



Glare Mitigation using a Differentiable Unified Glare Rating

Linus Beresna  and Eugene Fiume 

Simon Fraser University, School of Computing Science, Canada

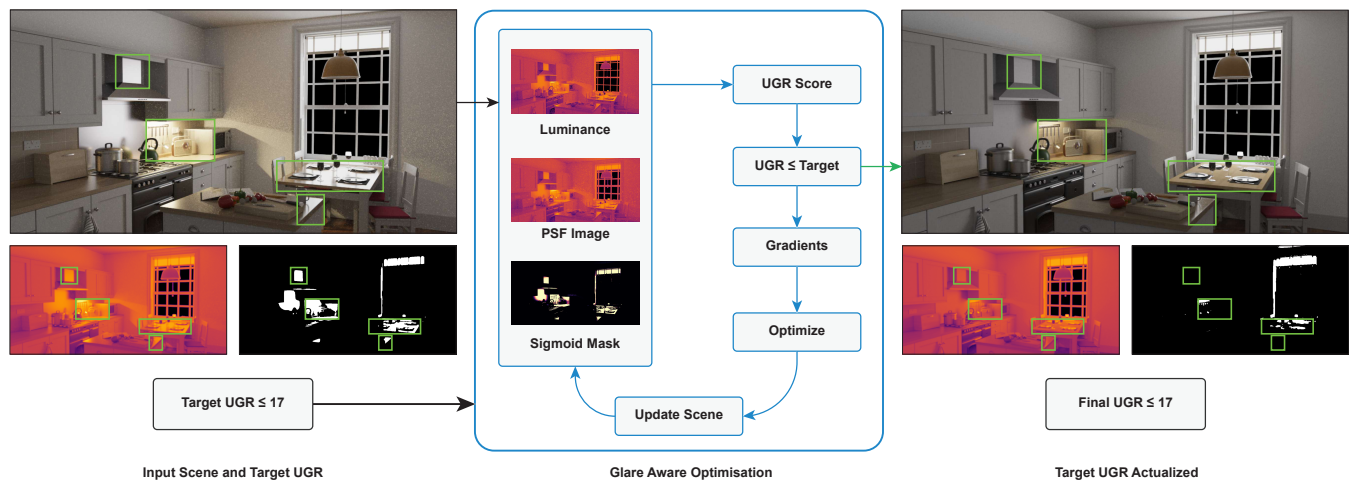


Figure 1: Starting from an initial high glare state (left), our differentiable rendering framework jointly optimises material parameters (α, η) and light source distributions to reduce the Unified Glare Rating. By utilising a point spread function and a soft sigmoid to formulate a continuous proxy, the framework smoothly guides the scene until a target ergonomic threshold (e.g., $UGR \leq 17$) is reached (right). Our method selectively mitigates localised glare hotspots while preserving the global illumination and aesthetic intent of the original indoor scene.

Abstract

Recent research in differentiable light transport extends the utility of computer graphics algorithms beyond traditional image generation, offering powerful tools for physical inverse design. In architectural and automotive applications, visual discomfort from glare is a critical design rating, traditionally quantified by the discrete CIE Unified Glare Rating (UGR). The standard UGR formulation relies on strict binary thresholds, making it fundamentally incompatible with smooth gradient-based inverse rendering. In this paper, we introduce a continuous, fully differentiable proxy for UGR. To resolve the severe optimisation instabilities caused by Monte Carlo variance at low sample densities, we introduce a differentiable optical scattering pass that simulates the Point Spread Function (PSF) of the human eye to heal fractured evaluation masks. We replace the discrete UGR step function with a tunable sigmoid boundary, enabling gradients to flow smoothly from the psychophysical measure back to the physical scene parameters. We deploy this differentiable framework to systematically reduce glare across three radiometric domains: surface-side microgeometry roughening, boundary-side index of refraction (IOR) optimisation, and source-side emitter gobo masking. By transforming a passive perceptual evaluation into an active loss landscape, our framework provides a robust, physics-based pipeline for optimizing visual comfort in complex global illumination environments.

CCS Concepts

• **Computing methodologies** → Ray tracing; Reflectance modeling; • **Mathematics of computing** → Continuous optimization;

1. Introduction

Architectural and automotive lighting design is fundamentally constrained by the need for human visual comfort. One major artifact that compromises this comfort is glare, defined as a psychophysical perception of excessive brightness that causes annoyance, discomfort, or loss in visual performance. In industry, discomfort glare is standardly quantified by the Unified Glare Rating (UGR), originally established by the International Commission on Illumination (CIE) [CIE95] and formalized in modern workplace lighting standards such as EN 12464-1 [Eur21, Arc24]. The UGR evaluates the discomfort caused by light sources relative to the ambient environment. From a computational perspective, it is critical to define this measure as a function of both the observer viewpoint v and the physical scene description S . Formally, the measure $\text{UGR}(v, S)$ is calculated as:

$$\text{UGR}(v, S) = 8 \log_{10} \left(\frac{0.25}{L_b} \sum_{i=1}^n \frac{L_i^2 \omega_i}{p_i^2} \right) \quad (1)$$

where L_b is the background adaptation luminance, L_i is the luminance of the i^{th} glare source in the direction of v , ω_i is the solid angle subtended by the source, and p_i is the Guth position index relative to v (see Figure 2 for a visualization of these spatial weights in camera space). Crucially, all material and lighting parameters that we adjust during optimization lie within the scene description S .

