Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Shared Damage-Field Control for Real-Time Appearance Modeling of Severely Damaged Hair

Min-Song Park
Inha University
Jong-Hyun Kim
jonghyunkim@inha.ac.kr

Inha University

Article

Keywords:

Posted Date: November 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11006999/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Shared Damage-Field Control for Real-Time Appearance Modeling of Severely Damaged Hair

Min-Song Park[1] **and Jong-Hyun Kim**[2,*]

[1]College of AI and Convergence (Department of Design Technology), Inha University, Michuhol-gu, Incheon 22212, South Korea
[2]College of AI and Convergence (Department of Artificial Intelligence, Design Technology), Graduate School of Electrical and Computer Engineering, Inha University, Michuhol-gu, Incheon 22212, South Korea
[*]corresponding author, jonghyunkim@inha.ac.kr

ABSTRACT

Severely damaged hair exhibits coupled changes in strand organization, color, highlight structure, and tip integrity. We present a real-time procedural appearance framework that represents these changes through a shared spatial damage field. The field combines root-to-tip position, strand exposure, and deterministic strand-level variation, and drives four response operators for geometric disorder, tip-biased discoloration, reflectance degradation, and strand-end fragmentation. Because the operators share the same local control signal, geometric and optical changes remain spatially correlated while retaining component-specific response rates. The method is implemented entirely in the GPU shader pipeline and requires neither dense auxiliary geometry nor high-resolution damage textures. Experiments on multiple hairstyles and damage levels demonstrate interactive control, predictable scaling with hair complexity, and limited runtime overhead. The method is intended as an artist-controllable appearance model rather than a physical predictor of chemical or biological hair damage.

Introduction

Real-time hair rendering has been developed primarily for healthy or idealized fibers, with emphasis on anisotropic reflection, transmission, and multiple scattering. Severely damaged hair poses a different appearance problem because geometric disorder, discoloration, reduced gloss, and strand-end breakage often appear together. Representing these cues with independent controls is possible, but it makes spatial consistency between geometry and shading an additional authoring task.

Fig. 1 illustrates the four appearance cues considered in this work: frizzy volume expansion, tip-biased discoloration, degraded highlights, and irregular strand termination. Our goal is not to simulate the chemical or biological processes that produce hair damage. Instead, we seek a compact real-time model that controls these visible consequences coherently on an existing strand- or hair-card asset.

The central idea is a shared local damage field defined over the hair. The field determines where damage is expressed, while component-specific response functions determine how strongly geometry, color, reflectance, and opacity react at each location. This separation keeps the individual operators lightweight while enforcing spatial correlation among the visible damage cues.

The resulting pipeline runs entirely on the GPU. The vertex stage evaluates the damage field and applies continuous strand deformation, whereas the fragment stage handles discoloration, reflectance modification, and strand-end erosion. Asset-dependent quantities, including strand exposure, persistent random seeds, and the rest-space strand frame, are precomputed once and reused during rendering.

The contributions of this work are:

- a spatial damage-field formulation that separates the location of damage from component-specific appearance responses;
- a GPU pipeline that uses the same local field to coordinate geometric disorder, discoloration, reflectance degradation, and strand-end fragmentation; and
- an evaluation of interactive control, component behavior, runtime overhead, and scalability across several hairstyle configurations.

Related Work

Hair Reflectance and Real-Time Shading

Hair rendering must account for anisotropic reflection, transmission, and multiple scattering from many thin fibers. Kajiya and Kay[1] introduced an early empirical fiber model, while Marschner et al.[2] measured scattering from human hair and decomposed it



Figure 1. Healthy and severely damaged hair motivating the modeled appearance cues: geometric disorder, tip-biased discoloration, degraded highlights, and fragmented strand ends.

into reflection and transmission paths. Subsequent work improved physical accuracy and practicality through multiple-scattering approximations[3], energy-conserving fiber models[4], production-oriented shading[5], anatomically motivated reflectance[6], and analysis of noncircular cross sections[7].

Interactive systems place additional emphasis on controllability and efficiency. Sadeghi et al.[8] exposed artist-friendly controls for hair shading, Xu et al.[9] supported interactive appearance editing under environment lighting, and recent work has explored efficient polygonal shading, GPU-generated hair meshes, and strand-level LOD[10–12]. These methods provide the optical and rendering foundations used here, but they primarily target healthy, stylized, or general-purpose hair rather than coordinated appearance degradation caused by severe damage.

Procedural Grooming and Strand Variation

Procedural displacement, noise, clumping, and strand-level variation are standard mechanisms for constructing and editing hair grooms. Such operators can readily produce local geometric irregularity, so the geometric deformation used in this work is not treated as an isolated contribution. Instead, it serves as one response to a shared spatial damage field, allowing the affected geometric regions to remain correlated with color, reflectance, and breakage responses.

Image-Based and Neural Hair Modeling

Image-based and neural approaches reconstruct detailed hair geometry or appearance from captured data. Neural Strands[13], HairStep[14], Neural Haircut[15], CT2Hair[16], and Dr.Hair[17] recover strand structure from images, video, or volumetric observations. More recent methods represent hair with Gaussian primitives or combine Gaussian and strand representations[18–21]. Their objective is typically reconstruction or relighting of captured hair. Our focus is complementary: procedural control of damage progression on an existing hair asset at interactive rates.

Damaged Hair and Procedural Appearance Degradation

Hair appearance depends on both internal pigmentation and the condition of the cuticle. Work on hair color digitization[22] and studies from cosmetic science and dermatology report that ultraviolet exposure, bleaching, and chemical treatments can be associated with color degradation, increased roughness, reduced gloss, split ends, and breakage[23–26]. These observations motivate the appearance cues modeled in this paper; they are not used as a physical damage simulator.

Procedural weathering models similarly describe spatially varying degradation through changes in color, reflectance, and geometry[27]. Hair requires a fiber-specific treatment because its damage cues vary along strands and across exposed regions. We therefore define a hair-specific damage field and use it as the common control signal for several lightweight shader responses.

Symbol	Range	Description
D_g	$[0,1]$	User-specified global damage level
s	$[0,1]$	Normalized root-to-tip coordinate
e_i	$[0,1]$	Exposure level of the i-th hair strand
$n(\xi)$	$[0,1]$	Normalized scalar procedural noise based on texture coordinates or strand ID
$P(s,e_i,\xi)$	$\mathbb{R}^+$	Spatial damage distribution pattern
w_s	$\mathbb{R}^+$	Contribution weight of the root-to-tip damage distribution
w_e	$\mathbb{R}^+$	Contribution weight of the strand exposure level
w_n	$\mathbb{R}^+$	Contribution weight of local stochastic variation
γ_s	$\mathbb{R}^+$	Exponent controlling damage concentration near hair tips

Table 1. Main parameters of the proposed damage field.

Damage-Conditioned Procedural Model

Model Overview and Inputs

The model has two layers. A spatial field $D_{\text{local}} \in [0,1]$ specifies where damage is expressed, and a set of response functions maps that field to geometry, color, reflectance, and opacity changes. The responses may have different onset rates and strengths, but they share the same spatial support.

The input is an existing strand- or hair-card asset with strand identifier i, normalized root-to-tip coordinate s, base color $\mathbf{C}_0$, tangent $\mathbf{T}$, and a stable rest-space strand frame. The user provides a global damage level D_g and parameters controlling the field and its responses. The base hair topology is unchanged.

