DOI: 10.1111/cgf.70348

EUROGRAPHICS 2026 / B. Masia and J. Thies
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 45 (2026), Number 2

Neural Local Inter-reflection Modeling for Garment Fold Rendering

Jooeun Son[1] Nuri Ryu[1] Gyoonseo Kim[1] Joo Ho Lee[2] Seungyong Lee[1†]

[1]POSTECH, South Korea
[2]Sogang University, South Korea

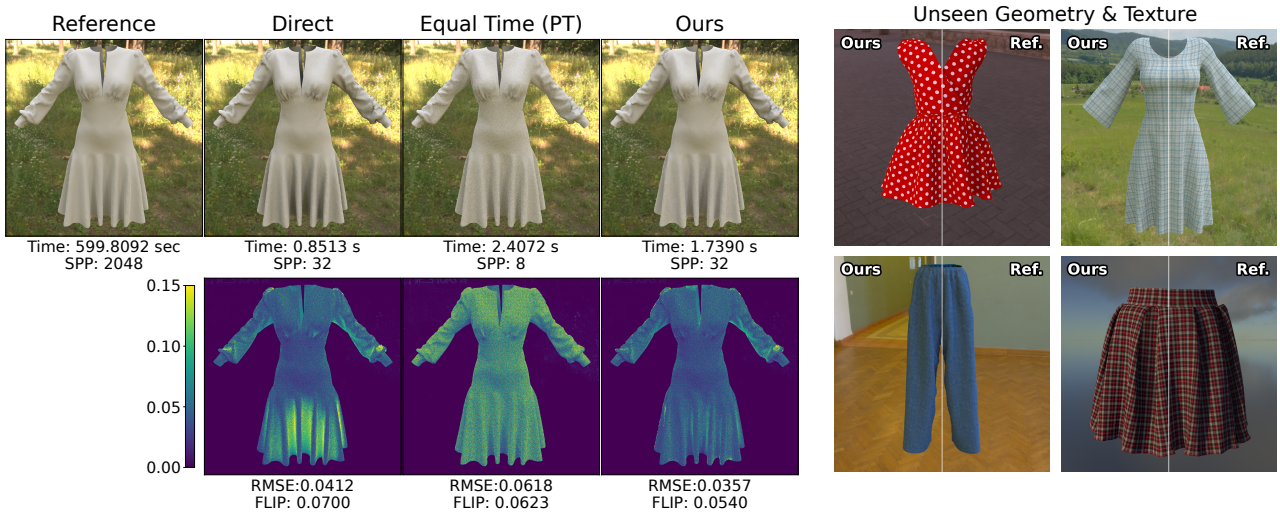

Figure 1: *Performance comparison (left): Our neural local inter-reflection model yields lower error (RMSE, FLIP [ANA21]) compared to the direct illumination integrator ("Direct") and path tracer ("PT") in an equal time budget. Generalization to unseen garments (right): By learning inter-reflection effects conditioned on local geometric features and material properties, our model synthesizes high-fidelity multi-bounce effects for garments with unseen geometry and textures.*

Abstract

Realistic garment rendering requires simulating complex multi-bounce light paths within intricate fold geometries. In these regions, conventional path tracing is computationally expensive as light becomes trapped, necessitating high bounce counts for convergence. We observe that these local inter-reflections are highly localized and exhibit radiance patterns strongly correlated with local fold shapes. Based on these insights, we propose a neural local inter-reflection model that factorizes light transport into overall intensity and directional distribution. By learning the relationship between incident light, material properties, and a novel fold shape descriptor, our model approximates multi-bounce effects using a compact Spherical Harmonics representation. Our approach demonstrates strong generalization to unseen geometries and various fabric textures without retraining. Compared to full path tracing, our method significantly reduces rendering time while preserving high visual fidelity.

CCS Concepts

• ***Computing methodologies*** → ***Rendering;***

1. Introduction

Realistic and efficient garment rendering is crucial in applications such as digital fashion and virtual try-ons. Achieving photorealism involves accurately representing cloth across its inherent scales, from the microscopic fiber, ply, and yarn structure to macroscopic fold and wrinkle geometry. For micro-scale appearance, curve-based or volumetric models [ZJMB11; ZWDR16; KSZ*15] can capture intricate details, but their high computational and spatial costs often limit practical use without acceleration or aggrega-

† Corresponding author.

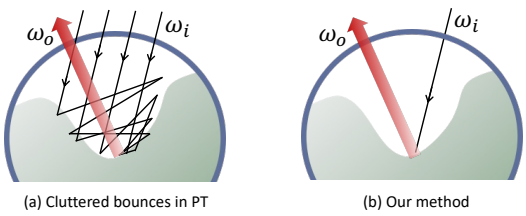

Figure 2: *Motivation. (a) We observe that multi-bounce light inter-reflections within garment folds are highly localized to small neighborhoods. (b) Our neural local inter-reflection model approximates these effects by condensing the internal scattering history into a single-pass shading operation, enabling efficient rendering of such complex fold structures.*

tion [MGZJ20b; SM24]. For better efficiency, recent methods often employ surface-based representations paired with sophisticated micro-, meso-scale appearance models, such as specialized Bidirectional Scattering Distribution Functions (BSDFs) [WJHY22; JWH*22; ZJA*23; ZHB*24; CWW24]. This combination balances detail with efficiency by effectively conveying intricate material properties on simpler geometry.

However, while these BSDF models define micro- and meso-scale light-material interactions, accounting for macro-scale interactions within folds and wrinkles of garments still necessitates global illumination algorithms like path tracing, which incurs a significant bottleneck. In such path tracing simulations, the trapping of light paths within the concave structures of folds leads to numerous inter-reflections, necessitating a high number of light bounce simulations for convergence, increasing rendering times. Existing cloth-specific accelerations are mostly for geometric level-of-detail or mesostructure appearance [WY17; WZYR19; ZHB*24], and do not inherently resolve this challenge of *macroscopic* multi-bounce light transport. Therefore, a core challenge remains to efficiently capture the detailed appearance produced by local inter-reflections within concave fold structures.

In this work, we address the challenge of costly multi-bounce simulation by proposing an efficient neural model that represents local inter-reflections within garment folds as a learnable *Bidirectional Local Inter-reflection Distribution Function (BLIDF)*. Our approach is grounded in two key observations regarding light transport in folds, as visualized in Fig. 2: first, multi-bounce events are highly localized within a small neighborhood; and second, the resulting radiance patterns are strongly correlated to the local fold geometry.

To leverage these insights, we introduce a novel fold shape descriptor that encodes the geometric features most salient to local light transport within garment folds. This descriptor, combined with incident light direction and material parameters (diffuse albedo), serves as the input to our factorized neural architecture. Specifically, we decouple the local inter-reflection distribution into an overall intensity scale and a directional distribution represented by Spherical Harmonics (SH) coefficients. This design choice allows the model to approximate the complex cumulative effect of multiple light bounces within a single inference pass, effectively replacing the computationally intensive tracing of high-depth paths.

Furthermore, because our model is trained on position-invariant geometric features rather than absolute coordinates, it demonstrates remarkable generalization to unseen garment geometries. By training on a diverse multi-material dataset, the network also learns to account for varying diffuse reflectances at runtime, enabling seamless application to new garment models and varied textures without any per-scene retraining or fine-tuning.

In summary, our main contributions are as follows:

- *Factorized neural inter-reflection model:* We propose a neural local inter-reflection model that approximates complex multi-bounce light transport within folds by decoupling radiance into intensity and SH-based directional distributions.
- *Generalizable Fold Shape Descriptor:* We introduce a compact fold shape descriptor encoding position-invariant geometric information to facilitate generalization across unseen garments.
- *Efficient and Practical Integration:* We demonstrate that our BLIDF accelerates convergence compared to path tracing, while enabling runtime material and geometry changes without per-scene retraining or fine-tuning.

2. Related Work

2.1. Cloth Appearance Models

High-fidelity cloth rendering requires balancing structural detail with efficiency [CLA19]. For structural detail, fine-scale models utilize micro-CT data [SKZ11; ZJMB11] or explicit fiber curves [ZLB16; KSZ*15] to represent 3D fabric composition. To manage these complex structures, researchers have proposed exemplar-based volumetric synthesis [ZJMB12], downsampled scattering parameters [ZWDR16], and recently, automatic reconstruction of fiber-level woven structures from single images [WKZ*25].

Despite their fidelity, these models are prohibitively expensive, requiring massive memory and suffering from slow path-tracing convergence. To mitigate these costs, prior methods have proposed aggregating fiber behavior into coarser structures [MGJZ21; MGZJ20a] or utilizing neural networks [SM24]. Other approaches focus on performance gains through efficient procedural fiber generation on the GPU [WY17] or precomputation-based algorithms for multi-bounce effects [KWN*17].