Traditionally, the CIE definition of UGR is applied strictly to primary light sources shining directly at the eye. In practice, however, industry tools also evaluate reflected glare over specific spatial domains like a desk or whiteboard, typically by placing virtual evaluation planes on the objects themselves. Because our framework is built on multi-bounce path tracing, essentially all surfaces in S become potential luminaires: we evaluate the UGR natively from the camera viewpoint v , treating any surface that reflects a harsh highlight into the observer’s field of view as an apparent glare source. This naturally unifies the traditional point-to-point measure with these spatial-domain evaluations.

In standard industry workflows (e.g., utilizing lighting design software such as DIALux [DIA26]), evaluating this measure is typically an iterative, forward-rendering task. Designers must manually adjust light placements, material properties, or architectural layouts, re-render the scene, and re-calculate the UGR until the value falls below a targeted discomfort threshold. This trial-and-error process is highly inefficient for complex, multi-bounce light transport scenarios. Differentiable rendering offers a systematic alternative by enabling inverse design: computing the gradients of a loss function with respect to scene parameters to automatically optimize materials and lighting. However, the standard UGR formulation presents mathematical discontinuities that can prevent standard gradient-based optimization. Specifically, identifying which pixels constitute the active glare sources L_i relies on a strict binary threshold (e.g., $L > 5L_b$). Mathematically, this formulation acts as a Heaviside step function. Consequently, the resulting gradients behave as Dirac delta functions, evaluating to zero almost everywhere and yielding infinite singularities precisely at the threshold boundary. This can prevent stable gradient descent and may yield unreliable convergence.

In this paper, we bridge the gap between psychophysical light-




Figure 2: Spatial components of the UGR camera model. (a) The Guth position index P maps human foveal sensitivity; dark central regions represent a value of ≈ 1.0 (highest sensitivity), while bright peripheral regions indicate lower sensitivity as P increases. (b) The pixel solid angle ω accounts for geometric perspective distortion, with peak values at the optical center. (c) The combined spatial weighting ω/P^2 highlights the central foveal region.

ing ratings and inverse rendering by introducing a continuous, fully differentiable proxy for UGR. By replacing discrete step functions with tunable, continuous approximations, we allow gradients to flow smoothly from the human-vision value, backward through the path-traced light transport, and into the physical parameters of the scene. To our knowledge, no prior work integrates a visual discomfort rating directly into a differentiable rendering loss landscape for inverse scene design.

We deploy this differentiable UGR rating to systematically evaluate automated glare mitigation across three distinct radiometric domains: source-side emission masking, boundary-side refraction (anti-reflective coatings), and surface-side microgeometry scattering (anti-glare roughening). Through this framework, we uncover a critical regularization trap in standard industry mitigation practices. We empirically demonstrate that while spatial roughening successfully diffuses mild glare, excessive roughening can inadvertently starve the surrounding environment of ambient light. Because the UGR formulation is inversely proportional to the background luminance L_b , this ambient darkening paradoxically spikes the discomfort rating. We show that retaining L_b within the computational graph provides a powerful self-regularizing effect that prevents this over-frosting failure state. Furthermore, we demonstrate that modifying the refractive index and source emission yields inherently stable loss landscapes, providing universally effective pathways for glare reduction.

In summary, our core contributions are:

- A continuous, differentiable formulation of the CIE UGR that enables gradient-based optimization of visual discomfort.
- An automated inverse design framework that establishes a methodological bridge between differentiable light transport and psychovisual computing, directly optimizing physical scene parameters against empirical perceptual models.
- An empirical analysis of light transport mitigation strategies, demonstrating the stability of anti-reflective and emission masking versus the ambient-starvation risks of microfacet roughening under complex global illumination environments.

2. Related Work

Our framework sits at the intersection of differentiable light transport, inverse material design, and psychophysical glare evaluation;

in this section, we analyze the foundational techniques in these domains and highlight how our approach differs from traditional photometric optimization.

2.1. Differentiable Rendering

Differentiable rendering (DR) spans rasterization-based methods [LB14, LLCL19, KUH18, GKM20] through to physically based path tracing [LADL18, ZRJ23, LHJ19]. Quantifying UGR requires strict radiometric accuracy, as glare frequently manifests through secondary light transport effects such as inter-reflections off glossy monitors, architectural glass, or metallic surfaces [CIE95, War94, CCDP15, PWB18]; rasterization-based approaches are therefore insufficient. We build on physically based path-tracing DR, including recent advances in volumetric and translucent transport [NDMKJ22, DLW*22, WRC*25]. By operating on fixed scene geometry, our framework avoids the primary visibility discontinuities that plague shape optimization, and we use Path Replay Backpropagation [VSJ21] to maintain constant memory while optimizing high-resolution spatial textures and emitter parameters.