The pipeline consists of three steps: (1) construct the local damage field from root-to-tip position, strand exposure, and deterministic variation; (2) map the field to component-specific response intensities; and (3) evaluate geometry in the vertex shader and appearance changes in the fragment shader.

Spatial Damage Field

For strand i, let $s \in [0,1]$ denote the normalized root-to-tip coordinate. We define a spatial pattern

$$P(s,e_i,\xi) = w_s s^{\gamma_s} + w_e e_i + w_n n(\xi), \qquad n(\xi) \in [0,1], \tag{1}$$

and scale it by the user control D_g:

$$D_{\text{raw}}(s,e_i,\xi) = D_g P(s,e_i,\xi), \tag{2}$$

$$D_{\text{local}}(s,e_i,\xi) = \text{clamp}(D_{\text{raw}}(s,e_i,\xi),0,1)\,. \tag{3}$$

Here, s^{γ_s} biases damage toward the tip, e_i modulates the contribution of externally exposed strands, and $n(\xi)$ introduces deterministic strand-level variation. The clamp guarantees that every downstream response receives a bounded control signal.

Strand exposure. e_i is stored as a *per-strand asset attribute*. For a groom with an available outer-layer or guide classification, e_i is assigned directly from that classification and normalized to $[0,1]$. When such metadata is unavailable, a practical default is to precompute e_i once from the rest groom by ranking strands according to their mean normalized distance from the local groom centerline or envelope; outer strands receive values close to one and interior strands values close to zero. The value is fixed in rest space and therefore does not fluctuate during animation. An artist-authored exposure mask can replace this automatic estimate without changing the shader.

Stable offset direction. The deformation direction must also be temporally consistent. For each strand, we construct a rest-space frame once. Let $\mathbf{T}_{i,j}$ be the tangent at sample j and let $\mathbf{r}_i$ be a reference direction at the root (for example, the scalp normal projected onto the plane orthogonal to the root tangent). We initialize

$$\mathbf{N}_{i,0} = \text{normalize}\,(\mathbf{r}_i - (\mathbf{r}_i \cdot \mathbf{T}_{i,0})\mathbf{T}_{i,0})\,, \tag{4}$$

and propagate the frame along the strand using the rest-shape tangent sequence (parallel transport or the equivalent frame construction used by the host hair representation). The resulting $\mathbf{N}_{i,j}$ is stored or reconstructed from rest-space attributes and transformed with the animated strand. Thus, the procedural deformation is tied to the strand frame rather than to camera- or frame-dependent random directions.

The noise seed ξ is derived from persistent strand identifiers and rest-space texture coordinates. Consequently, e_i, $n(\xi)$, and the offset frame are deterministic for a fixed asset. Because e_i, $n(\xi)$, and the strand frame are persistent, D_{local} is temporally deterministic for fixed D_g. If D_g is animated, smoothing can be applied to that user control at the application level.

Table 1 summarizes the parameters used to shape the field.

Phase-Shift-Based Geometry Disorder

The geometry operator converts D_{local} into a continuous displacement along the stable strand frame. The operator is deliberately simple: its role is not to introduce a new grooming primitive, but to provide a temporally coherent response whose spatial amplitude is correlated with the other damage cues. A harmonic basis is used instead of uncorrelated per-vertex white noise because it preserves continuity along each strand.



Figure 2. Comparison of hair volume expansion using white-noise-based displacement(a) and the proposed procedural-noise-based deformation(b). White noise produces irregular geometric artifacts, whereas the proposed method preserves strand continuity while generating damaged-hair volume expansion.

Fig. 2 shows the difference between simple white-noise-based vertex displacement and the proposed procedural deformation. When white noise is directly applied to vertex positions, the local direction of hair strands changes abruptly, producing artifacts that look more like geometric tearing or discontinuous clumping than natural frizziness of damaged hair. In contrast, the proposed method uses a continuous harmonic function with respect to the root-to-tip coordinate s and applies strand-level phase shifts, thereby preserving the overall curved flow of the hair. As a result, damage-induced volume expansion and disorder appear near the hair tips and boundary regions, while the continuity of individual strands is maintained, producing a more natural damaged-hair appearance.

Let $\mathbf{V}_0$ be the original vertex position of the i-th hair strand and $\mathbf{V}'$ be the deformed position. The vertex update in the vertex shader is defined as follows:

$$\mathbf{V}' = \mathbf{V}_0 + \mathbf{N}\left(A_{\text{vol}} D_{\text{local}} s^{\gamma_v} \sin\left(\omega_v s + 2\pi R_i\right)\right), \tag{5}$$

where A_{vol} is the maximum amplitude of volume expansion, γ_v controls how strongly the deformation increases toward the hair tips, and ω_v is the root-to-tip oscillation frequency. $R_i \in [0,1]$ is a random value assigned to the i-th hair strand, which varies the phase of each strand and reduces repetitive patterns. As D_{local} increases, the amplitude of volume expansion increases, but the deformation varies continuously along the hair length due to $\sin(\cdot)$, preserving the curved flow of the hair.

To represent more complex disorder, two or more frequencies can be combined:

$$\Delta_{\text{geo}}(s) = A_{\text{vol}} D_{\text{local}} s^{\gamma_v} \left[\eta_1 \sin(\omega_1 s + 2\pi R_i) + \eta_2 \sin(\omega_2 s + 2\pi R_i')\right], \tag{6}$$

$$\mathbf{V}' = \mathbf{V}_0 + \mathbf{N}\Delta_{\text{geo}}(s). \tag{7}$$

Here, η_1 and η_2 are blending weights for the low- and high-frequency components. The low-frequency component produces overall volume expansion, while the high-frequency component generates local disorder characteristic of damaged hair. However, if the high-frequency component is too strong, the hair may appear torn or jittery; therefore, we set $\eta_2 < \eta_1$ in this paper.

Damage-Aware Ombre Discoloration

In damaged hair, color degradation occurs due to the loss of internal melanin pigments, and discoloration is more prominent near the hair tips, which are more frequently exposed to external contact and heat stimuli. To represent this effect, we use procedural ombre interpolation based on the damage field and the length parameter, without requiring a separate discoloration texture.

Let $\mathbf{C}_0$ be the base color of healthy hair and $\mathbf{C}_{\text{bleach}}$ be the bleached color. The discoloration weight W_{dis} is computed as follows:

$$W_{\text{dis}}(s) = \text{smoothstep}\,(1 - D_{\text{local}}, 1, s)\,. \tag{8}$$

Eq. (8) makes the onset position of discoloration expand from the hair tips toward the root as D_{local} increases. The final color is computed as follows:

$$\mathbf{C}_{\text{hair}} = (1 - W_{\text{dis}})\mathbf{C}_0 + W_{\text{dis}}\mathbf{C}_{\text{bleach}}. \tag{9}$$

To produce more natural damaged color variation, the bleached color can be defined by controlling saturation and brightness according to the damage level, rather than using a single fixed color:

$$\mathbf{C}_{\text{bleach}} = \text{mix}\left(\mathbf{C}_0, \mathbf{C}_{\text{pale}}, D_{\text{local}}^{\gamma_c}\right), \tag{10}$$

where $\mathbf{C}_{\text{pale}}$ is a bright and desaturated hair color with significant melanin loss, and γ_c controls the nonlinearity of color degradation. This process procedurally represents the phenomenon in which discoloration becomes stronger toward the hair tips and the discolored region expands as the damage level increases.