In contrast, surface-based models [ZJA*23; ZHB*24; KZYM25] offer greater efficiency by representing cloth as a macroscopic surface [WAT92]. To convey fine-scale detail, these methods employ specialized Bidirectional Texture Functions (BTFs) [SSK03; KBD07] or BSDFs [IM12a; SBDJ13]. To more realistically capture the volumetric nature of cloth, modern BSDFs often leverage microflake theory [JAM*10; HDCD15] or the Spongecake model [WJHY22], with recent extensions addressing non-local effects [ZJA*23] and sheen and parallax effects [ZHB*24]. However, while these models are computationally lightweight for local evaluation, they neglect the complex light transport induced by macroscopic fold geometry.

Consequently, even with these efficient material models, path tracing remains computationally demanding. The primary bottleneck stems from macroscopic geometry, specifically folds and wrinkles [LSGV18; CKSV23], rather than micro-scale detail. In

these regions, light becomes trapped within concave structures, leading to extensive inter-reflections. Since the computational cost of simulating this exhaustive light transport far outweights that of local material evaluation, achieving efficiency requires a model that explicitly accounts for the macroscopic fold shape.

2.2. Efficient Inter-reflection Rendering

Various approaches have been proposed to mitigate the prohibitive costs of unbiased path tracing by accelerating or approximating global illumination (GI). However, most general-purpose techniques rely on assumptions that mismatch the specific characteristics of garment fabric.

Classical radiosity methods [GTGB84] assume Lambertian surfaces, conflicting with the specular nature of common fabrics [IM12b]. Similarly, illumination caching methods, such as irradiance [WRC88] or radiance caching [KGPB05], often oversmooth the sharp, view-dependent highlights and shadows prevalent in tight creases. Other approximations, including Virtual Point Lights (VPLs) [Kel97; WFA*05] and screen-space methods [Mit07; RGS09], are typically optimized for low-bounce global effects; they struggle to resolve the high-order scattering events that are characteristic of garment folds. Finally, while Precomputed Radiance Transfer (PRT) [SLS05; PWPB07; IDYN07] effectively models complex visibility, the storage and precomputation demands make it impractical for the diverse and dynamic derformations inherent to clothing.

Our work is conceptually aligned with V-groove microfacet models [XH18; LJJ*18], which provide an analytical framework for multiple scattering within specular cavities. These models, alongside foundational work in rough surface reflectance [CT82; WMLT07], describe how geometry traps light. Our proposed neural local appearance model functions as a data-driven, macroscopic analogue to these theories; it directly predicts the outcome of localized multi-bounce events, offering a targeted and efficient solution tailored to the arbitrary, non-analytical geometries of garment folds.

2.3. Neural Approaches for Complex Light Transport

Neural networks have emerged as powerful tools to accelerate global illumination and model complex scattering phenomena. Prominent techniques such as Neural Radiance Caching (NRC) [MRNK21; DKHD25] learn radiance fields to speed up path tracing. However, these methods require online adaptation, necessitating continuous re-training for dynamic scenes and often exhibiting high variance and error before stabilization. In contrast, our approach utilizes a pre-trained representation that generalizes immediately. By conditioning on local features, our model delivers high-fidelity results on dynamic scenes without the latency or computational overhead of online adaptation.

Domain-specific neural models have also succeeded in bypassing expensive light transport simulations at specific scales. For instance, Vicini et al. [VKJ19] trained a shape-adaptive network to learn Bidirectional Scattering Surface Reflectance Distribution Functions (BSSRDFs) [JMLH01] for translucent materials, effectively replacing internal volumetric scattering with a neural approximation. Similarly, neural approaches have been deployed to aggregate scattering in fur fiber bundles [ZZW*22] and yarns [SM24], predict high-order scattered radiance within dense hair volumes [KJA*23], or represent multi-scale weave patterns [CWW24]. While these methods effectively capture appearance arising from a material's intrinsic micro- or meso-structure, they typically yield local surface models that remain agnostic to the light transport governed by the macroscopic 3D configuration of the folds.

3. Local Fold Inter-reflection Model

This section presents our framework for efficiently modeling multi-bounce light transport within garment folds. We first mathematically derive the Bidirectional Local Inter-reflection Distribution Function (BLIDF) (3.1). We then identify the modeling requirements for its neural representation (3.2) and introduce geometric and material descriptors (3.3) to capture relevant features. Finally, we detail our factorized neural architecture (3.4) that decouples light transport into intensity and directional distributions.

Scope and assumptions. To ensure computational tractability and physical consistency, our framework is built on several key assumptions regarding material properties, geometry and lighting environments. *Material Properties:* We model garments as opaque thin-shells and do not account for light transmission through fabric. We assume the fabric is primarily diffuse; highly specular or complex anisotropic materials such as satin are outside the scope of this work. *Geometric Smoothness:* The macroscopic fold shapes are assumed to be smooth to ensure they are well-approximated by our second-degree tri-variate polynomial descriptors. *Locally Constant Illumination:* We assume the incident radiance field L_i is locally constant across the small spherical neighborhood $\mathcal{N}(x_\mathrm{o})$. This approximation is most suitable for directional or distant emitters, where lighting variation across the neighborhood is negligible.

3.1. Mathematical Derivation of BLIDF

We seek a mathematical model that summarizes the internal scattering history within a local spherical neighborhood $\mathcal{N}(x_\mathrm{o})$ of radius r at surface point x_o. We begin by partitioning the total outgoing radiance $L_\mathrm{o}^{\mathrm{local}}(x_\mathrm{o},\omega_\mathrm{o})$ into local single-scatter and multi-bounce components:

$$L_\mathrm{o}^{\mathrm{local}}(x_\mathrm{o},\omega_\mathrm{o}) = L^{local}_{o,\,single}(x_\mathrm{o},\omega_\mathrm{o}) + L^{local}_{o,\,multi}(x_\mathrm{o},\omega_\mathrm{o}). \quad (1)$$

The *single-scatter* component $L^{local}_{o,single}$ describes light reflecting directly from x_o without prior interaction inside the neighborhood. This is evaluated using the standard reflection equation:

$$\int_\Omega L_i(x_\mathrm{o},\omega_\mathrm{i}) V(x_\mathrm{o},\omega_\mathrm{i}) S_{base}(x_\mathrm{o},\omega_\mathrm{o},\omega_\mathrm{i}) |\cos\theta_i| d\omega_\mathrm{i}, \quad (2)$$

where $L_i(x_\mathrm{o},\omega_\mathrm{i})$ is the incident radiance at x_o from direction ω_i, $V(x_\mathrm{o},\omega_\mathrm{i})$ is the visibility term, and $|\cos\theta_i|$ is the foreshortening factor.

The *multi-bounce* component represents the aggregate radiance from light paths that enter $\mathcal{N}(x_\mathrm{o})$ at any point x_i, undergo one or more internal reflections, and exit at x_o. To formalize this, we draw

(a) Shape descriptors ϕ_S (b) Normal directions (c) Illumination descriptor ϕ_L (d) Material descriptor ϕ_M

Figure 3: *Visualization of descriptors. We visualize the input features and auxiliary information used to condition our neural inter-reflection model. (a) Shape $\phi_S \in \mathbb{R}^9$: The Hill-Valley descriptor $\phi_{hv} \in \mathbb{R}$ captures local concavity and is defined for both front and back surfaces due to the thin-shell nature of cloth. It is complemented by the polynomial descriptor $\phi_{poly} \in \mathbb{R}^8$, which captures structural asymmetries essential for modeling the directional distribution. (b) Normals: The shading normal $\vec{n}$ and neighborhood normal $\vec{n_{\mathcal{N}}}$ provide local and regional geometric context. These are auxiliary inputs required to compute the illumination descriptor. (c) Incident illumination $\phi_L \in \mathbb{R}^2$: The illumination descriptor encodes the local light direction with respect to the shading and neighborhood normals. (d) Material $\phi_M \in \mathbb{R}$: The diffuse albedo dictates energy attenuation across multiple light bounces.*

an analogy from BSSRDFs [JMLH01], which summarize internal volumetric scattering by characterizing the transport between an entry point x_i and exit point x_o. Under this framework, the multi-bounce component is expressed as a double integral over the hemisphere Ω and the neighborhood surface area:

$$\int_{\Omega}\int_{\mathcal{N}(x_o)} L_i(x_i,\omega_i)V(x_i,\omega_i)S_{\text{surf}}(x_o,\omega_o,x_i,\omega_i)|\cos\theta_i|dA_i d\omega_i, \quad (3)$$

where S_{surf} is a scattering kernel representing the cumulative inter-reflection transport.