DR is increasingly applied to source-side manipulation, from computational optics [TVB*25] to view-independent adjoint luminaire optimization for architectural targets [LHEN*24]. However, because discomfort glare is inherently view-dependent, tightly coupled to the observer’s retinal projections and the solid angle of each source, view-independent heuristics are insufficient. Our emitter optimization builds on these source-side principles while strictly maintaining a camera-space evaluation.

2.2. Inverse Material Design

Inverse material design has historically been formulated as an appearance matching problem, recovering SVBRDFs from real-world captures [AAL16, DAD*18, GLD*19, NDDJK21] via pixel-wise photometric losses. Our method departs from this paradigm: rather than mimicking a target image, we optimize against a global psychophysical objective, allowing the scene to deviate radiometrically as long as it converges toward human visual comfort. Unlike post-processing neural networks that hallucinate glare removal from 2D images [WHX*21, DLZ*22] without providing actionable design data, our gradients directly reflect physically accurate light transport.

Beyond appearance matching, inverse rendering has been applied to the functional design of optics and layered materials [JdJM14, GHZ18, TMM*21, SWF*21]. Selectively optimizing scene parameters corresponds directly to real-world architectural interventions, for example, index-matching anti-reflective thin films and chemically etched anti-glare surfaces. Furthermore, our framework redirects this machinery toward a human-centric visual comfort rating rather than imaging fidelity.

2.3. Psychophysical Glare Measures

Quantifying visual discomfort relies on empirically derived psychophysical measures that aggregate complex physical light transport into a single scalar rating. Within this domain, the Daylight

Glare Probability (DGP) is widely adopted for daylight-driven scenarios featuring massive, high-contrast exterior sources (e.g., the sun) [WC06]. However, for artificially lit indoor environments featuring multiple point or area luminaires and complex multi-bounce reflections, the UGR remains the formal international standard [CIE95, Eur21]. Because our framework targets the optimization of architectural materials, secondary specular bounces, and active indoor luminaires, UGR is the radiometrically appropriate objective.

Crucially, UGR evaluates discomfort as a ratio between localized glare sources and global adaptation luminance: naively minimizing scene brightness yields an unusable, darkened room, while flooding the scene with background luminance washes out contrast. Traditionally, these ratings have been evaluated *post-hoc* using forward-simulation rendering tools, most notably the Radiance system [War94, Jon17].

Recent research has focused on accelerating these evaluations using simplified screen-space heuristics or deep convolutional neural networks [PWB18], yet uniformly treats the measure as a read-only evaluation or forward prediction task. Such static, localized methods often fail to align with subjective perception under complex ambient lighting [EGS*26]. In contrast, by introducing a soft, continuous proxy that natively incorporates global background adaptation (L_b) into the computational graph, we transform UGR from a *post-process* evaluation tool into a dynamic, differentiable objective, dynamically balancing localized glare against the total ambient environment.

3. Differentiable Glare Formulation

To utilize the Unified Glare Rating as an active objective function for inverse design, we must translate its mathematical formulation into a robust, fully differentiable computational pipeline. In this section, we detail the mathematical and stochastic challenges of standard image-space evaluation, and introduce our two-stage continuous formulation designed to resolve them.

3.1. Image-Space Evaluation

As introduced in Section 1, the standard UGR formula aggregates the localized discomfort produced by all individual glare sources in the visual field against the background luminance. In practical lighting design, these calculated UGR values correspond directly to subjective human discomfort. A UGR of 10 or below is generally considered imperceptible, 19 is the universally accepted maximum threshold for standard office environments, and values exceeding 28 are deemed visually intolerable [CIE95, Eur21].

In the context of image-space evaluation, this formulation is calculated on a per-pixel basis using a rendered High Dynamic Range image [Bla13]. The summation iterates over contiguous clusters of pixels that are classified as glare sources. The standard methodology dictates that a pixel belongs to a glare source if and only if its luminance strictly exceeds a predefined adaptation threshold. In lighting design software, this boundary is typically defined by a binary inequality condition: $L_i > 5L_b$ [WC06]. Throughout our method, we choose to use a hard limit of 5 times the background



Figure 3: Visualizing spatial glare from a rendered scene. (a) A forward path-traced render of a bathroom environment featuring multiple reflective surfaces. (b) The photometric luminance distribution, capturing the extreme dynamic range of the incident light. (c) The resulting spatial glare mask, isolating active glare pixels that exceed the relative background adaptation threshold ($L > 5L_b$).

luminance, but any limit can be chosen. Figure 3 visualizes how a rendering compares against the luminosity of the view and the glare mask.

While binary thresholding can yield converging gradients in some stochastic setups, the discrete classification creates a harsh step in the loss landscape: a pixel crossing from $4.99L_b$ to $5.01L_b$ is abruptly hit with the full UGR penalty rather than a proportional one, causing the optimizer to overcorrect and produce unnatural material artifacts. We explore this empirically in Section 5. To traverse the landscape smoothly, we require a continuous approach that penalizes glare before it triggers a strict violation (Section 3.3).

3.2. Optical Scattering

While Monte Carlo path tracing computes accurate radiometric quantities, discomfort glare is fundamentally perceptual. Suppressing the stochastic variance that destabilizes optimization through high sample counts is impractical, as gradient-based inverse design requires hundreds of iterative passes; operating at low sample counts (e.g., 16 spp) is therefore necessary for practical computation times. By accommodating the eye’s biological response to this variance, our pipeline operates robustly on stochastic inputs without disrupting the gradient flow.