Damage-Conditioned Reflectance Response

We use a Marschner-based hair shader as the healthy-hair baseline[2]. The method does not introduce a new scattering model. Instead, D_{local} modifies the orientation and width/strength of the existing lobes so that damage produces less coherent highlights.

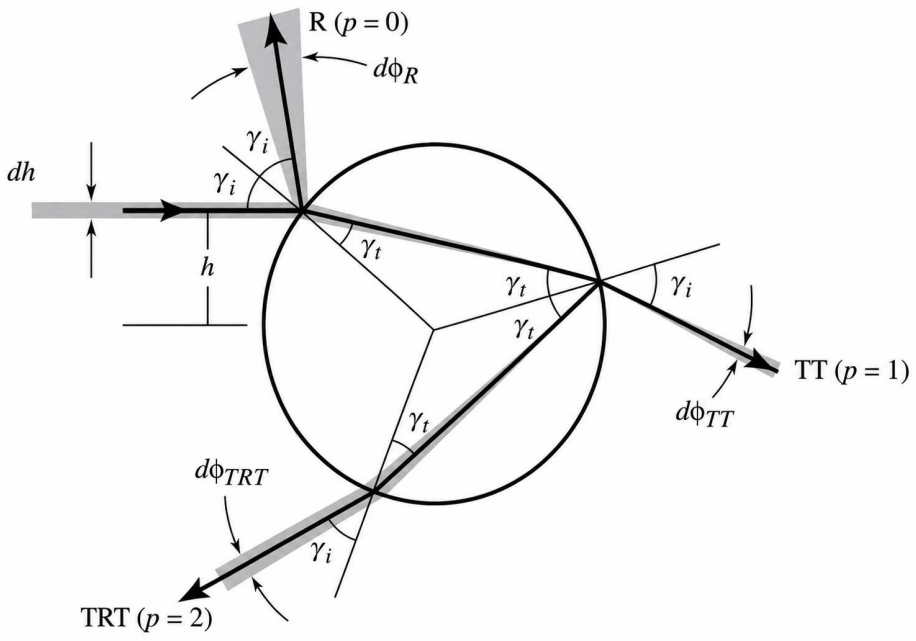


Figure 3. Marschner hair scattering paths used by the baseline shader[2].

For notational simplicity, let $D = D_{\text{local}}$. We first perturb the local tangent with deterministic rest-space noise,

$$S_{\text{cuticle}} = \alpha_c + DN_h(\mathbf{u}), \tag{11}$$

$$\mathbf{T}' = \text{normalize}\,(\mathbf{T} + \beta_c S_{\text{cuticle}}\mathbf{B} + \beta_n DN_r(\mathbf{u})\mathbf{N})\,, \tag{12}$$

where $\mathbf{B}$ and $\mathbf{N}$ belong to the stable strand frame and N_h, N_r are deterministic procedural signals. This operation decorrelates highlight orientation locally without introducing frame-to-frame noise.

To represent the loss of a narrow, coherent highlight, the damage response combines lobe broadening with peak reduction. Let $\sigma_R^0, \sigma_{TRT}^0, \sigma_{TT}^0$ denote the baseline lobe widths. We use

$$\sigma_q(D) = \sigma_q^0(1+\mu_q D), \qquad q \in \{R, TRT, TT\}, \tag{13}$$

with $\mu_q \geq 0$. If the host shader exposes lobe widths, Eq. (13) is applied directly. If it exposes only lobe weights, the same qualitative response can be approximated through tangent perturbation and conservative peak attenuation.

The damage-dependent lobe weights are written as

$$k_R(D) = k_R^0(1-\rho_R D), \tag{14}$$

$$k_{TRT}(D) = k_{TRT}^0(1-\rho_{TRT} D), \tag{15}$$

$$k_{TT}(D) = k_{TT}^0(1-\rho_{TT} D), \tag{16}$$

$$k_d(D) = k_d^0(1+\rho_d D). \tag{17}$$

The coefficients are clamped to nonnegative values in implementation. The final radiance is

$$\begin{aligned} L_{\text{hair}} = {} & k_R(D)L_R(\mathbf{T}', \sigma_R(D)) + k_{TRT}(D)L_{TRT}(\mathbf{T}', \sigma_{TRT}(D)) \\ & + k_{TT}(D)L_{TT}(\mathbf{T}', \sigma_{TT}(D)) + k_d(D)L_d. \end{aligned} \tag{18}$$

This formulation should be interpreted as an appearance approximation. Without measured damaged-hair BRDF/BSDF data, we do not claim that the particular attenuation or broadening rates are physically predictive. Their purpose is to reproduce the observed tendency toward less coherent, less glossy highlights while retaining compatibility with a real-time hair shader.

Procedural Alpha Erosion for Strand Breakage

Severe damage weakens the structural integrity of hair tips, and hair is not cleanly cut at a uniform length but instead becomes thinner, split, and irregularly torn. Directly modeling this geometrically would require subdividing hair strands or generating additional tip geometry, which is costly for real-time rendering. In this paper, we represent this breakage phenomenon by combining high-frequency-noise-based alpha erosion with tip tapering at the fragment shader stage.



Figure 4. Comparison between a conventional 3D line-rendering mask(a) and the proposed alpha-erosion mask for damaged hair strand termination(b). The proposed mask gradually thins and erodes strand tips, producing a more fragmented and frayed appearance.

Fig. 4 shows the difference between a conventional 3D line-rendering mask and the proposed alpha-erosion mask. The conventional method simply cuts hair tips at a fixed position, resulting in relatively blunt cross-sections that appear as if they

were cut with scissors. Such results cannot sufficiently represent the thinning, splitting, and irregular tearing observed in severely damaged hair.

In contrast, the proposed alpha-erosion mask applies both high-frequency noise and transparency tapering near the hair tips. This allows the hair to gradually become thinner along the length direction and be partially removed, instead of being cut with a uniform cross-section, resulting in fragmented tip shapes. Therefore, the proposed method can represent damaged-hair strand termination that appears broken or torn without generating additional geometry, entirely at the fragment shader stage.

First, we define a strand-specific tip damage value D_i^{tip} for each hair strand.

$$D_i^{\text{tip}} = D_{\text{local}}(1, e_i, \xi_i), \tag{19}$$

where D_i^{tip} is the local damage evaluated at the tip of the i-th strand. Using this value, the strand-specific cut threshold C_i is defined as follows.

$$C_i = 1 - R_i^{\text{cut}} \beta_{\max} D_i^{\text{tip}}, \tag{20}$$

As D_i^{tip} increases, the cut threshold moves toward the root, removing a larger portion of the tip region. Since both R_i^{cut} and D_i^{tip} vary across strands, the method avoids the artificial appearance of all hair strands being cut to the same length.

$$W_{\text{fray}}(s) = \text{smoothstep}\left(C_i - \delta_f, C_i, s\right), \tag{21}$$

where δ_f is the width of the lengthwise interval in which erosion occurs. W_{fray} increases as the fragment approaches the cut position and controls the region where high-frequency noise is applied.

