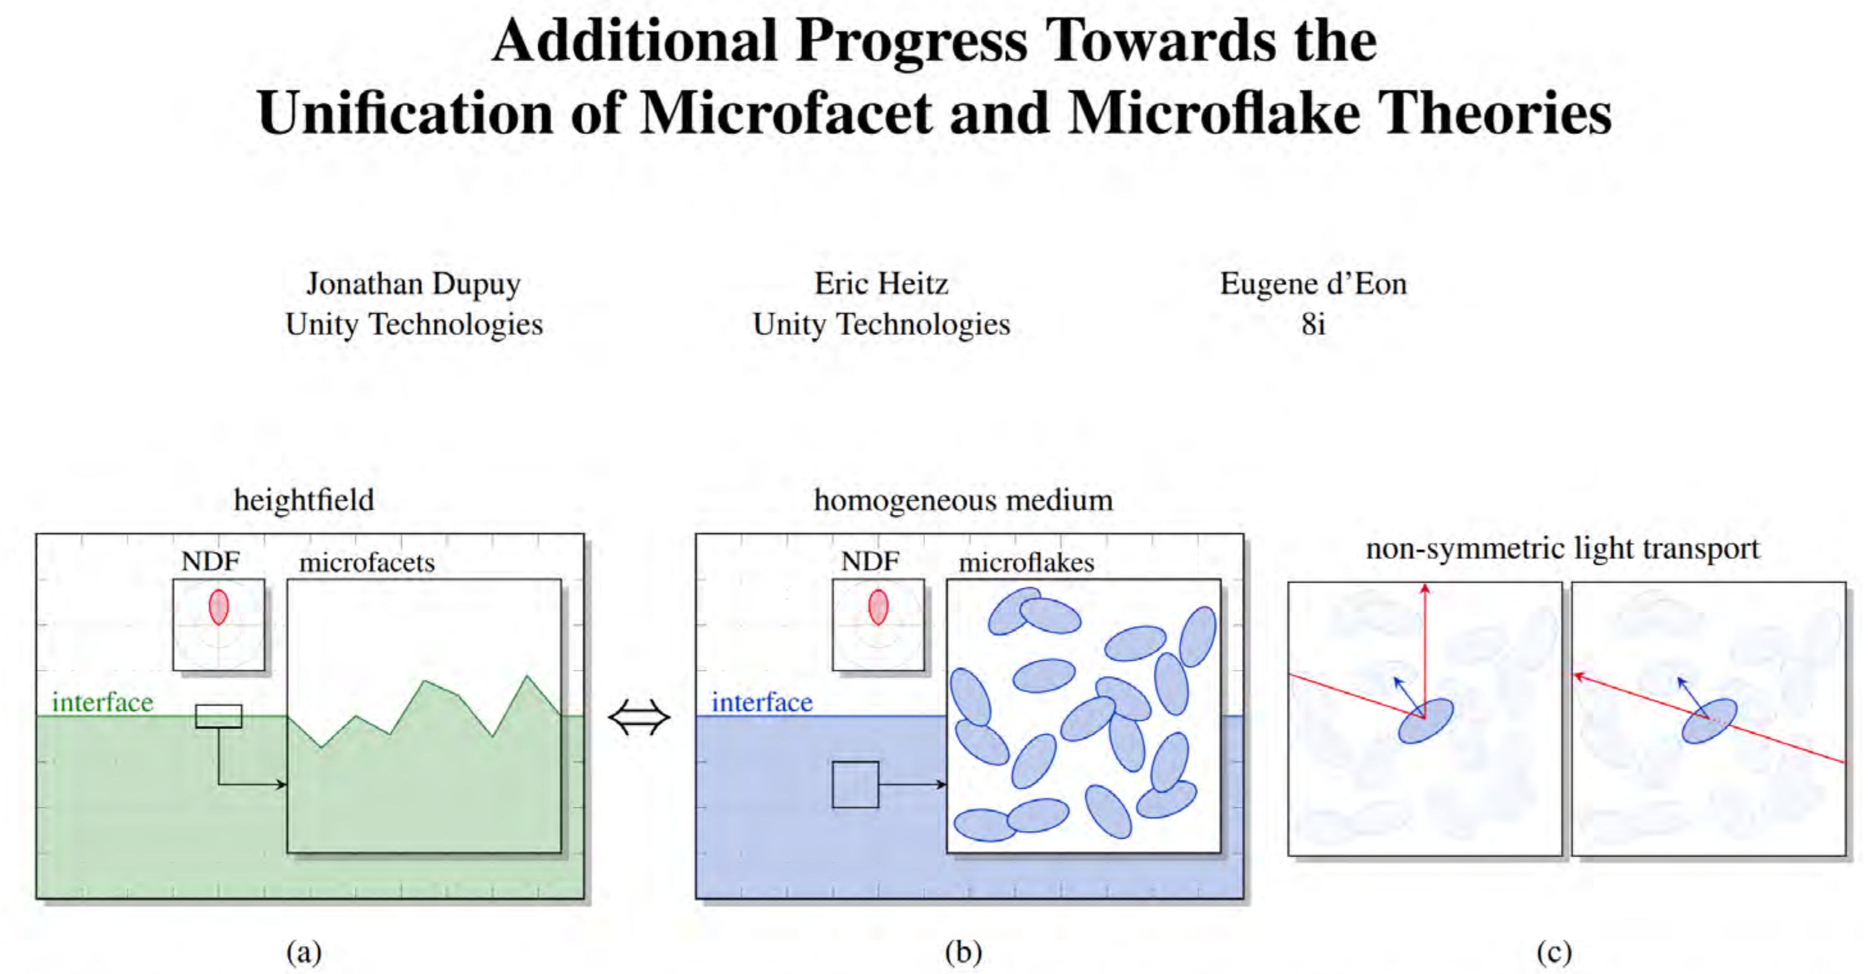

Figure 1: We show that light transport due to (a) a rough microfacet surface is the same as the light transport due to (b) a semi-infinite homogeneous microflake volume consisting of (c) non-symmetric microflakes that are reflective on one side and transparent on the other side, which have the same NDF as (a).

Future Work

- Machine learning + appearance modeling
  - Next talk!
  - Neural BTF Compression and Interpolation [Rainer et al. 2019]
  - Unified Neural Encoding of BTFs [Rainer et al. 2020]
  - …