Directly evaluating the double integral of Eq. (3) is computationally prohibitive as it requires dense sampling of the entry points x_i. To make this tractable, we apply a common approximation used to simplify BSSRDFs into point-based BRDFs [**jensen2001bssrdf**]: we assume the incident radiance field L_i is locally constant ($L_i(x_i,\omega_i) \approx L_i(x_o,\omega_i)$) over the small neighborhood $\mathcal{N}(x_o)$. This allows us to factor L_i out of the area integral:

$$\int_{\Omega} L_i(x_o,\omega_i)\left[\int_{\mathcal{N}(x_o)} V(x_i,\omega_i)S_{\text{surf}}(x_o,\omega_o,x_i,\omega_i)|\cos\theta_i|,dA_i\right]d\omega_i. \quad (4)$$

We define the bracketed area integral as the product of the *Bidirectional Local Inter-reflection Distribution Function (BLIDF)* $\bar{S}(x_o,\omega_o,\omega_i)$ and the *neighborhood visibility* $\bar{V}_{\mathcal{N}}(x_o,\omega_i)$, outlined in Section 4.1. This yields our final formulation:

$$L^{local}_{o,\,multi} \approx \int_{\Omega} L_i(x_o,\omega_i)\bar{V}_{\mathcal{N}}(x_o,\omega_i)\bar{S}(x_o,\omega_o,\omega_i)d\omega_i. \quad (5)$$

We refer to this formulation as the *local rendering equation*, as it maintains the structure of the standard rendering integral while restricting light transport context to the neighborhood $\mathcal{N}(x_o)$. By condensing the spatial scattering history into the BLIDF, we transform a complex path-tracing problem into an efficient single-point shading operation similar to a standard BSDF.

3.2. Modeling Requirements of Neural BLIDF

The BLIDF $\bar{S}(x_o,\omega_o,\omega_i)$ encapsulates the aggregate effect of multi-bounce light paths confined within the local neighborhood $\mathcal{N}(x_o)$. To represent this function neurally, we must account for three primary factors: 1) the local fold geometry $\phi_S(x_o)$, which dictates path trajectories and occlusions; 2) the material properties ϕ_M, which govern energy loss at each surface interaction within the fold; and 3) the incident light direction ω_i, which determines the initial energy distribution. The non-linear coupling between these high-dimensional parameters makes a general closed-form analytic expression for $\bar{S}$ intractable. This motivates the use of a learned neural approximation capable of capturing these intricate relationships.

A naïve neural approach would be to model $\bar{S}$ as a function of the absolute world-space position x_o. However, such a model would overfit to the training geometry and fail to generalize to new garment shapes. To achieve geometry-agnostic generalization, we require the model to operate exclusively on local, position-invariant features. By extracting geometric features relative to the local surface frame at x_o, we ensure that the network learns the relationship between fold shape and radiance patterns rather than memorizing specific coordinates.

We thus formulate our neural BLIDF as:

$$\bar{S}(x_o,\omega_i,\omega_o) \approx S_\theta(\omega_i,\omega_o,\phi_S(x_o),\phi_M(x_o),\phi_L(x_o)), \quad (6)$$

where S_θ is a neural network with learnable parameters and all input descriptors are defined within the local coordinate system at x_o. We detail the specific construction of these descriptors in Section 3.3 and the architecture of the neural module in Section 3.4.

3.3. Geometric and Material Descriptors

To evaluate the BLIDF neurally, we represent th
hood through three position-invariant descriptor
material $\phi_M(x_o)$, and illumination $\phi_L(x_o)$ (Fig. 3)

3.3.1. Shape Descriptors

The geometry within $\mathcal{N}(x_o)$ fundamentally dete
erning how light is occluded and trapped. We
shape descriptor $\phi_S(x_o) = (\phi_{hv}(x_o), \phi_{poly}(x_o))$ t
coarse concavity and fine surface variation.

Hill-Valley descriptor. The hill-valley descriptor ϕ
value that quantifies local concavity, which is a
determining inter-reflection intensity. While ligl
convex or flat regions ("hills"), it becomes trapp
gions ("valleys"), leading to high-magnitude mul
To capture this, we fit a virtual sphere tangent tc
the maximum radius r_{sphere} that avoids surface intersections within
$\mathcal{N}(x_o)$ (Fig. 4a). As can be seen in Fig. 3a (left, "Front side"), valleys are thus mapped to small positive radii, while hills yield large values, clipped to the neighborhood radius r.

Since garment fabrics are thin and two-sided, we pre-compute these radii for both the front and back sides of the mesh. At render time, the descriptor $\phi_{hv}(x_o)$ is dynamically assigned the value corresponding to the side currently undergoing shading.

Polynomial surface descriptor. While $\phi_{hv}(x_o)$ captures the concavity for intensity estimation, complex fold structures require a more detailed representation to model the directional distribution of light. Following [VKJ19], we fit a second-degree tri-variate polynomial to the surface points within $\mathcal{N}(x_o)$ relative to the local coordinate frame. For computational efficiency, we select only the 8 most significant coefficients out of the entire 10 coefficients. These coefficients serve as high-dimensional features capturing structural asymmetries and surface variations (Fig. 3a, right) that a scalar radius cannot describe.

Similar to the hill-valley descriptor, these coefficients must adapt depending on which side of the surface is undergoing shading. Due to the thin-shell nature of garments, the back side is locally antiparallel to the front, allowing the derivation of back side coefficients via a coordinate reflection across the local xy-plane without a costly re-fitting process (Fig. 4b). Algebraically, this reflection is realized by flipping the signs of coefficients for terms with odd powers of z (i.e., the z, xz, yz terms). This ensures geometric consistency across both sides of the garment, with negligible computational overhead.

3.3.2. Material and Light Descriptors

Material descriptor. The material descriptor $\phi_M(x_o)$ is the 1-dimensional diffuse albedo of the surface. We prioritize albedo over other material parameters because it dictates energy attenuation across multiple bounces, thereby dominating the cumulative brightness of light escaping from folds. In contrast, parameters governing micro-scale details (e.g., ply orientation) have comparatively negligible impacts on macroscopic inter-reflection intensity.

Incident illumination descriptor. The illumination descriptor $\phi_L(x_o)$ is a 2D vector representing the incident light direction ω_i

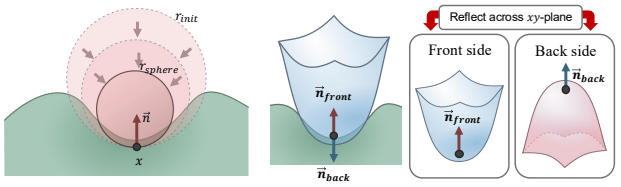

(a) Virtual sphere fitting (b) Symmetry transformation for back side ϕ_{poly}

Figure 4: *Shape descriptors. (a) The Hill-Valley descriptor is computed via an iterative sphere fitting algorithm to find the maximum radius r_{sphere} free from surface intersections. (b) Back side polynomial descriptors are derived via coordinate reflection across the local xy-plane.*

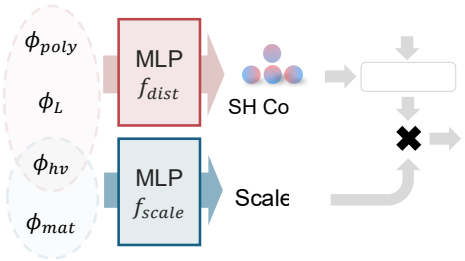

Figure 5: *Factorized network architecture. To enable efficient evaluation of the local rendering equation, we factorize local light transport into intensity f_{scale} and directional distribution f_{dist}. f_{scale} predicts the total energy scaling κ based on the Hill-Valley descriptor ϕ_{hv} and albedo ϕ_M. f_{dist} predicts the estimated directional distribution encoded as SH-coefficients, which are evaluated at ω_o and scaled by κ to yield the final color c.*

(Fig. 3c). It contains the cosines of the angles between ω_i and two distinct vectors: the local shading normal $\vec{n}$ and a *neighborhood normal* $\vec{n_{\mathcal{N}}}$ (Fig. 3b). The neighborhood normal is computed by averaging surface normals within $\mathcal{N}(x_o)$, effectively capturing the low-frequency orientation of the fold. This dual-normal design allows the network to distinguish between local incidence x_o and the global context of the light relative to the fold.

3.4. Factorized Network Architecture

Directly mapping high-dimensional geometric and material descriptors to a complex multi-bounce distribution requires a complex architecture that is both computationally expensive and difficult to train. We therefore adopt a factorized architecture that decomposes the inter-reflection process into two physically intuitive components: overall *intensity* and *directional distribution*. Our model S_θ implements this architecture using two compact sub-networks: an intensity network f_{scale} and a distribution network f_{dist}.

Intensity network f_{scale}. The intensity network f_{scale}, predicts a scalar multiplier κ representing the total energy reflected after multiple internal bounces. While the diffuse albedo $\phi_M(x_o)$ defines the energy reflected from a single interaction, the cumulative intensity in a fold arises from a complex history of light paths dictated by the local geometry. We therefore condition the network on both the albedo $\phi_M(x_o)$ and the Hill-Valley descriptor to estimate the aggre-

gate inter-reflection magnitude:

$$\kappa = f_{\text{scale}}(\phi_{\text{M}}(x_{\text{o}}), \phi_{\text{hv}}(x_{\text{o}})). \tag{7}$$

Distribution network f_{dist}. The distribution network f_{dist} models the directional profile of the outgoing radiance. It is conditioned on the incident light descriptor $\phi_{\text{L}}(x_{\text{o}})$ and the shape descriptor $\phi_{\text{S}}(x_{\text{o}})$. The network outputs a set of Spherical Harmonics (SH) coefficients ϕ'_{sh}:

$$\phi'_{sh} = f_{\text{dist}}(\phi_{\text{S}}(x_{\text{o}}), \phi_{\text{L}}(x_{\text{o}})). \tag{8}$$

The final color c for direction ω_{o} is obtained by evaluating the SH basis at ω_{o} and modulating the result by the intensity scale κ:

$$c = \kappa * \text{sh\_eval}(\phi'_{sh}). \tag{9}$$

4. Integrating Local Scattering into Differentiable Rendering

While BLIDF defines internal scattering within a local neighborhood, the final radiance is determined by coupling $\bar{S}$ with incident illumination and local visibility within the rendering integral, as given in Eq. (5). Bridging the gap between a local model to a complete light transport solution requires a tractable model for neighborhood occlusion and a differentiable training strategy. This section details the integration of BLIDF into a full differentiable rendering pipeline and the optimization of its parameters.