The high-frequency alpha erosion mask is computed as follows.

$$M_\alpha(s, \mathbf{u}) = (1 - W_{\text{fray}}(s)) + W_{\text{fray}}(s) N_{\text{high}}(\mathbf{u}), \tag{22}$$

where $N_{\text{high}}(\mathbf{u})$ is high-frequency noise based on fragment texture coordinates. This term irregularly erodes the alpha value toward the hair tips, producing fragmented breakage shapes.

In addition, we apply a tip tapering function to avoid blunt cut surfaces.

$$F_{\text{tip}}(s) = 1 - \text{smoothstep}\left(C_i - \varepsilon_t, C_i, s\right), \tag{23}$$

where ε_t is a very narrow interval in which opacity rapidly decreases immediately before the cut position. The final alpha value is computed as follows.

$$\alpha_{\text{final}} = \alpha_0 M_\alpha(s, \mathbf{u}) F_{\text{tip}}(s), \tag{24}$$

where α_0 is the original opacity of the hair fragment. During rendering, the fragment is discarded if α_{final} is smaller than the threshold τ_α.

$$\text{discard} \quad \text{if} \quad \alpha_{\text{final}} < \tau_\alpha. \tag{25}$$

This process represents damaged tips that appear thin and torn without cutting the actual geometry.

GPU Shader Pipeline

The run-time pipeline is split between vertex and fragment stages. Asset-dependent quantities—strand exposure e_i, persistent random seeds, and the rest-space offset frame—are precomputed once on the host or during groom preprocessing and stored as compact per-strand/per-vertex attributes. This avoids recomputing values that do not change during editing. The vertex shader evaluates D_{local} and geometry disorder; the fragment shader evaluates color, reflectance, and alpha responses. Table 2 summarizes this decomposition.

Algorithm 1 summarizes the overall execution process of the proposed damage-field-based procedural hair appearance shader. At the vertex stage, the local damage level is computed using strand-level information and the root-to-tip coordinate, and vertex positions are updated accordingly. At the fragment stage, color, reflectance, and opacity changes are computed using the interpolated damage and geometric information, and the final damaged-hair appearance is output.

Stage	Operation	Main equations
Vertex shader	Damage field evaluation	Eqs. (2)–(3)
Vertex shader	Volume expansion	Eqs. (5)–(7)
Fragment shader	Ombre discoloration	Eqs. (8)–(9)
Fragment shader	Tangent perturbation	Eqs. (11)–(12)
Fragment shader	Reflectance remapping	Eqs. (14)–(18)
Fragment shader	Alpha erosion	Eqs. (20)–(25)

Table 2. Shader-stage decomposition of the proposed method.

Algorithm 1 Damage-field-based procedural hair appearance shader.

Require: Hair geometry $\mathscr{H}$, global damage D_g, base color $\mathbf{C}_0$, base shading parameters
Ensure: Rendered damaged hair appearance

for each hair vertex v **do**
  Read strand ID i, root-to-tip coordinate s, offset direction $\mathbf{N}$
  Read precomputed exposure e_i and persistent seed; evaluate $n(\xi)$
  Compute spatial damage pattern $P(s, e_i, \xi)$ using Eq. (1)
  Compute D_{raw} using Eq. (2)
  Compute D_{local} using Eq. (3)
  Compute geometry disorder Δ_{geo} using Eq. (6)
  Update vertex position $\mathbf{V}' \leftarrow \mathbf{V}_0 + \mathbf{N}\Delta_{\text{geo}}$
end for
for each hair fragment p **do**
  Interpolate s, $\mathbf{T}$, $\mathbf{N}$, $\mathbf{u}$, e_i, and D_{local}
  Compute discoloration weight W_{dis} using Eq. (8)
  Compute damaged color $\mathbf{C}_{\text{hair}}$ using Eq. (9)
  Perturb tangent direction $\mathbf{T}'$ using Eq. (12)
  Remap reflectance coefficients using Eqs. (14)–(17)
  Compute hair radiance L_{hair} using Eq. (18)
  Compute tip damage D_i^{tip} using Eq. (19)
  Compute strand cut threshold C_i using Eq. (20)
  Compute alpha erosion α_{final} using Eq. (24)
  if $\alpha_{\text{final}} < \tau_\alpha$ **then**
    Discard fragment
  end if
  Output final color $\mathbf{C}_{\text{hair}} L_{\text{hair}}$ and alpha α_{final}
end for

Damage-to-Response Mapping

In the proposed algorithm, each damage effect is controlled by the same D_{local}, but the sensitivity of each appearance change is set differently. For example, volume expansion may appear relatively early even at low damage levels, whereas strand breakage can be designed to become prominent only at severe damage levels. To this end, each component uses the following response functions. Because the new damage field definition scales D_{local} by the global damage level D_g, all component response functions remain at their healthy-hair state when $D_g = 0$.

$$R_{\text{vol}}(D) = D^{\gamma_{\text{vol}}}, \tag{26}$$

$$R_{\text{col}}(D) = D^{\gamma_{\text{col}}}, \tag{27}$$

$$R_{\text{ref}}(D) = 1 - \exp(-\kappa_{\text{ref}} D), \tag{28}$$

$$R_{\text{cut}}(D) = \text{smoothstep}(D_{\text{min}}^{\text{cut}}, D_{\text{max}}^{\text{cut}}, D). \tag{29}$$

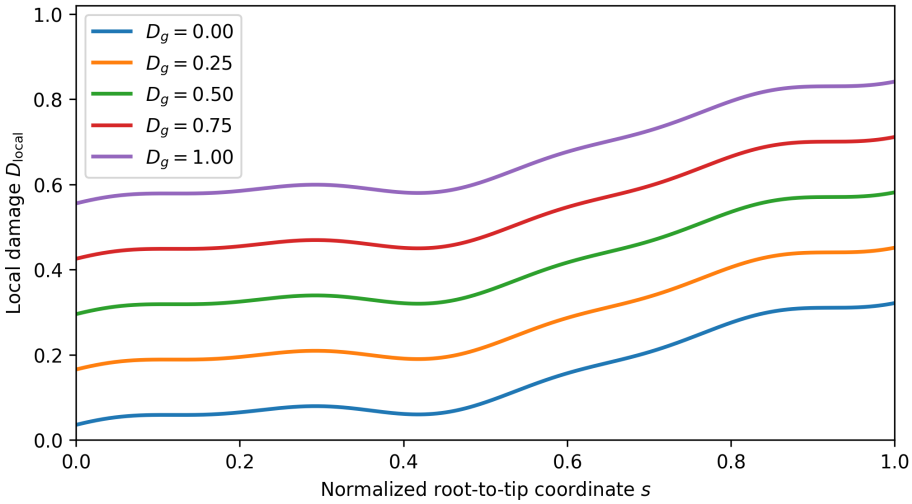


Figure 5. Damage field profiles along the normalized root-to-tip coordinate. As the global damage D_g increases, local damage expands from the hair tip toward the root while preserving strand-level variation.