The human eye is not a perfect optical instrument. As light passes through the ocular media, specifically the cornea, the crystalline lens, and the vitreous humor, a significant portion of the incoming light undergoes intraocular scattering [SSZG95]. This biological scattering acts as a spatial low-pass filter, distributing the energy of intense localized light sources across a wider area of the retina. This phenomenon, often referred to as “veiling luminance” or “bloom,” is characterized by the Point Spread Function (PSF) of the human eye.

In the context of differentiable rendering, modeling this intraocular scattering provides a dual benefit. Physically, convolving the Monte Carlo radiance estimate with an approximation of the eye’s Point Spread Function (PSF) translates incident scene radiance into perceived retinal luminance [KMN*04]. While classical light transport defines this sensor response via the measurement equation [Vea98], our framework extends it into the differentiable do-

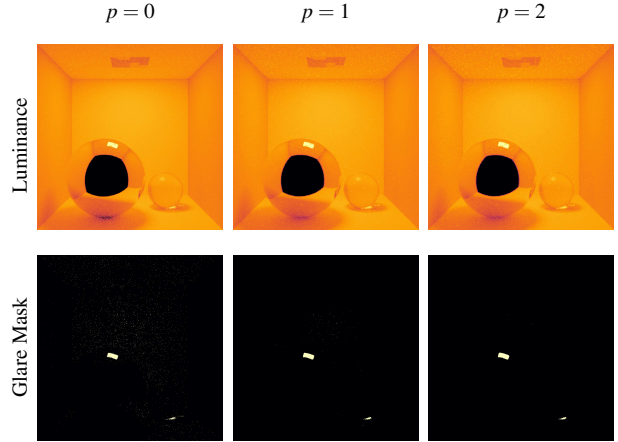


Figure 4: (Top) As the filter strength increases ($p = 0 \rightarrow 2$), the raw luminance image becomes progressively blurred, spatially distributing the energy of high-variance Monte Carlo samples. (Bottom) This energy distribution stabilizes the corresponding glare masks. Note the isolated white pixels scattered throughout the background at $p = 0$; as the filter diffuses their energy, these false-positive spikes fall below the evaluation threshold and are suppressed by $p = 2$. This isolates primary glare source and prevents stochastic noise from destabilizing the optimization gradients.

main, so the biological PSF becomes an active gradient pathway rather than a forward-evaluation filter.

Concretely, we implement this PSF as p sequential applications of a 3×3 Gaussian kernel, with corner, edge, and center weights of $\frac{1}{16}$, $\frac{1}{8}$, and $\frac{1}{4}$ respectively, applied to the luminance image prior to mask computation. The parameter p reported in Figure 4 denotes the number of scattering passes. Mathematically, this physically-grounded scattering stabilizes the inverse rendering process, by absorbing outlier radiance values.

3.3. Soft UGR Proxy

To enable gradient-based optimization, we must replace the discrete $L > 5L_b$ classification with a continuous, differentiable proxy. We achieve this by formulating a soft masking function that assigns a fractional glare weight, $w_{glare} \in [0, 1]$, to every pixel in the image based on its relative luminance.

First, to establish a stable adaptation threshold without allowing the optimizer to trivially scale the global exposure to artificially minimize the loss, we compute the global mean luminance of the scene, \bar{L} , and detach it from the computational graph. The glare threshold is then defined as $L_{\text{thresh}} = 5\bar{L}$.

We define the relative luminance difference for a given pixel i as:

$$\Delta_i = \frac{L_i - L_{\text{thresh}}}{L_{\text{thresh}}}. \quad (2)$$

This formulation is mathematically equivalent to the *Weber contrast* [Pel90], a fundamental measure in psychophysics that describes the perceived difference between a stimulus and its background. By framing the glare source identification in terms of con-

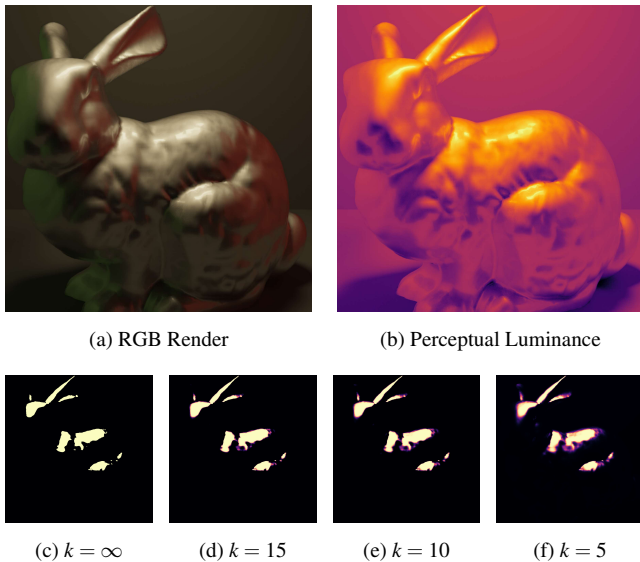


Figure 5: Differentiable boundary formulation via Sigmoid softening. (a, b) The baseline render and its perceptual luminance. (c) Even after resolving Monte Carlo variance, a standard binary threshold produces a Heaviside step function, preventing gradient propagation outside the immediate glare source. (d-f) By relaxing the sharpness parameter (k), we expand the functional penumbra into a continuous gradient. This ensures that scene parameters just outside the threshold receive meaningful optimization signals, smoothly guiding them toward compliance.

trast rather than raw radiance, our framework aligns with established models of human visual perception.