4.1. Modeling Neighborhood Visibility

To evaluate the local rendering integral in Eq. (5), we must account for the neighborhood visibility $\bar{V}_{\mathcal{N}}(x_{\text{o}}, \omega_{\text{i}})$, representing the average visibility within $\mathcal{N}(x_{\text{o}})$ for light incident from direction ω_{i}. Explicitly computing this term via dense ray-casting during training is prohibitively expensive.

Inspired by methods for modeling meso-scale fabric shadowing [ZJA*23], we approximate this visibility using Anisotropic Spherical Gaussians (ASGs) [XSD*13]. While the previous work [ZJA*23] uses ASGs to capture fine-grained visibility in fabric, we observe that macro-scale fold visibility is similarly directional and smooth, making it well-suited for a Gaussian representation. We decompose the visibility into the maximum of a direct ray-occlusion test $V(x_{\text{o}}, \omega_{\text{i}})$ and a learned residual component $\bar{V}_{res}(x_{\text{o}}, \omega_{\text{i}})$ parameterized using ASGs:

$$\bar{V}(x_{\text{o}}, \omega_{\text{i}}) = \max(V(x_{\text{o}}, \omega_{\text{i}}), \bar{V}_{res}(x_{\text{o}}, \omega_{\text{i}})), \tag{10}$$

$$\bar{V}_{res}(x_{\text{o}}, \omega_{\text{i}}) = \text{ASG}(\omega_{\text{i}} | \eta_{\text{ASG}}(x_{\text{o}})). \tag{11}$$

By taking the maximum, we ensure high-frequency occlusions are handled by the direct ray test $V(x_{\text{o}}, \omega_{\text{i}})$, while ASGs capture the low-frequency regional context. The ASG parameters $\eta_{\text{ASG}}(x_{\text{o}})$ are jointly optimized with the BLIDF parameters θ, ensuring the visibility and scattering model $\bar{S}$ are mutually consistent.

Because the parameters η_{ASG} are scene-specific, they are not readily available for garments unseen during training. To ensure generalization, we utilize a heuristic fallback for inference on novel geometries: we approximate $\bar{V}_{\mathcal{N}}$ by performing direct visibility tests on surface samples within the neighborhood, obtained by casting probe rays around x_{o} in a manner similar to BSSRDF importance sampling [KKCF13]. This dual approach allows for high-fidelity joint learning during training while maintaining robust performance during general inference. Further details on the ASG formulation and the heuristic sampling algorithm are provided in the supplementary material.

4.2. Differentiable Training Framework

We supervise BLIDF using an image-based differentiable rendering strategy. Unlike point-based supervision, which necessitates exhaustive simulation across various light and view directions for every surface point, training on rendered images provides a spatially dense distribution of geometric features. This approach allows the model to more efficiently learn the mapping between local shape descriptors and the resulting inter-reflection field.

Local multi-bounce supervision. To isolate the inter-reflection component, we created a custom integrator LocalPT, a modified path tracer that simulates light transport confined within the neighborhood $\mathcal{N}(x_{\text{o}})$. When a ray first intersects the garment at x_{o}, we trace subsequent light paths but terminate any segment that exits the neighborhood. The multi-bounce ground truth is computed by subtracting the single-bounce contribution from the LocalPT's result.

Optimization. We optimize the network parameters θ and visibility parameters $\eta_{\text{ASG}}(x_{\text{o}})$ jointly by minimizing the L_2 distance between the multi-bounce component rendered via our local rendering equation and the LocalPT ground truth. This allows gradients to flow from image pixels to both network parameters θ and visibility parameters η_{ASG}.

5. Experimental Results

5.1. Implementation Details

Network architecture. The factorized networks, f_{scale} and f_{dist}, are lightweight MLPs featuring 4 fully connected layers with widths of 16 and 8, respectively. We employ ReLU activations, except for the final layer which uses a Tanh activation. These networks are integrated into the Mitsuba 3 renderer [JSR*22] as a custom BSDF plugin. We represent the directional distribution using 16 SH coefficients (up to the 3rd order). For visibility, we use 2 ASG lobes during training and a 4-sample probe-ray heuristic during inference.

Training configuration. We train on decorative garments (skirts and dresses) from the ClothesNet dataset [ZZL*23], with mesh resolutions ranging between 10K and 20K vertices. For each garment, we generate 300 monochrome images under randomly oriented directional lighting, camera position, and material parameters, using our custom LocalPT integrator. Per-vertex shape descriptors are also pre-computed for each garment, a process that takes roughly $2 \sim 4$ minutes. We provide details on the fitting algorithm in the supplementary. The neighborhood radius r used for both descriptor computation and LocalPT rendering is selected to encapsulate a single fold. Through our experiments, we found that a constant radius of 0.15 for skirts and 0.10 for dresses provides robust performance. For the material model, we use the surface-based cloth

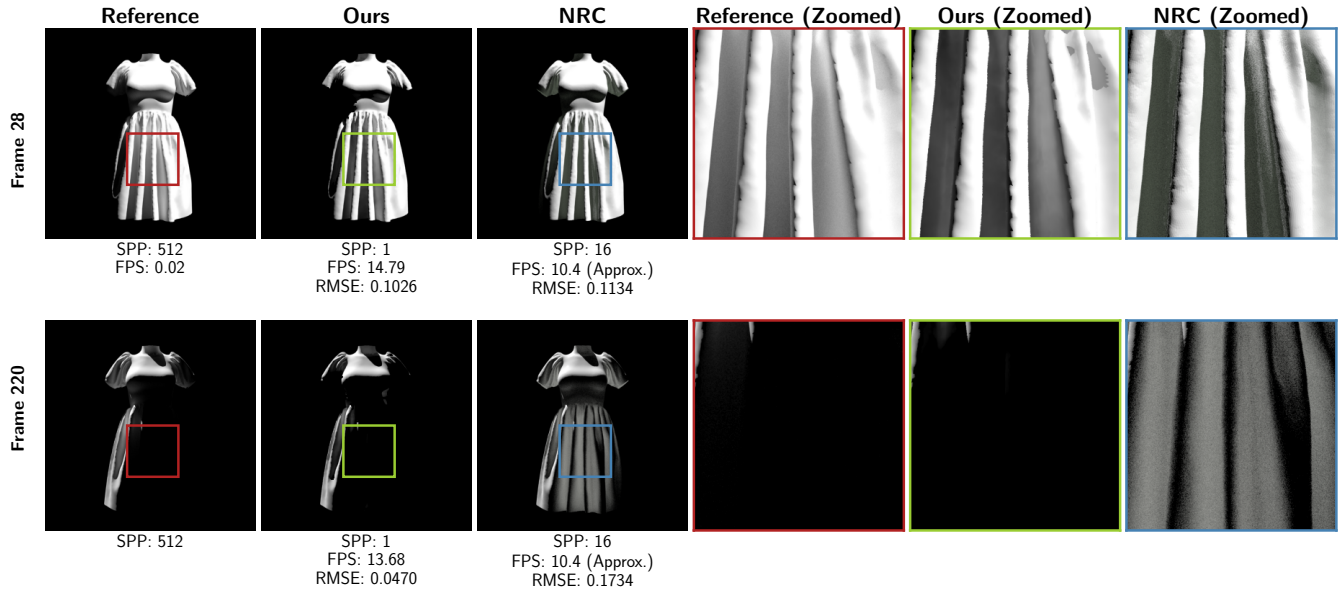

Figure 6: *Comparison with Neural Radiance Cache (NRC). We compare the rendering quality and performance for dynamic scenes by varying the directional light's azimuthal angle. At* 2048×2048 *resolution, our method yields temporally smooth and plausible results at 1 SPP by approximating aggregate multi-path effects in a single pass. In contrast, NRC fails to adapt instantaneously, causing inaccurate fold colors or light leaks. We report average FPS for the animation; note the device differences make FPS comparisons approximate.*

BSDF [ZMA*23], except for comparisons in Fig. 6 where a Lambertian model is used for a fair comparison. During training, we disable rendering of the back side due to its more complex occlusion patterns than the front side, which can harm training stability.