Parameter	Default value	Role
A_{vol}	0.02–0.08	Maximum amplitude of volume expansion
ω_v	6.0–16.0	Frequency of root-to-tip harmonic deformation
γ_v	1.0–2.0	Tip concentration of volume expansion
γ_c	1.0–2.5	Nonlinearity of discoloration progression
ρ_R	0.4–0.8	Attenuation rate of surface reflection
ρ_{TRT}	0.3–0.7	Attenuation rate of the secondary highlight
ρ_{TT}	0.2–0.6	Attenuation rate of the transmission component
ρ_d	0.5–1.5	Increase rate of diffuse-like response
μ_q	implementation dependent	Damage-dependent broadening of host hair-shading lobes
$\beta_{\max}$	0.1–0.45	Maximum ratio of strand shortening
δ_f	0.03–0.12	Width of the alpha erosion region
ε_t	0.005–0.03	Width of the tip tapering region
τ_α	0.1–0.3	Fragment discard threshold

Table 3. Suggested default parameter settings used in our experiments.

R_{vol}, R_{col}, R_{ref}, and R_{cut} are the damage response intensities for volume expansion, color degradation, reflectance attenuation, and strand breakage, respectively. Using these response functions allows all effects to be linked to a single damage field while independently controlling the onset and intensity of each damage phenomenon.

Fig. 5 shows how the local damage level D_{local} changes along the hair length as the global damage level D_g increases. Since hair damage generally appears more strongly near the tips, which are more frequently exposed to external contact, heat, ultraviolet light, and chemical treatments, the proposed damage field is designed to produce larger damage values as s increases. In addition, as the global damage level D_g increases, the high-damage region progressively expands from the hair tips toward the root.

Fig. 6 shows the response curves of each damage component controlled by the same local damage level D_{local}. Although all effects are derived from a single damage field, the stage at which each damage phenomenon becomes visually prominent differs. Volume expansion appears relatively early even at low damage levels, while discoloration gradually becomes stronger as the damage level increases. Reflectance attenuation increases rapidly from the early damage stage, representing the weakening of continuous highlights in healthy hair. In contrast, strand breakage is designed to increase sharply at medium to high damage levels, emphasizing fragmented and torn tips in severely damaged states.

In the actual implementation, each damage response function is controlled by a small number of intuitive shader parameters. Table 3 summarizes the default parameter ranges used in our experiments. These ranges were set to generate stable appearance changes across various hairstyles while avoiding excessive geometric distortion, over-discoloration, or excessive alpha removal.

More detailed mathematical stability and approximation error analyses are provided in Appendix .

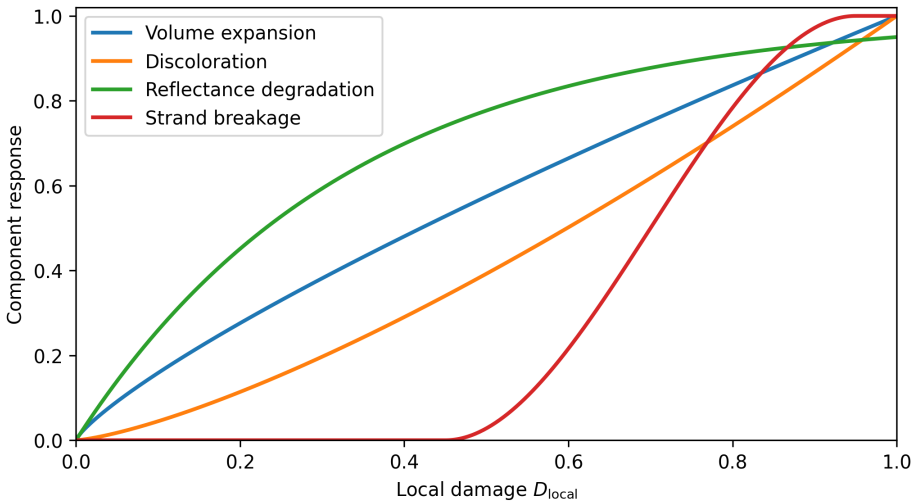


Figure 6. Component response curves driven by the local damage field. Volume expansion and reflectance degradation respond relatively early, while fragmented strand breakage becomes dominant at severe damage levels.

Evaluation

We evaluate the method from three perspectives: qualitative behavior across different hair assets, continuous control with the global damage parameter, and computational cost. All rendered results use the same field construction and component operators; hairstyle-specific behavior arises from the input geometry, root-to-tip coordinate, exposure attribute, and persistent strand variation.

Qualitative Results

Fig. 7 compares the healthy input and a severe-damage setting for several hairstyles. Across the examples, the shared field produces the same four classes of change: boundary and tip disorder, progressive lightening toward the tips, less coherent highlights, and irregular strand termination. The spatial expression differs with the asset: long hair emphasizes root-to-tip progression, whereas short and layered styles show stronger changes around exposed silhouette regions.

The result is not obtained by applying one uniform distortion to the entire groom. The length coordinate biases damage toward the tips, e_i emphasizes exposed strands, and persistent strand variation breaks up repeated patterns. The four rendering operators then respond to the same D_{local}, so geometry and shading changes remain co-located without requiring separate damage textures or masks.

Damage Progression

Figs. 8–10 show continuous control obtained by increasing D_g on long, short, and wavy assets. As D_g grows, the high-damage region expands while each response follows its own transfer function: volume disorder appears early, discoloration spreads gradually, reflectance becomes broader and weaker, and strand-end fragmentation becomes prominent at higher damage levels.

For long hair, the progression is most visible along the strand length: changes begin near the tips and extend upward as D_g increases. The overall silhouette remains recognizable because the geometric operator is continuous along the strand and the breakage operator modifies opacity rather than the base topology.

Short hair exhibits the same control logic but a different spatial appearance. Because strand lengths are shorter and more of the outer groom is exposed, disorder and fragmentation are concentrated around the silhouette, bangs, neckline, and layered boundaries rather than forming a long tip-to-root gradient.

The wavy example tests whether the deformation disrupts pre-existing curvature. The stable strand frame preserves the large-scale wave pattern while adding damage-dependent displacement along the original flow. Together, these examples show that the same control model produces asset-dependent spatial behavior without changing the algorithm or introducing hairstyle-specific masks.