To map this relative luminance difference into a continuous spatial mask, we apply a logistic sigmoid function, following the sigmoid-based luminance remapping used in photographic tone reproduction [Rei02]:

$$m_i = \frac{1}{1 + \exp(-k\Delta_i)}, \quad (3)$$

where k is a hyperparameter dictating the steepness of the transition threshold. We set $k = 10.0$ to ensure non-zero derivatives during optimization while maintaining a strictly localized boundary around high-radiance sources. Consequently, a pixel whose luminance significantly exceeds the evaluation threshold approaches a mask value of $m_i \approx 1.0$, whereas pixels below the threshold approach 0.0.

Using this continuous mask, we can smoothly separate the scene into glare and background components. The effective glare luminance for pixel i becomes:

$$L_{s,i} = L_i \cdot m_i. \quad (4)$$

The background adaptation luminance L_b must also be evaluated continuously to maintain energy conservation and prevent gradient discontinuities. We compute this global adaptation level by weighting the incident scene luminance against the complementary spatial

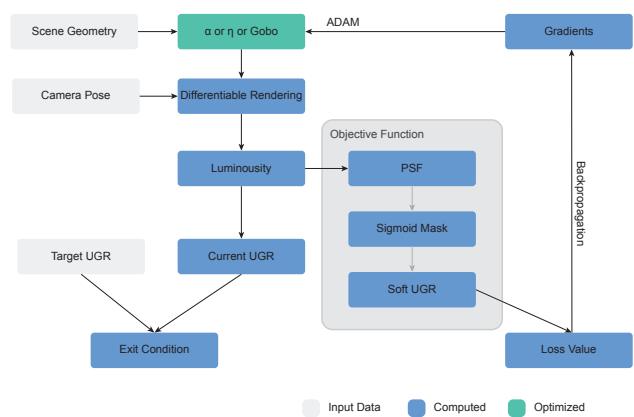


Figure 6: Our framework interfaces with the Mitsuba 3 renderer to compute forward-pass luminance and glare measures. Gradients are backpropagated through a soft UGR proxy to update scene parameters (α , η , and emitter masks) using the Adam optimizer, subject to TV regularization and physical boundary constraints. The optimization exits once the target UGR has been achieved or up to a maximum iteration count.

mask, $(1 - m_i)$:

$$L_b = \frac{\sum_{i=1}^n L_i (1 - m_i)}{\sum_{i=1}^n (1 - m_i) + \epsilon}, \quad (5)$$

where ϵ is a small numerical constant preventing division by zero.

By substituting these continuous formulations of $L_{s,i}$ and L_b back into the original UGR equation, we obtain a differentiable UGR proxy. This formulation converges to the standard discrete measure at its boundary conditions, specifically, when a pixel's luminance is either decisively below or drastically above the evaluation threshold ($m_i \rightarrow 0$ and $m_i \rightarrow 1$, respectively). Between these boundaries, the proxy provides a continuous, differentiable loss landscape, allowing gradients to propagate directly from the psychophysical evaluation back to the physical scene parameters.

4. Implementation and Optimization

With our continuous glare formulation established, this section details its integration into a fully differentiable rendering pipeline and its systematic application to physical scene parameters.

4.1. Pipeline and Loss Formulation

A high-level overview of our optimization pipeline is illustrated in Figure 6. We implemented our automated glare mitigation framework in Mitsuba 3 [JSR*22], leveraging the Dr.Jit compiler [JSRV22] to evaluate our continuous UGR proxy natively on the GPU. Global illumination gradients were computed using PRB [VJS21].

The gradient descent process was driven by the Adam optimizer [KB14] with the base learning rate dependent on the scene and the optimization parameter. Because inverse rendering via Monte Carlo