We optimize the network and ASG parameters jointly for 3K epochs using the Adam optimizer [Kin14]. We set a learning rate of 4×10^{-3} for the network and 1×10^{-2} for the ASG parameters. The total training time per garment is under 30 minutes.

Runtime and inference. At inference time, we support back side rendering by selecting shape descriptors corresponding to the side of the current shading frame, as mentioned in Section 3.3. Additionally, we enable recursive tracing (e.g., up to depth 3) by leveraging the single-scattering BSDF's sample function for importance sampling of outgoing directions. Whenever a path continues after evaluating our local model, we offset the ray origin to the boundary of the neighborhood $\mathcal{N}(x_o)$ to prevent redundant computation of local effects that are already captured by BLIDF.

The system is implemented using CUDA and OptiX for hardware-accelered ray tracing on the GPU. All experiments were conducted on a NVIDIA RTX 4090 GPU and an Intel i9-10900K CPU. Our model adds approximately a $0.5\times$ computational overhead compared to the BRDF alone, making it suitable for practical rendering pipelines.

5.2. Comparisons

Comparison with a generalized global illumination method. We compare our method against NRC [MRNK21], a state-of-the-art neural method for global illumination designed to handle dynamic scenes in real time. The fundamental distinction lies in the approach to generalization: while NRC utilizes rapid online adaptation (i.e., continuous retraining on the dynamic scene), our method relies on pre-training with local, position-invariant features.

This distinction leads to a fundamental trade-off illustrated in Fig. 6 under dynamic lighting. NRC's continuous optimization can introduce temporal artifacts, such as flickering or subtle oscillations, as the network adapts to changing illumination. Conversely, our pre-trained model provides perfectly stable results across frames because its weights remain fixed at runtime. As shown in Fig. 6, our method renders dynamic effects consistently even at 1 SPP, avoiding the temporal instability inherent in online training. Due to Windows OS dependency in the official NRC implementation, this comparison was performed on a laptop featuring NVIDIA RTX 4050 GPU and an Intel Ultra 7 155h CPU.

Comparison under environment map lighting. Fig. 7 demonstrates the performance of our method under environment map lighting, comparing rendering results from a direct integrator (i.e., accounting for direct illumination only), our BLIDF, and a full path tracer (PT). Notably, despite being trained solely on directional lighting, our model is compatible with complex environment illumination. By efficiently approximating local multi-bounce effects, our method can utilize more samples per pixel (SPP) than PT within an equal time budget, resulting in lower variance and RMSE. The quantitative results presented in Fig. 8 demonstrate that our method provides a compelling trade-off between quality and efficiency, achieving lower Root-Mean-Square Error (RMSE) than direct illumination while being computationally more efficient than full path tracing.

Figure 7: *Rendering with environment map lighting. We compare our method to direct illumination ("Direct") and full path tracing ("PT"). Our BLIDF, trained only on directional lights, is compatible with environment map lighting and accurately captures the intensity of folds of various scales. By efficiently approximating local multi-bounce effects, our method can afford a higher SPP within an equal time budget compared to PT, resulting in lower variance and RMSE. We provide quantitative analysis in Fig. 8. For this experiment, a separate network is trained for each garment to utilize the ASG-based visibility. (Image resolution:* 2048×2048*)*

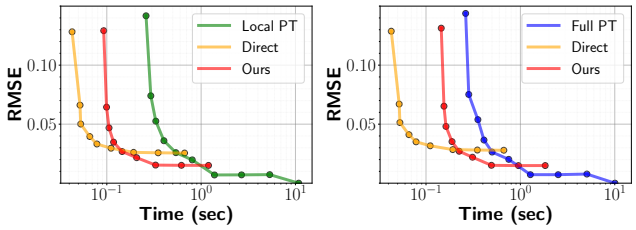

Figure 8: *Quantitative evaluation of environment map lighting. Our method achieves a lower error than direct rendering for a given time budget and a faster result than path tracing for the same SPP. The plots compare errors and times on a* 512×512 *scene with environmental lighting at varying SPPs. The path depth of our method was set to 2 for the comparison with* LocalPT *and 3 for the comparison with* PT.

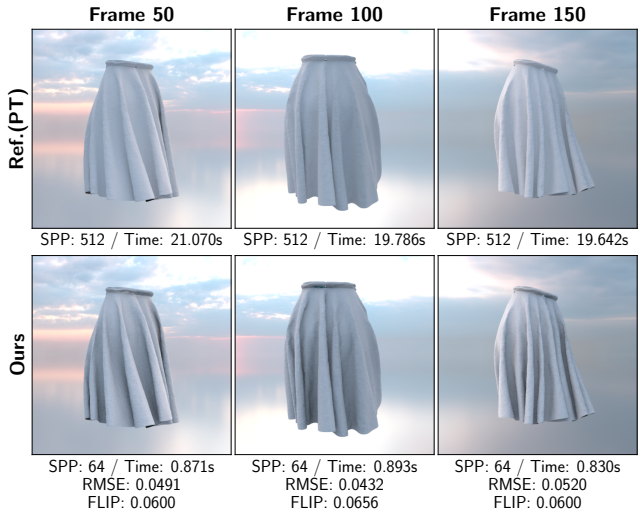

Figure 9: *Garment animation with precomputed descriptors. We apply our method to deformed garments, by re-using the shape descriptors of the initial frame for all frames. Our method effectively captures the dynamic behavior of garment folds by leveraging the relative stability of their fundamental shapes over time. (Image resolution:* 1024×1024*)*

5.3. Analysis and Applications

We demonstrate our method's performance on various applications such as generalization to unseen geometry, texture control, and animations. In the supplementary, we provide detailed analysis on network architecture, shape descriptors, visibility method, and the effect of intrinsic fold shape.

Rendering unseen geometries. A significant advantage of our approach is its ability to generalize to unseen garment geometries without retraining. Because the network is conditioned exclusively on local, position-invariant geometric features rather than absolute world-space coordinates, it can effectively recognize fold structures across disparate garment types. Fig. 11 demonstrates rendering results for various new garment configurations using a model trained only on a single dress. For this experiment, we utilize the heuristic visibility fallback described in Section 4.1.

Figure 10: *Effect of diffuse albedo parameter* ρ*. We analyze the rendering accuracy across different albedo values. The bottom row ("Only High-Order") shows high-order components, computed by subtracting the direct illumination component from the path-traced full image. Our network successfully handles the gradual intensity increase with higher albedo. It also ensures that the intensity values of the high-order components in convex regions stay near zero, appearing as black regions in the hills of the skirt folds.*

Colored and textured garments. Our framework effectively generalizes to garments with complex RGB albedo textures, as shown in Fig. 12. A naïve implementation would require three network evaluations per light-surface interaction to handle each color channel, tripling the computational overhead. To maintain efficiency, we employ a stochastic channel sampling strategy. At each interaction, we randomly evaluate a single color channel and scale the resulting radiance by the reciprocal of the sampling probability (i.e., 3). This ensures the model remains robust across diverse texture patterns while preserving the high rendering efficiency of our original formulation.

Animation rendering. In Fig. 9, we demonstrate that our local inter-reflection model can efficiently render dynamic garment animations. We utilize a cloth simulation software [CLO25] to generate a skirt geometry sequence when the underlying human model performs a catwalk motion. To maximize performance, we compute the per-vertex shape descriptors only for the initial frame. While the exact local geometry of each point evolves during the sequence, we leverage the observation that the fundamental characteristics of garment folds remain relatively stable throughout the animation, even under large deformations. As shown in Fig. 9 and the video in the supplementary material, our method produces realistic multi-bounce effects and maintains high temporal stability without flickering artifacts, demonstrating its practicality.

Effect of material parameter ρ. In Fig. 10, we analyze the sensitivity of our network to variations in the material property by evaluating the model across a range of diffuse albedo values ($\rho \in [0.25, 1.0]$). As the albedo ρ increases, the network correctly predicts an increase in multi-bounce intensity within the folds. Crucially, the model maintains near-zero intensity values for convex regions, demonstrating that the learned BLIDF successfully enforces geometric constraints regardless of the material parameter.