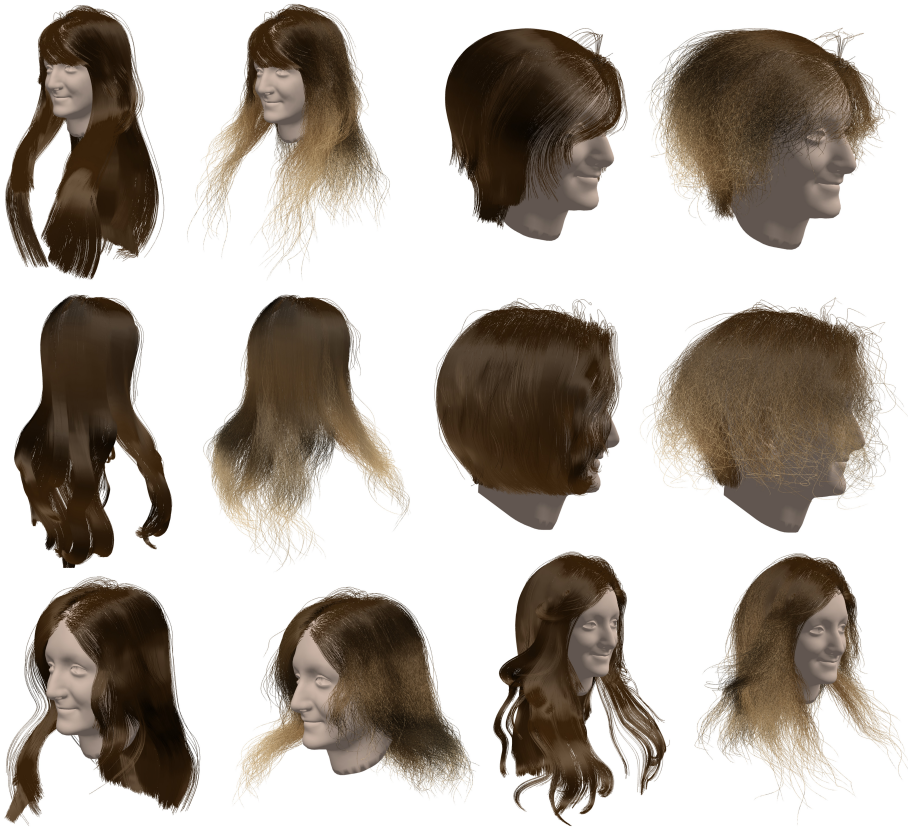


Figure 7. Qualitative results of the proposed damage-field-based procedural appearance model on various hairstyles. For each hairstyle, the left image shows the original healthy-hair appearance, while the right image shows the severely damaged result generated by the proposed method. The proposed shader consistently produces volume expansion, tip discoloration, reduced gloss, and fragmented strand breakage across different hair lengths and shapes.

Analysis and Discussion

Component Roles

The four operators serve distinct visual roles under the common field. Geometry deformation introduces frizzy disorder; discoloration controls tip-biased color degradation; reflectance modification reduces highlight coherence; and alpha erosion creates irregular strand termination. Table 4 summarizes the cumulative configurations used to isolate these functions.

Variant	Vol.	Color	Refl.	Break
Baseline hair shader	–	–	–	–
Volume expansion	✓	–	–	–
Discoloration	✓	✓	–	–
Reflectance remapping	✓	✓	✓	–
Alpha erosion	✓	✓	✓	✓

Table 4. Component configurations obtained by cumulatively adding damage-conditioned appearance operators. Each row activates an additional component to isolate its contribution to the final damaged-hair appearance.

The table is a configuration analysis rather than a perceptual ablation metric. Its purpose is to make the division of responsibility explicit: no single operator is intended to represent severe damage by itself, and the shared field provides the spatial coupling among them.

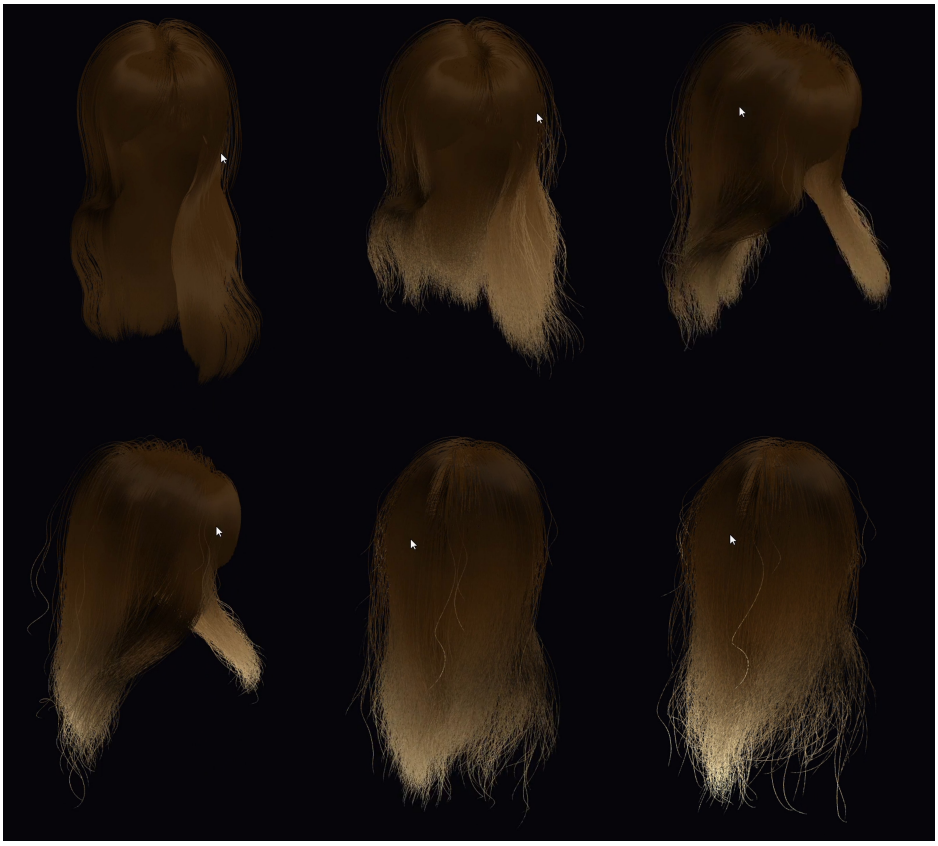


Figure 8. Real-time damage progression on a long-hair model. As the global damage level increases, the proposed method progressively produces tip discoloration, volume expansion, reduced gloss, and fragmented strand breakage while preserving the overall hairstyle structure.

Runtime Performance

All damage operations are local GPU computations. Table 5 reports average frame time and FPS on an AMD Ryzen Threadripper PRO 3955WX workstation with 256 GB RAM and an NVIDIA RTX 3090 at 1920×1080. Values were averaged over 300 frames after initialization under identical camera, lighting, and asset conditions.

The full method increases frame time from 6.59 ms to 7.42 ms, corresponding to a 12.6% overhead. Geometry deformation and discoloration add relatively little cost; the larger increments come from tangent perturbation, reflectance adjustment, high-frequency erosion, and fragment discard. Fig. 11 shows the cumulative trend.

Scalability

Because the method performs no inter-strand search or physical interaction solve, its cost primarily follows the number of processed vertices and fragments. Table 6 and Fig. 12 show the measured scaling with strand count.

The relative overhead rises from 13.7% at 25K strands to 19.5% at 200K strands. This increase is consistent with the additional per-fragment work, while the overall trend remains close to that of the baseline renderer.

Limitations

The method is an appearance model rather than a physical predictor of hair damage. The exposure attribute e_i is a groom-level proxy and does not model treatment history, heat, ultraviolet exposure, humidity, or contact. Likewise, the reflectance response uses a controllable broadening and attenuation approximation rather than measured damaged-hair BRDF/BSDF data.

Alpha erosion modifies visibility without changing the underlying strand geometry. Close-up views and physically based shadowing can therefore reveal differences from true split or broken fibers. The current evaluation also uses synthetic assets and a limited range of colors, lengths, and curl patterns. Direct comparison with photographs or measured damaged-hair appearance, as well as an explicit comparison against independently randomized geometry and shading controls, would provide stronger validation of the shared-field design.