Conclusion

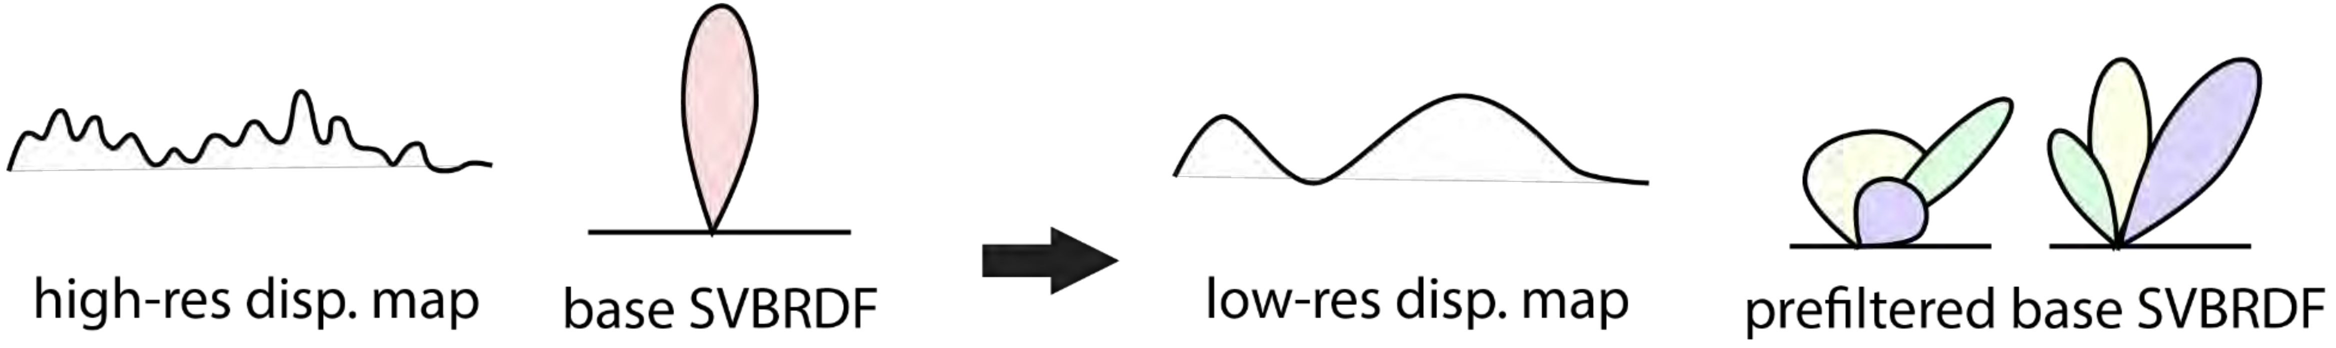

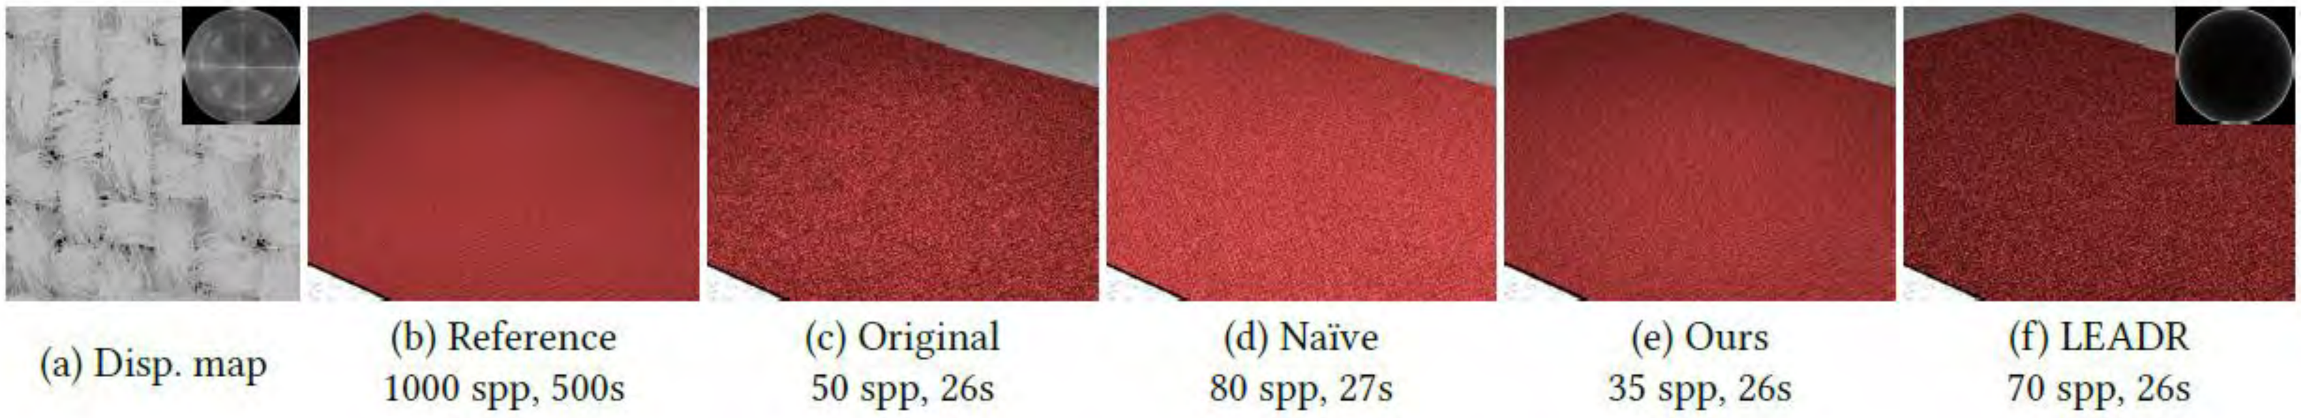

- Anti-aliased
- Accurate
- General surface

Thank you!