Figure 11: *Generalization to unseen garments. Our model, trained solely on the Dress item in Fig. 7, generalizes well to render plausible results for entirely new garment types without retraining. This is achieved by learning from position-invariant features and utilizing a probe-ray heuristic for visibility on novel geometries. Renders were produced at* 1024×1024 *resolution with a constant albedo of* 0.9*. Our method and the direct-only baseline use 128 samples-per-pixel (SPP), while the reference was generated using LocalPT with 2048 SPP.*

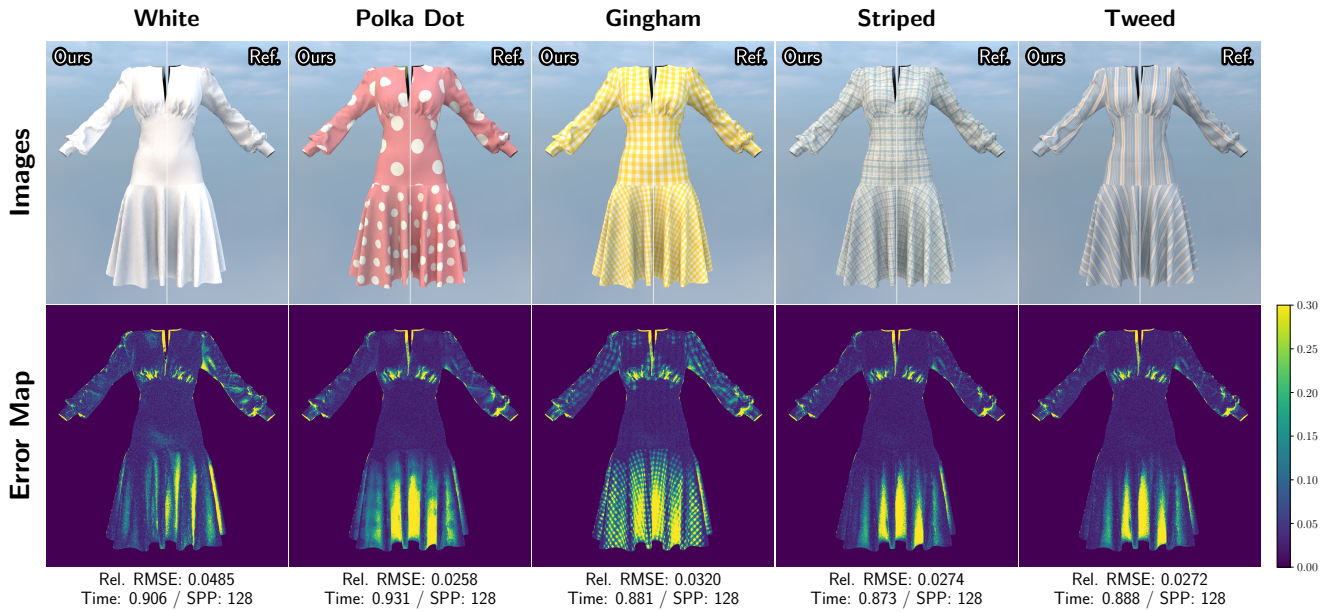

Figure 12: *Rendering garments of various textures. By learning material-dependent effects, our model accurately renders various albedo textures despite being trained on monochromatic fabrics. We compare our results to path-traced references ("Ref.") and report the relative RMSE. By utilizing stochastic RGB channel sampling, our method maintains high efficiency and visual realism, demonstrating robust performance across diverse texture patterns.*

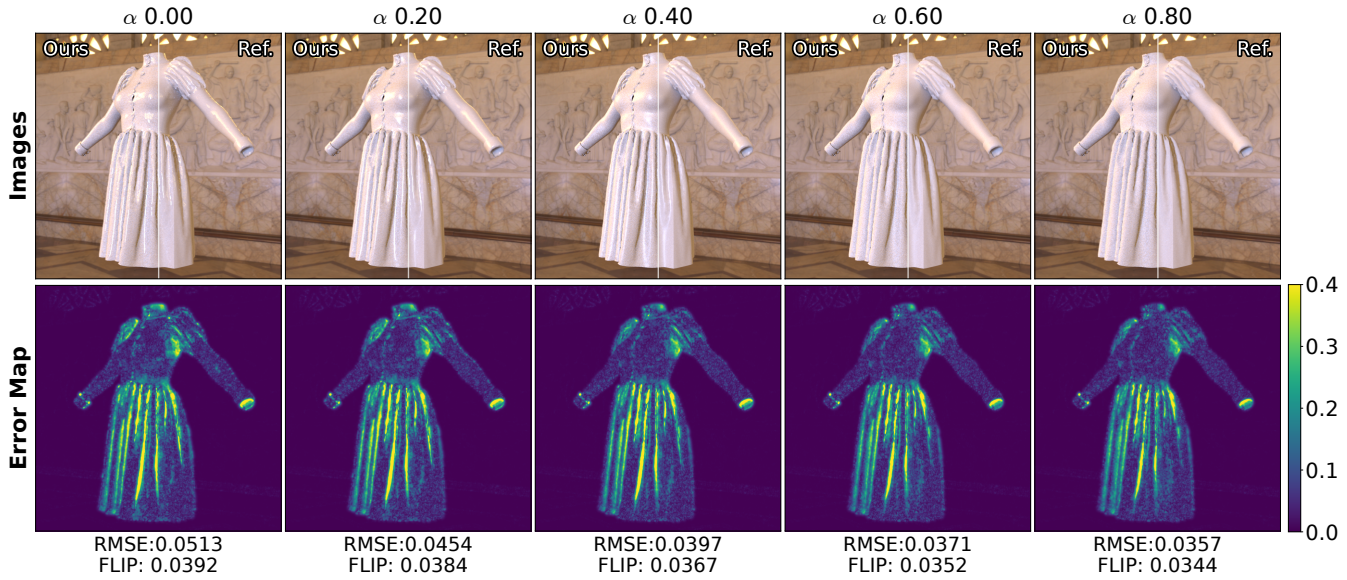

Figure 13: *Supporting roughness control for specular materials. We can extend our neural model to support specular materials by adding specular reflectance and the roughness parameter* α *to the material descriptor* ϕ_M*. (top) Qualitative comparison between our method and the reference ground truth across a range of roughness values (*$\alpha \in [0.0, 0.8]$*). (bottom) Corresponding FLIP error maps. Our model accurately reproduces the characteristic broadening of inter-reflection profiles and diminished specular highlights as the roughness* α *increases. Quantitatively, both RMSE and FLIP errors exhibit a downward trend as* α *increases, demonstrating that our SH-approximation is particularly effective at representing lower-frequency inter-reflection distributions inherent in more diffuse materials. The SPP is 128 and 1024 for ours and the reference, respectively.*

5.4. Extending to Specular Materials

In Fig. 13, we extend our framework to support basic specular materials by expanding the material descriptor ϕ_M to include the specular reflectance and roughness parameters. We randomly sampled from all three material parameters for training, and evaluated the renderings for various roughness values while maintaining the diffuse and specular reflectance at 0.7 and 1.0, respectively. Our model faithfully reproduces the characteristic appearance of glossy fabrics across varying roughness scales. As the roughness increases, the garments exhibit a broader, softer appearance with diminished sharp specular highlights. Then, the multi-bounce distribution becomes lower frequency, allowing for the RMSE and FLIP errors to decrease. While our current model effectively handles a majority of common garment fabrics, we observe that highly view-dependent materials (e.g., polished leather, silk) introduce high-frequency features that remain challenging for low-order SH approximations. Further enhancing the support for highly specular materials remains interesting future work.

6. Conclusions

We presented BLIDF, a neural framework that accelerates the rendering of inter-reflections in garment folds. By adapting the BSSRDF formulation to local fold geometry, we condense complex internal scattering into a compact local distribution function. Our method relies on a novel hybrid shape descriptor ϕ_S, which enables generalization across diverse fabric configurations. We employ a factorized network architecture that separates overall intensity from directional distribution, resolving multi-bounce effects in a single evaluation. Integration of BLIDF into standard path tracers achieves significant speedup in convergence compared to the baselines without compromising visual quality.

6.1. Limitations and Future Work

While our framework effectively captures local inter-reflections in garment folds, several limitations remain that offer directions for future research.

Descriptors and structural generalization. Our handcrafted hybrid descriptors may not fully capture highly irregular or asymmetric folding patterns. Future work could explore data-driven descriptors to enhance generalization across such complex configurations. Furthermore, although our model supports back side rendering at inference, we excluded the back side samples during training to avoid instability caused by their complex occlusion patterns. While necessary for optimization stability, this exclusion limits the diversity of learned configurations, which could be addressed in future work. Finally, our method occasionally exhibits artifacts such as localized light leaks or darkening, which can be caused by insufficient mesh resolution compared to the fixed neighborhood radius r. An adaptive neighborhood size could mitigate these artifacts by disentangling fold scale from the descriptor, allowing the model to focus on intrinsic shape features.

Expanding beyond model assumptions. Our current framework is optimized for opaque thin-shells and lacks a transmission component for translucent fabrics (e.g., lace). Extending our framework beyond diffuse and isotropic materials to handle anisotropic fabric models (e.g., satin) remains an open challenge. From an illumination perspective, our model assumes locally constant lighting, which simplifies computation but makes it less suitable for point lights or highly localized emitters. Finally, developing a specialized importance sampling strategy for the BLIDF distribution, rather than relying on the base BSDF, would significantly enhance recursive tracing efficiency in complex scenes.

Expanding to other meso-scale materials. The principles of the BLIDF could be extended to other materials exhibiting meso-scale inter-reflections, such as skin wrinkles or crumpled paper. Adapting the model to these materials would require specialized descriptors that are capable of capturing organic or non-manifold geometric signatures. Future research could establish a unified neural approach to local light transport that spans these diverse materials, providing a general-purpose solution for meso-scale rendering.