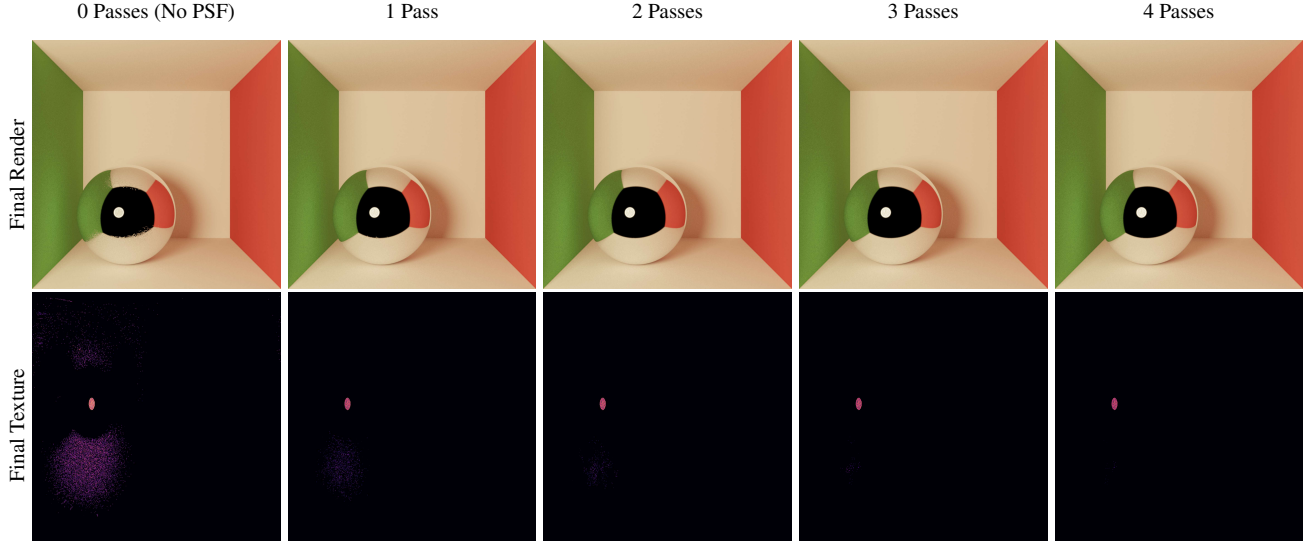


Figure 7: Optimizing the roughness texture map of the sphere to reduce glare; the optimized roughness map is shown in the bottom row, with the resulting render above. Optimizing without optical scattering (0 Passes) causes the optimizer to overfit to stochastic “firefly” pixels: isolated Monte Carlo noise spikes are misclassified as glare sources, so the optimizer roughens scattered, incoherent texels attempting to suppress glare that is not physically present, destroying the spatial coherence of the resulting texture map. Applying sequential Point Spread Function passes acts as a low-pass spatial filter, distributing this high-intensity energy to heal the evaluation mask. This ensures stable gradient flow and yields a cleanly converged, physically plausible roughness map.

integration is inherently noisy, optimizing high-resolution spatial textures is prone to severe overfitting. To prevent the optimizer from embedding high-frequency noise patterns into the materials, our total loss function incorporates Total Variation (TV) regularization for spatially varying parameters (emitter masks and AR blending distributions):

$$\mathcal{L}_{\text{total}} = \text{UGR}_{\text{soft}}(\theta) + \lambda_{TV} \mathcal{R}_{TV}(\theta_{\text{spatial}}), \quad (6)$$

where θ represents the full set of optimizable scene parameters, $\theta_{\text{spatial}} \subseteq \theta$ denotes the subset of spatially-varying 2D textures subject to spatial regularization, \mathcal{R}_{TV} is the TV penalty, and λ_{TV} dictates the regularization strength. While we found TV regularization to be generally effective for enforcing manufacturable, piecewise-smooth physical boundaries, the framework is agnostic to the specific penalty; any appropriate spatial regularizer (e.g., Laplacian smoothing or sparsity constraints) can be chosen depending on the specific architectural context and material being optimized.

Because the UGR proxy applies non-linear functions of the rendered luminance L to compute the rating, both the squared source term and the outer logarithm in Equation 1, care is needed when estimating the gradient of $\mathcal{L}_{\text{total}}$ in Equation 6. For a term $f(L)$, the gradient with respect to scene parameters θ is $df/dL \cdot dL/d\theta$; when f is non-linear, evaluating df/dL at a Monte Carlo estimate of L yields a biased gradient, since only L itself is estimated without bias [NRN*23]. We mitigate the effect by using a higher sample count in the forward pass to reduce the variance of L (consistent with the deeper forward-pass path lengths discussed in Section 5.3); debiasing constructions such as Russian-roulette estimators [MBGJ22] offer a complementary alternative.

Finally, the gradient descent loop implements an early stopping criterion. The optimization halts immediately once the calculated UGR of the scene falls below a user-defined ergonomic threshold (e.g., $\text{UGR} \leq 19$ for standard office spaces). This prevents the optimizer from needlessly darkening the environment or over-altering materials once visual comfort has been achieved.

Across all experiments we apply $p = 4$ PSF scattering passes and a maximum of 100 optimization iterations, with early stopping once the target UGR (typically ≤ 17) is reached. The Adam learning rate was set in the range 0.05–0.1 depending on the scene and parameter domain, and the TV regularization weight λ_{TV} was set in the range 0.01–0.1 depending on the spatial resolution of the optimized texture.

4.2. Radiometric Mitigation Domains

With a fully continuous loss landscape established and our pipeline configured, we can systematically apply gradient descent to the physical scene parameters responsible for visual discomfort. We categorize our automated mitigation strategies into three distinct radiometric domains.

Source Side: Emitter Radiance Masking

The most direct intervention occurs at the light source. We parameterize the spatially varying emission profile of the luminaires as high-resolution textures. The optimizer selectively dims specific directional solid angles, effectively generating an automated, view-dependent *computational gobo* that maximizes comfort without altering scene geometry.