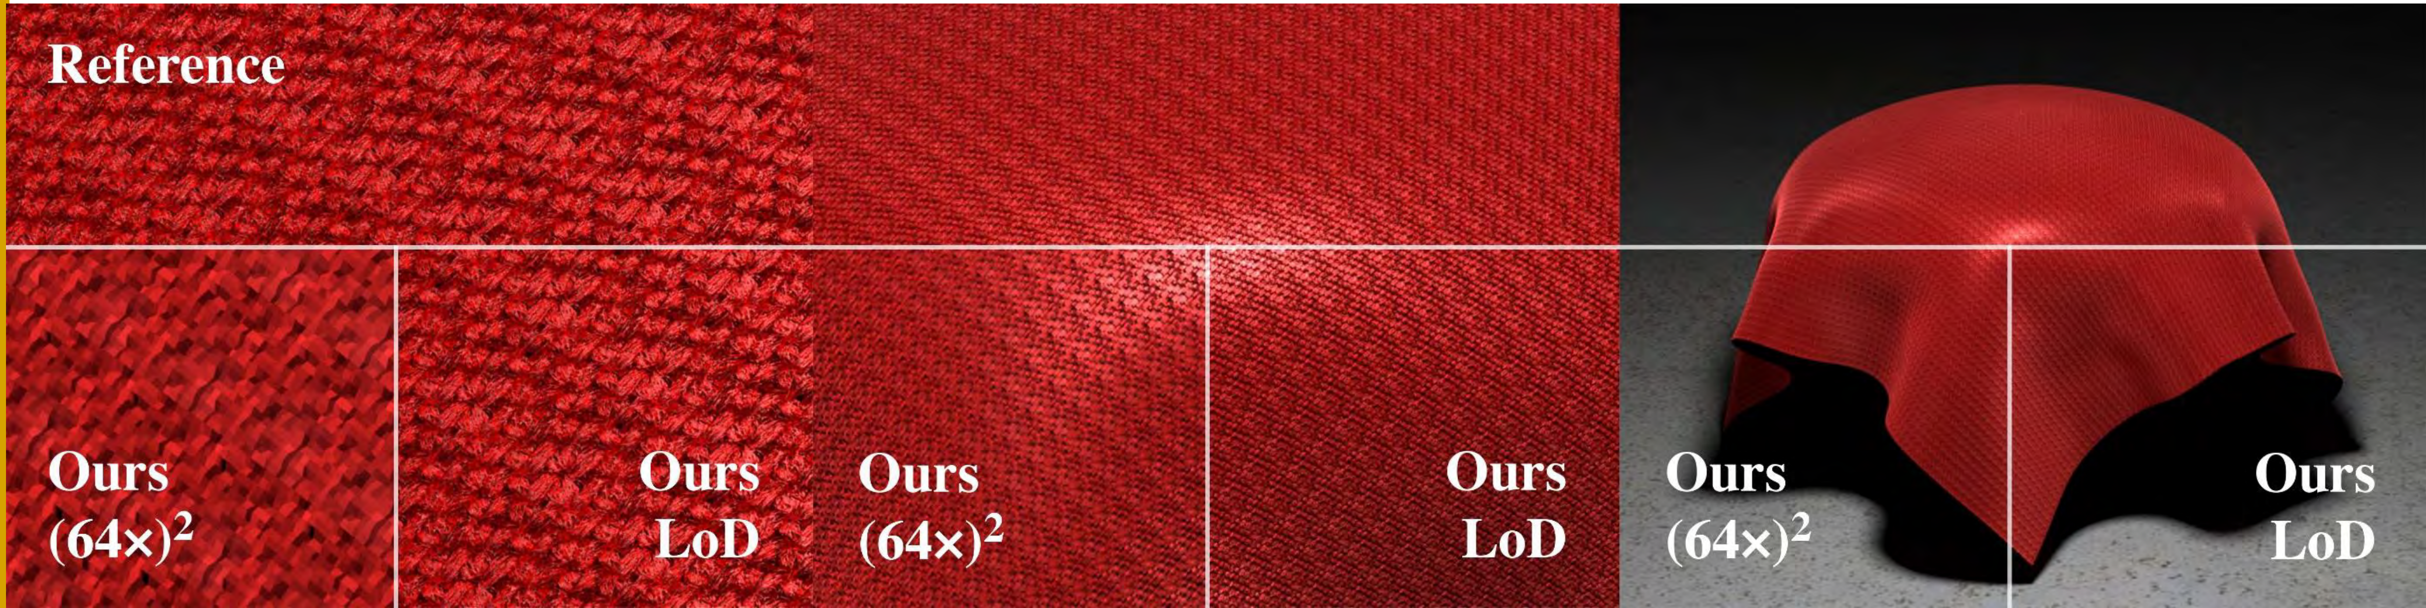
Reference
Ours (64×)²
Ours LoD
Ours (64×)²
Ours LoD
Ours (64×)²
Ours LoD