Acknowledgments

We thank anonymous reviewers for their valuable feedback. This work was supported by NRF grants (RS-2025-02216257, RS-2024-00451947, RS-2023-00212828) and IITP grants (RS-2022-II220290, RS-2024-00437866) funded by the Korean government (MSIT).

References

[ANA21] ANDERSSON, PONTUS, NILSSON, JIM, and AKENINE-MÖLLER, TOMAS. "Visualizing and Communicating Errors in Rendered Images". *Ray Tracing Gems II*. Ed. by MARRS, ADAM, SHIRLEY, PETER, and WALD, INGO. 2021. Chap. 19, 301–320 1.

[CKSV23] CHEN, ZHEN, KAUFMAN, DANNY, SKOURAS, MÉLINA, and VOUGA, ETIENNE. "Complex Wrinkle Field Evolution". *ACM Transactions on Graphics* 42 (4 Aug. 2023). ISSN: 15577368. DOI: 10.1145/3592397 2.

[CLA19] CASTILLO, CARLOS, LÓPEZ-MORENO, JORGE, and ALIAGA, CARLOS. "Recent advances in fabric appearance reproduction". *Computers & Graphics* 84 (2019), 103–121 2.

[CLO25] CLO VIRTUAL FASHION. *Marvelous Designer*. Software Version 2025.0.133.54297. Seoul, South Korea, 2025. URL: https://www.marvelousdesigner.com/ 9.

[CT82] COOK, ROBERT L and TORRANCE, KENNETH E. "A reflectance model for computer graphics". *ACM Transactions on Graphics (ToG)* 1.1 (1982), 7–24 3.

[CWW24] CHEN, XIANG, WANG, LU, and WANG, BEIBEI. "Real-time Neural Woven Fabric Rendering". Association for Computing Machinery (ACM), July 2024, 1–10. ISBN: 9798400705250. DOI: 10.1145/3641519.3657496 2, 3.

[DKHD25] DEREVIANNYKH, MIKHAIL, KLEPIKOV, DMITRII, HANIKA, JOHANNES, and DACHSBACHER, CARSTEN. "Neural Two-Level Monte Carlo Real-Time Rendering". *Computer Graphics Forum* 44.2 (2025), e70050. DOI: https://doi.org/10.1111/cgf.70050. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.70050. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.70050 3.

[GTGB84] GORAL, CINDY M, TORRANCE, KENNETH E, GREENBERG, DONALD P, and BATTAILE, BENNETT. "Modeling the interaction of light between diffuse surfaces". *ACM SIGGRAPH computer graphics* 18.3 (1984), 213–222 3.

[HDCD15] HEITZ, ERIC, DUPUY, JONATHAN, CRASSIN, CYRIL, and DACHSBACHER, CARSTEN. "The SGGX microflake distribution". *ACM Transactions on Graphics (TOG)* 34.4 (2015), 1–11 2.

[IDYN07] IWASAKI, KEI, DOBASHI, YOSHINORI, YOSHIMOTO, FUJICHI, and NISHITA, TOMOYUKI. "Precomputed radiance transfer for dynamic scenes taking into account light interreflection". *Proceedings of the 18th Eurographics Conference on Rendering Techniques*. 2007, 35–44 3.

[IM12a] IRAWAN, PITI and MARSCHNER, STEVE. "Specular reflection from woven cloth". *ACM Trans. Graph.* 31.1 (Feb. 2012). ISSN: 0730-0301. DOI: 10.1145/2077341.2077352. URL: https://doi.org/10.1145/2077341.2077352 2.

[IM12b] IRAWAN, PITI and MARSCHNER, STEVE. "Specular reflection from woven cloth". *ACM Transactions on Graphics* 31 (1 Jan. 2012). ISSN: 07300301. DOI: 10.1145/2077341.2077352 3.

[JAM*10] JAKOB, WENZEL, ARBREE, ADAM, MOON, JONATHAN T, et al. "A radiative transfer framework for rendering materials with anisotropic structure". *ACM SIGGRAPH 2010 papers*. 2010, 1–13 2.

[JMLH01] JENSEN, HENRIK WANN, MARSCHNER, STEPHEN R., LEVOY, MARC, and HANRAHAN, PAT. "A practical model for subsurface light transport". *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '01. New York, NY, USA: Association for Computing Machinery, 2001, 511–518. ISBN: 158113374X. DOI: 10.1145/383259.383319. URL: https://doi.org/10.1145/383259.383319 3, 4.

[JSR*22] JAKOB, WENZEL, SPEIERER, SÉBASTIEN, ROUSSEL, NICOLAS, et al. *Mitsuba 3 renderer*. Version 3.1.1. https://mitsuba-renderer.org. 2022 6.

[JWH*22] JIN, WENHUA, WANG, BEIBEI, HASAN, MILOS, et al. "Woven Fabric Capture from a Single Photo". *Proceedings - SIGGRAPH Asia 2022 Conference Papers*. Association for Computing Machinery, Inc, Nov. 2022. ISBN: 9781450394703. DOI: 10.1145/3550469.3555380 2.

[KBD07] KAUTZ, JAN, BOULOS, SOLOMON, and DURAND, FRÉDO. "Interactive editing and modeling of bidirectional texture functions". *ACM Trans. Graph.* 26.3 (July 2007), 53–es. ISSN: 0730-0301. DOI: 10.1145/1276377.1276443. URL: https://doi.org/10.1145/1276377.1276443 2.

[Kel97] KELLER, ALEXANDER. "Instant radiosity". *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 1997, 49–56 3.

[KGPB05] KRIVÁNEK, JAROSLAV, GAUTRON, PASCAL, PATTANAIK, SUMANTA, and BOUATOUCH, KADI. "Radiance caching for efficient global illumination computation". *IEEE Transactions on Visualization and Computer Graphics* 11.5 (2005), 550–561 3.

[Kin14] KINGMA, DIEDERIK P. "Adam: A method for stochastic optimization". *arXiv preprint arXiv:1412.6980* (2014) 7.

[KJA*23] KT, AAKASH, JARABO, ADRIAN, ALIAGA, CARLOS, et al. "Accelerating Hair Rendering by Learning High-Order Scattered Radiance". *Computer Graphics Forum* 42.4 (2023), e14895. DOI: https://doi.org/10.1111/cgf.14895. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14895. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14895 3.

[KKCF13] KING, ALAN, KULLA, CHRISTOPHER, CONTY, ALEJANDRO, and FAJARDO, MARCOS. "BSSRDF importance sampling". *ACM SIGGRAPH 2013 Talks*. 2013, 1–1 6.

[KSZ*15] KHUNGURN, PRAMOOK, SCHROEDER, DANIEL, ZHAO, SHUANG, et al. "Matching real fabrics with micro-appearance models". *ACM Transactions on Graphics* 35 (1 Dec. 2015). ISSN: 15577368. DOI: 10.1145/2818648 1, 2.

[KWN*17] KHUNGURN, PRAMOOK, WU, RUNDONG, NOECKEL, JAMES, et al. "Fast rendering of fabric micro-appearance models under directional and spherical gaussian lights". *ACM Transactions on Graphics (TOG)* 36.6 (2017), 1–15 2.

[KZYM25] KHATTAR, APOORV, ZHU, JUNQIU, YAN, LING-QI, and MONTAZERI, ZAHRA. "A Texture-Free Practical Model for Realistic Surface-Based Rendering of Woven Fabrics". *Computer Graphics Forum* 44.1 (2025), e15283. DOI: https://doi.org/10.1111/cgf.15283. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.15283. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.15283 2.

[LJJ*18] LEE, JOO HO, JARABO, ADRIAN, JEON, DANIEL S., et al. "Practical Multiple Scattering for Rough Surfaces". *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2018)* 37.6 (2018), 275:1–12. DOI: 10.1145/3272127.3275016. URL: http://dx.doi.org/10.1145/3272127.3275016 3.

[LSGV18] LI, MINCHEN, SHEFFER, ALLA, GRINSPUN, EITAN, and VINING, NICHOLAS. "Foldsketch: Enriching garments with physically reproducible folds". *ACM Transactions on Graphics (TOG)* 37.4 (2018), 1–13 2.

[MGJZ21] MONTAZERI, ZAHRA, GAMMELMARK, SOREN, JENSEN, HENRIK W., and ZHAO, SHUANG. "A Practical Ply-Based Appearance Modeling for Knitted Fabrics". *Eurographics Symposium on Rendering* (July 2021). URL: http://arxiv.org/abs/2105.02475 2.

[MGZJ20a] MONTAZERI, ZAHRA, GAMMELMARK, SØREN B, ZHAO, SHUANG, and JENSEN, HENRIK WANN. "A practical ply-based appearance model of woven fabrics". *ACM Transactions on Graphics (TOG)* 39.6 (2020), 1–13 2.

[MGZJ20b] MONTAZERI, ZAHRA, GAMMELMARK, SØREN B., ZHAO, SHUANG, and JENSEN, HENRIK WANN. "A practical ply-based appearance model of woven fabrics". *ACM Transactions on Graphics* 39 (6 Nov. 2020). ISSN: 15577368. DOI: 10.1145/3414685.3417777 2.