Figure 8: Ablation of the evaluation threshold. Compared to the glaring initial state (a), the discontinuous Hard Step function (b) forces the optimizer to over-compensate during gradient descent, resulting in an excessively blurred, unnatural surface finish. Our continuous Soft Sigmoid formulation (c) achieves a visually comfortable convergence without over-altering the material.

Boundary Side: Anti-Reflective (AR) Optimization

When light sources are fixed (e.g., daylighting), we target the physical boundary of the glaring object.

- *Scalar Refraction*: We expose the relative index of refraction to the optimizer, guiding it toward minimizing Fresnel reflections.
- *Spatial Blending*: For composite surfaces, we optimize a high-resolution mask interpolating between specular and AR-coated materials. To ensure physically plausible, contiguous patterns and prevent high-frequency noise, we apply the aforementioned TV regularization loss to the spatial gradients of the mask.

Surface Side: Anti-Glare (AG) Scattering

We simulate industry standard scattering (e.g., frosted glass) by exposing the microfacet roughness parameter (α).

5. Results and Discussion

We evaluate our automated glare mitigation framework through a series of controlled ablations and complex architectural scenarios. Because inverse rendering performance is fundamentally bound by Monte Carlo sample counts and scene complexity, our empirical analysis prioritizes the robustness of the optimization landscape over raw wall-clock execution time.

5.1. Differentiable Formulation and Stability

Our initial experiments validate the necessity of modifying the standard UGR formulation for gradient-based optimization, specifically regarding Monte Carlo noise and discontinuous thresholding.

Resolving Variance via Optical Scattering. To evaluate the impact of Monte Carlo noise on spatial material optimization, we optimized a roughness texture to a target UGR using identical setups, isolating the presence of our PSF scattering passes as the sole variable. As demonstrated in Figure 7, optimizing without PSF scattering causes the objective function to misidentify stochastic “fireflies” as true glare sources. Consequently, the optimizer updates incorrect regions of the texture map, attempting to mitigate non-existent glare and destroying spatial coherence. Simulating intraoc-

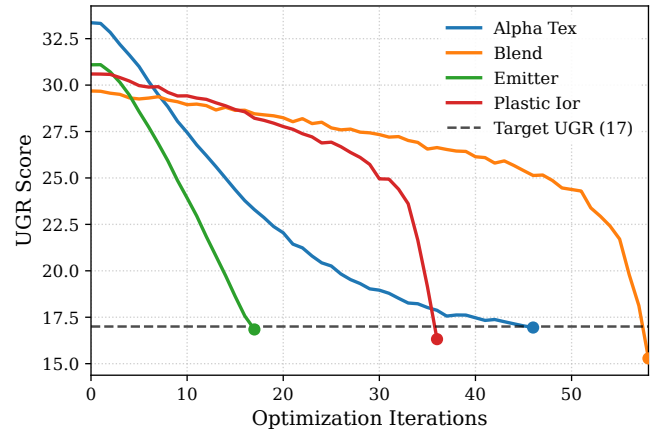


Figure 9: We plot the UGR reduction over time for each parameter isolated in Figure 10. Our framework demonstrates consistent, smooth convergence regardless of whether the target is a surface property (Roughness, IOR) or a light control parameter (Radiance, Blend), reaching the target UGR

ular scattering acts as a critical low-pass filter, dilating these noise spikes and yielding cleanly converged texture maps.

Continuous Gradients via Soft Sigmoid. We also performed a controlled ablation to compare our continuous Soft Sigmoid formulation against a standard Hard Step function. In Figure 8 both optimizations were initialized from the identical high-glare Cornell Box state and halted upon reaching the same target UGR. While the Hard Step function technically succeeds in reducing the UGR in this extreme scenario, it produces noticeable visual artifacts. Because each iteration causes new pixels to abruptly cross the binary glare threshold, those pixels are suddenly hit with the full force of the Adam optimizer’s gradient update. This discontinuity causes the optimizer to over-compensate, resulting in excessive and unnecessary surface blurring. Conversely, our Soft Sigmoid formulation provides a smooth gradient penumbra, reaching the target UGR while preserving the intended physical specularly of the material.

5.2. Validation Across Radiometric Domains

To demonstrate the versatility of our differentiable UGR framework, we validated its performance across four distinct physical domains within the controlled environment of the Cornell Box. Our objective was to prove that visual discomfort can be systematically mitigated regardless of the specific architectural constraint; whether by altering surface microgeometry, boundary refraction, or source emission.

As illustrated in Figure 10, we successfully optimized the following parameters from an initially high-glare state down to an ergonomic target ($UGR \leq 17$):

- **Surface Scattering:** Increasing the microfacet roughness (α) of a metallic conductor to diffuse the specular highlight.
- **Boundary Refraction:** Adjusting the Index of Refraction (η) of a dielectric plastic to minimize Fresnel boundary reflections.