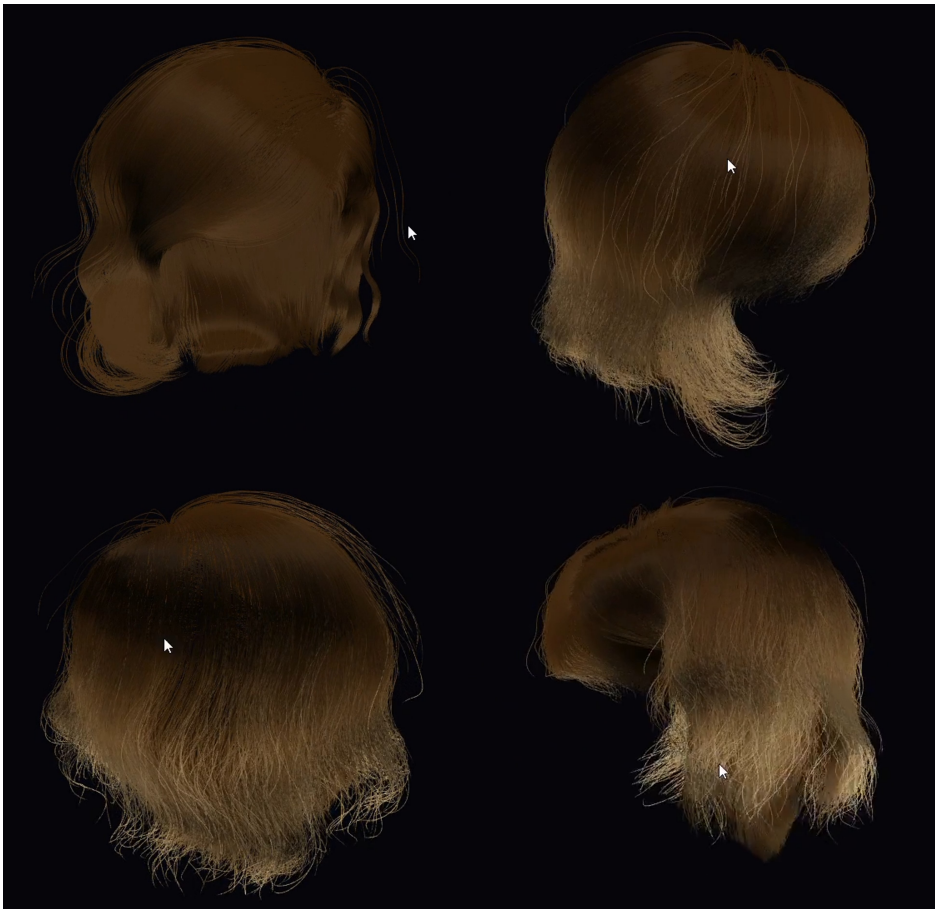


Figure 9. Real-time damage progression on a short-hair model. The proposed method generates short-hair-specific damage effects, including boundary frizz, localized discoloration, reduced gloss, and irregular strand breakage around the silhouette and hair tips.

Applications and Future Work

The model is suitable for interactive applications in which damage must be edited continuously, including digital humans, games, virtual production, and hair-care visualization. Future extensions could replace the procedural exposure term with image-derived or artist-painted damage maps, calibrate reflectance parameters from measured data, and combine shader-based fragmentation with lightweight geometry refinement for close-up rendering. Coupling the appearance field to mechanical properties such as stiffness, friction, or clumping would further extend the framework from static appearance control toward damage-aware hair dynamics.

Conclusion

We presented a shared damage-field framework for real-time appearance modeling of severely damaged hair. A spatial field determines where damage is expressed, while component-specific responses control geometric disorder, discoloration, reflectance degradation, and strand-end fragmentation. This separation provides coordinated spatial behavior without requiring dense auxiliary geometry or high-resolution damage textures.

The implementation combines precomputed strand attributes with local vertex- and fragment-shader operations and remains interactive across the tested hair complexities. The current results demonstrate controllable appearance progression on several synthetic hairstyles, while also indicating clear directions for stronger validation through real-hair references, measured reflectance data, and direct comparisons with independently controlled procedural effects.



Figure 10. Real-time damage progression on a wavy long-hair model. As the global damage level increases, the proposed method progressively generates tip discoloration, volume expansion, reduced gloss, and fragmented strand breakage while preserving the characteristic wavy hairstyle structure.

Variant	ms/frame	FPS	Overhead
Baseline hair shader	6.59	151.8	–
Volume expansion	6.73	148.6	2.1%
Discoloration	6.81	146.9	3.3%
Reflectance remapping	7.08	141.2	7.4%
Alpha erosion	7.27	137.5	10.3%
Full method	7.42	134.7	12.6%

Table 5. Runtime performance of cumulatively added shader components. The overhead is measured relative to the baseline hair shader.

Implementation Stability Notes

The local damage value is bounded by construction,

$$0 \leq D_{\text{local}} \leq 1, \tag{30}$$

so every downstream response receives a finite nonnegative control value. For fixed strand attributes, the harmonic deformation is continuous in the root-to-tip coordinate s and bounded by the selected amplitude parameters. The discoloration operator is a bounded interpolation between base and damaged colors, and the alpha operator only reduces the original opacity. These properties describe numerical behavior of the implementation; they do not constitute physical validation of real hair damage.

References

1. Kajiya, J. T. & Kay, T. L. Rendering fur with three dimensional textures. In *Proceedings of ACM SIGGRAPH*, 271–280, DOI: 10.1145/74334.74361 (1989).
2. Marschner, S. R., Jensen, H. W., Cammarano, M., Worley, S. & Hanrahan, P. Light scattering from human hair fibers. *ACM Transactions on Graph.* **22**, 780–791, DOI: 10.1145/882262.882345 (2003).
3. Moon, J. T., Walter, B. & Marschner, S. R. Efficient multiple scattering in hair using spherical harmonics. *ACM Transactions on Graph.* **27**, 31:1–31:7, DOI: 10.1145/1360612.1360631 (2008).
4. d'Eon, E., Francois, G., Hill, M., Letteri, J. & Aubry, J.-M. An energy-conserving hair reflectance model. *Comput. Graph. Forum* **30**, 1181–1187, DOI: 10.1111/j.1467-8659.2011.01976.x (2011).
5. Chiang, M. J.-Y., Bitterli, B., Tappan, C. & Burley, B. A practical and controllable hair and fur model for production path tracing. *Comput. Graph. Forum* **35**, 275–283, DOI: 10.1111/cgf.12830 (2016).

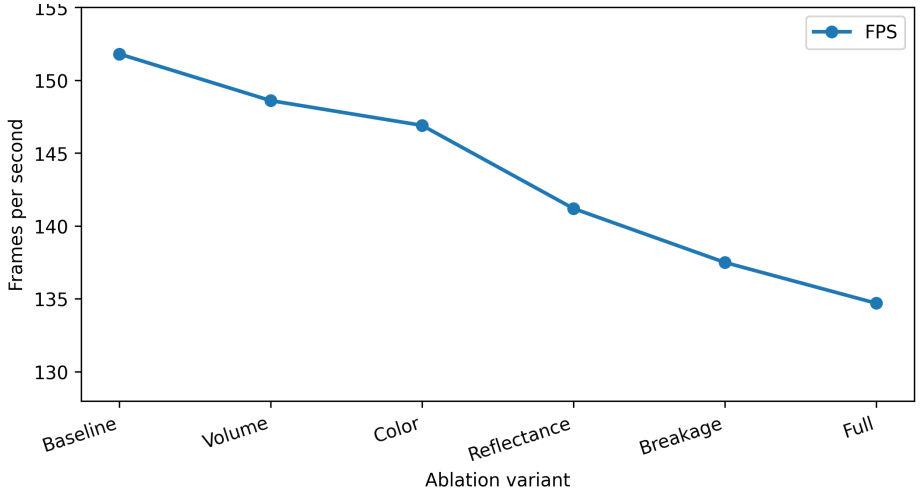


Figure 11. Runtime trend across shader components. Although each component adds a small computational cost, the full method remains suitable for real-time rendering.