[Mit07] MITTRING, MARTIN. "Finding next gen: CryEngine 2". *ACM SIGGRAPH 2007 Courses*. SIGGRAPH '07. San Diego, California: Association for Computing Machinery, 2007, 97–121. ISBN: 9781450318235. DOI: 10.1145/1281500.1281671. URL: https://doi.org/10.1145/1281500.1281671 3.

[MRNK21] MÜLLER, THOMAS, ROUSSELLE, FABRICE, NOVÁK, JAN, and KELLER, ALEXANDER. "Real-time neural radiance caching for path tracing". *ACM Transactions on Graphics (TOG)* 40.4 (2021), 1–16 3, 7.

[PWPB07] PAN, MINGHAO, WANG XINGUO LIU, RUI, PENG, QUNSHENG, and BAO, HUJUN. "Precomputed radiance transfer field for rendering interreflections in dynamic scenes". *Computer Graphics Forum*. Vol. 26. 3. Wiley Online Library. 2007, 485–493 3.

[RGS09] RITSCHEL, TOBIAS, GROSCH, THORSTEN, and SEIDEL, HANS-PETER. "Approximating dynamic global illumination in image space". *Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games*. 2009, 75–82 3.

[SBDJ13] SADEGHI, IMAN, BISKER, OLEG, DE DEKEN, JOACHIM, and JENSEN, HENRIK WANN. "A practical microcylinder appearance model for cloth rendering". *ACM Trans. Graph.* 32.2 (Apr. 2013). ISSN: 0730-0301. DOI: 10.1145/2451236.2451240. URL: https://doi.org/10.1145/2451236.2451240 2.

[SKZ11] SCHRODER, KAI, KLEIN, REINHARD, and ZINKE, ARNO. "A volumetric approach to predictive rendering of fabrics". *Computer Graphics Forum*. Vol. 30. 4. Wiley Online Library. 2011, 1277–1286 2.

[SLS05] SLOAN, PETER-PIKE, LUNA, BEN, and SNYDER, JOHN. "Local, deformable precomputed radiance transfer". *ACM Transactions on Graphics (TOG)* 24.3 (2005), 1216–1224 3.

[SM24] SOH, G. Y. and MONTAZERI, Z. "Neural Appearance Model for Cloth Rendering". *Computer Graphics Forum* 43 (4 July 2024). ISSN: 0167-7055. DOI: 10.1111/cgf.15156. URL: https://onlinelibrary.wiley.com/doi/10.1111/cgf.15156 2, 3.

[SSK03] SATTLER, MIRKO, SARLETTE, RALF, and KLEIN, REINHARD. "Efficient and realistic visualization of cloth". *Proceedings of the 14th Eurographics Workshop on Rendering*. EGRW '03. Leuven, Belgium: Eurographics Association, 2003, 167–177. ISBN: 3905673037 2.

[VKJ19] VICINI, DELIO, KOLTUN, VLADLEN, and JAKOB, WENZEL. "A learned shape-adaptive subsurface scattering model". *ACM Transactions on Graphics (TOG)* 38.4 (2019), 1–15 3, 5.

[WAT92] WESTIN, STEPHEN H, ARVO, JAMES R, and TORRANCE, KENNETH E. "Predicting reflectance functions from complex surfaces". *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques*. 1992, 255–264 2.

[WFA*05] WALTER, BRUCE, FERNANDEZ, SEBASTIAN, ARBREE, ADAM, et al. "Lightcuts: a scalable approach to illumination". *ACM SIGGRAPH 2005 Papers*. 2005, 1098–1107 3.

[WJHY22] WANG, BEIBEI, JIN, WENHUA, HAŠAN, MILOŠ, and YAN, LING QI. "SpongeCake: A Layered Microflake Surface Appearance Model". *ACM Transactions on Graphics* 42 (1 Sept. 2022). ISSN: 15577368. DOI: 10.1145/3546940 2.

[WKZ*25] WU, CHENGHAO, KHATTAR, APOORV, ZHU, JUNQIU, et al. "Automatic Reconstruction of Woven Cloth from a Single Close-up Image". *Computer Graphics Forum*. Vol. 44. 7. Wiley Online Library. 2025, e70243 2.

[WMLT07] WALTER, BRUCE, MARSCHNER, STEPHEN R, LI, HONGSONG, and TORRANCE, KENNETH E. "Microfacet Models for Refraction through Rough Surfaces." *Rendering techniques* 2007 (2007), 18th 3.

[WRC88] WARD, GREGORY J., RUBINSTEIN, FRANCIS M., and CLEAR, ROBERT D. "A ray tracing solution for diffuse interreflection". *Proceedings of the 15th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '88. New York, NY, USA: Association for Computing Machinery, 1988, 85–92. ISBN: 0897912756. DOI: 10.1145/54852.378490. URL: https://doi.org/10.1145/54852.378490 3.

[WY17] WU, KUI and YUKSEL, CEM. "Real-time fiber-level cloth rendering". *Proceedings - I3D 2017: 21st ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. Association for Computing Machinery, Inc, Feb. 2017. ISBN: 9781450348867. DOI: 10.1145/3023368.3023372 2.

[WZYR19] WU, LIFAN, ZHAO, SHUANG, YAN, LING QI, and RAMAMOORTHI, RAVI. "Accurate appearance preserving prefiltering for rendering displacement-mapped surfaces". *ACM Transactions on Graphics* 38 (4 July 2019). ISSN: 15577368. DOI: 10.1145/3306346.3322936 2.

[XH18] XIE, FENG and HANRAHAN, PAT. "Multiple scattering from distributions of specular v-grooves". *ACM Trans. Graph.* 37.6 (Dec. 2018). ISSN: 0730-0301. DOI: 10.1145/3272127.3275078. URL: https://doi.org/10.1145/3272127.3275078 3.

[XSD*13] XU, KUN, SUN, WEI-LUN, DONG, ZHAO, et al. "Anisotropic spherical Gaussians". *ACM Trans. Graph.* 32.6 (Nov. 2013). ISSN: 0730-0301. DOI: 10.1145/2508363.2508386. URL: https://doi.org/10.1145/2508363.2508386 6.

[ZHB*24] ZHU, JUNQIU, HERY, CHRISTOPHE, BODE, LUKAS, et al. "A Realistic Multi-scale Surface-based Cloth Appearance Model". *ACM SIGGRAPH 2024 Conference Papers*. SIGGRAPH '24. Denver, CO, USA: Association for Computing Machinery, 2024. ISBN: 9798400705250. DOI: 10.1145/3641519.3657426. URL: https://doi.org/10.1145/3641519.3657426 2.

[ZJA*23] ZHU, JUNQIU, JARABO, ADRIAN, ALIAGA, CARLOS, et al. "A Realistic Surface-based Cloth Rendering Model". *Proceedings - SIGGRAPH 2023 Conference Papers*. Association for Computing Machinery, Inc, July 2023. ISBN: 9798400701597. DOI: 10.1145/3588432.3591554 2, 6.

[ZJMB11] ZHAO, SHUANG, JAKOB, WENZEL, MARSCHNER, STEVE, and BALA, KAVITA. "Building volumetric appearance models of fabric using micro CT imaging". *ACM Transactions on Graphics (TOG)* 30.4 (2011), 1–10 1, 2.

[ZJMB12] ZHAO, SHUANG, JAKOB, WENZEL, MARSCHNER, STEVE, and BALA, KAVITA. "Structure-aware synthesis for predictive woven fabric appearance". *ACM Transactions on Graphics (TOG)* 31.4 (2012), 1–10 2.

[ZLB16] ZHAO, SHUANG, LUAN, FUJUN, and BALA, KAVITA. "Fitting procedural yarn models for realistic cloth rendering". *ACM Transactions on Graphics*. Vol. 35. Association for Computing Machinery, July 2016. DOI: 10.1145/2897824.2925932 2.

[ZMA*23] ZHU, JUNQUI, MONTAZERI, ZAHRA, AUBRY, J, et al. "A Practical and Hierarchical Yarn-based Shading Model for Cloth". *Computer Graphics Forum*. Vol. 42. 4. Wiley Online Library. 2023, e14894 7.

[ZWDR16] ZHAO, SHUANG, WU, LIFAN, DURAND, FRÉDO, and RAMAMOORTHI, RAVI. "Downsampling scattering parameters for rendering anisotropic media". *ACM Transactions on Graphics (TOG)* 35.6 (2016), 1–11 1, 2.

[ZZL*23] ZHOU, BINGYANG, ZHOU, HAOYU, LIANG, TIANHAI, et al. "Clothesnet: An information-rich 3d garment model repository with simulated clothes environment". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023, 20428–20438 6.

[ZZW*22] ZHU, JUNQIU, ZHAO, SIZHE, WANG, LU, et al. "Practical Level-of-Detail Aggregation of Fur Appearance". *ACM Transactions on Graphics* 41 (4 July 2022). ISSN: 15577368. DOI: 10.1145/3528223.3530105 3.