Strands	Baseline ms	Full ms	Overhead
25K	5.1	5.8	13.7%
50K	7.4	8.5	14.9%
100K	11.8	13.9	17.8%
150K	16.2	19.2	18.5%
200K	21.0	25.1	19.5%

Table 6. Runtime scalability with respect to the number of hair strands. The proposed method shows a moderate overhead while preserving real-time performance for practical strand counts.

6. Yan, L.-Q., Tseng, C.-W., Jensen, H. W. & Ramamoorthi, R. Physically-accurate fur reflectance: Modeling, measurement and rendering. *ACM Transactions on Graph.* **34**, 185:1–185:13, DOI: 10.1145/2816795.2818080 (2015).

7. Khungurn, P. & Marschner, S. Azimuthal scattering from elliptical hair fibers. *ACM Transactions on Graph.* **36**, 13:1–13:23, DOI: 10.1145/2998578 (2017).

8. Sadeghi, I., Pritchett, H., Jensen, H. W. & Tamstorf, R. An artist friendly hair shading system. *ACM Transactions on Graph.* **29**, 56:1–56:10, DOI: 10.1145/1778765.1778793 (2010).

9. Xu, K., Ma, L.-Q., Ren, B., Wang, R. & Hu, S.-M. Interactive hair rendering and appearance editing under environment lighting. *ACM Transactions on Graph.* **30**, 173:1–173:10 (2011).

10. Gerard, M. Fast real-time shading for polygonal hair. arXiv preprint arXiv:2302.06468 (2023).

11. Bhokare, G., Montalvo, E., Diaz, E. & Yuksel, C. Real-time hair rendering with hair meshes. In *ACM SIGGRAPH 2024 Conference Papers*, DOI: 10.1145/3641519.3657521 (2024).

12. Huang, T. *et al.* Real-time level-of-detail strand-based hair rendering. arXiv preprint arXiv:2405.10565 (2024).

13. Rosu, R. A. *et al.* Neural strands: Learning hair geometry and appearance from multi-view images. In *Proceedings of the European Conference on Computer Vision*, 73–89 (2022).

14. Zheng, Y. *et al.* Hairstep: Transfer synthetic to real using strand and depth maps for single-view 3d hair modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023).

15. Sklyarova, V. *et al.* Neural haircut: Prior-guided strand-based hair reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19762–19773 (2023).

16. Shen, Y. *et al.* Ct2hair: High-fidelity 3d hair modeling using computed tomography. *ACM Transactions on Graph.* **42**, 1–13 (2023).

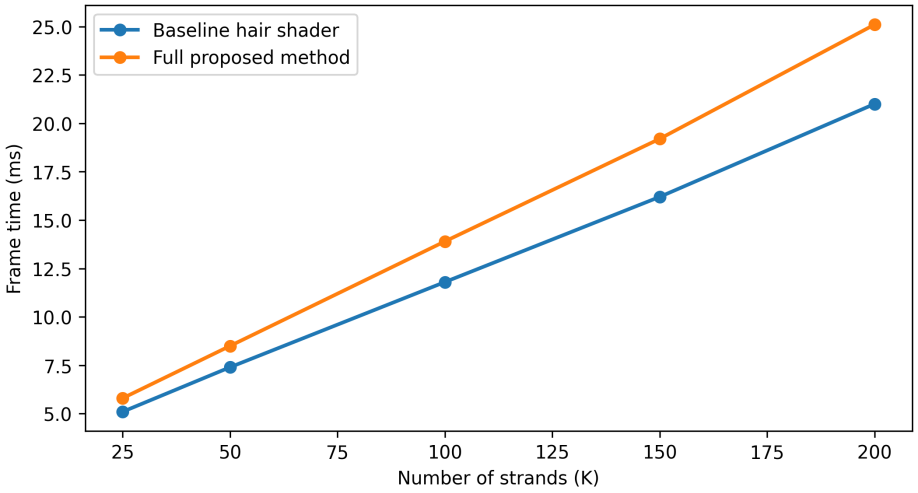


Figure 12. Runtime scalability with strand count. The proposed method follows the baseline trend because the damage operations are evaluated locally in the shader pipeline.

17. Takimoto, Y., Nagano, K., Saito, S., Nam, G. & Hodgins, J. Dr.hair: Reconstructing scalp-connected hair strands without pre-training via differentiable rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024).

18. Luo, H. *et al.* Gaussianhair: Hair modeling and rendering with light-aware gaussians. arXiv preprint arXiv:2402.10483 (2024).

19. Zakharov, E. *et al.* Human hair reconstruction with strand-aligned 3d gaussians. In *Proceedings of the European Conference on Computer Vision* (2024).

20. Pan, Y., Niessner, M. & Kirschstein, T. Hair strand reconstruction based on 3d gaussian splatting. arXiv preprint (2025).

21. Lazuardi, R. N., Sevastopolsky, A., Zakharov, E., Niessner, M. & Sklyarova, V. Geomhair: Reconstruction of hair strands from colorless 3d scans. arXiv preprint arXiv:2505.05376 (2025).

22. Kips, R., Bokaris, P.-A., Perrot, M., Gori, P. & Bloch, I. Hair color digitization through imaging and deep inverse graphics. arXiv preprint arXiv:2202.03723 (2022).

23. Lee, W.-S. Photoaggravation of hair aging. *Int. J. Trichology* **1**, 94–99 (2009).

24. Chandrashekara, M. N. & Ranganathaiah, C. Chemical and photochemical degradation of human hair. *J. Photochem. Photobiol. B: Biol.* **101**, 286–294 (2010).

25. Dario, M. F. *et al.* Effects of solar radiation on hair and photoprotection. *J. Photochem. Photobiol. B: Biol.* **153**, 240–246 (2015).

26. Maeda, K. Mechanism of cuticle hole development in human hair due to uv-radiation exposure. *Cosmetics* **5**, 24, DOI: 10.3390/cosmetics5020024 (2018).

27. Mérillou, S. & Ghazanfarpour, D. A survey of aging and weathering phenomena in computer graphics. *Comput. & Graph.* **32**, 159–174, DOI: 10.1016/j.cag.2008.01.003 (2008).

Data availability

All data generated or analysed during this study are included in this published article and its supplementary information files, including the source code used for the implementation.

Funding

This work was supported by INHA UNIVERSITY Research Grant. This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2026-25474535). This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) under the Artificial Intelligence Innovation Human Resources Development (IITP-2026-RS-2026-25548560) grant funded by the Korea government(MSIT)

Author contributions

J.H.K. conceived the study and developed the core methodology. J.H.K. and M.S.P. implemented the system and conducted the experiments. J.H.K. analyzed the results and wrote the manuscript. All authors reviewed and approved the final manuscript.

Ethics statement

This study involved only computer graphics simulations and rendering experiments using synthetic hair models. No human participants, human biological materials, animals, or personally identifiable data were involved. Therefore, institutional ethical approval and informed consent were not applicable to this study.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Code.zip