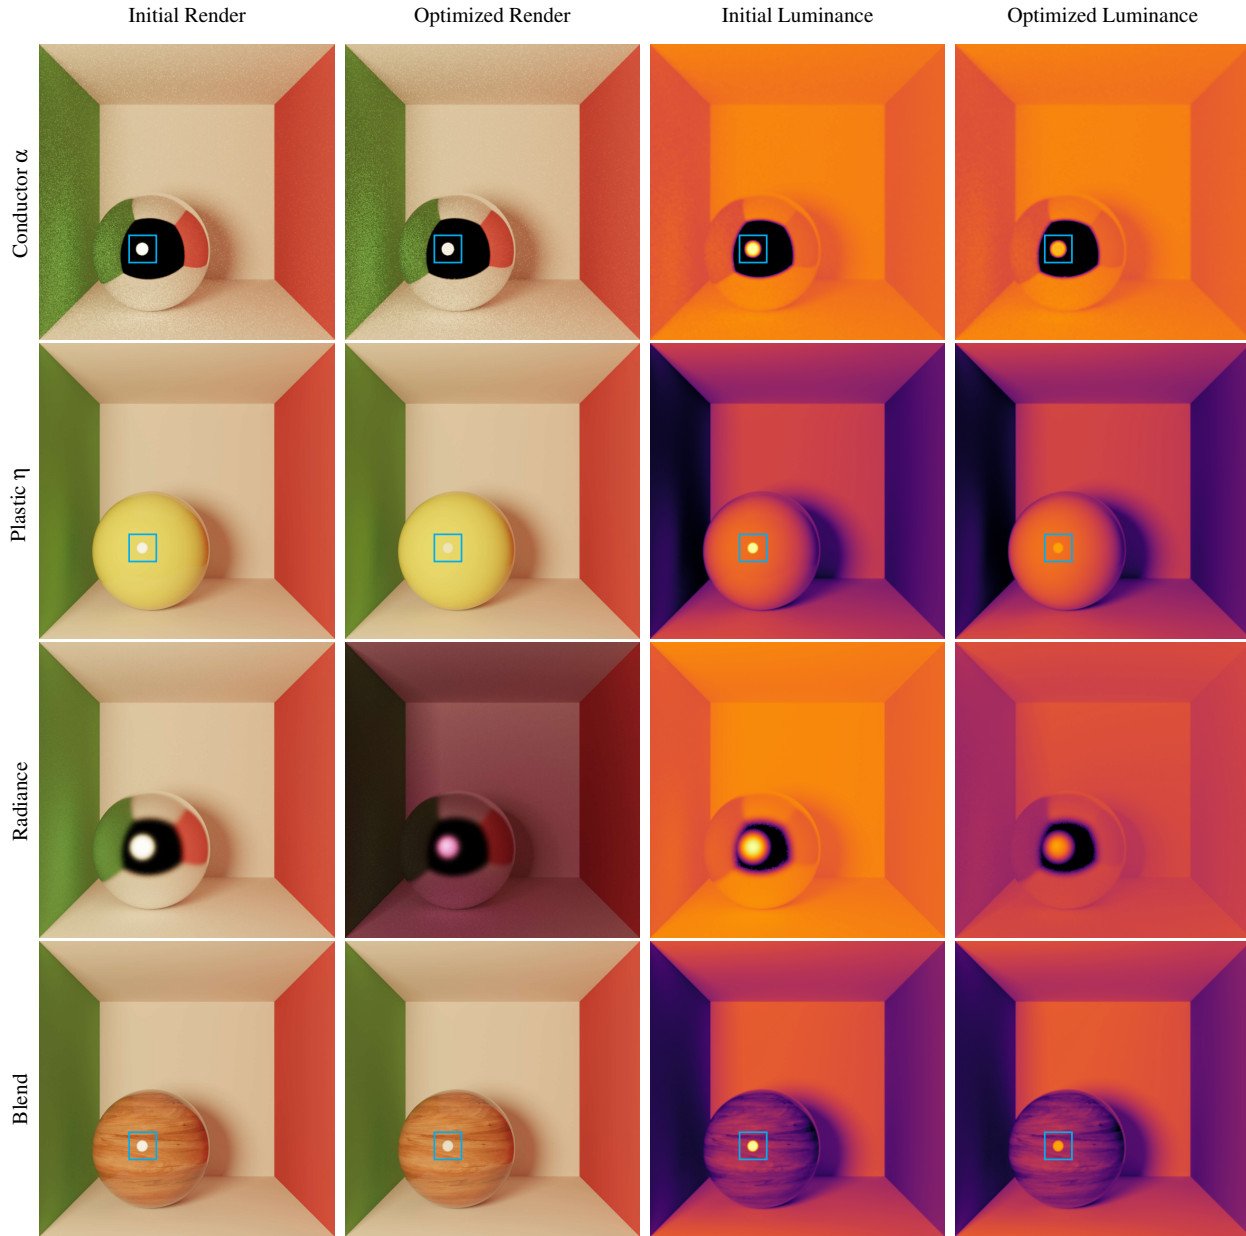


Figure 10: We evaluate our differentiable UGR framework across four distinct physical domains. Each row demonstrates the transition from a high-glare state to an optimized state. Pairing initial and optimized renders side-by-side highlights the subtle material and radiance adjustments, while the corresponding luminance maps demonstrate the resulting shift in high-intensity energy distributions used by the UGR metric. Notably, direct emitter optimization (row 3, optimized render) inherently darkens the overall scene and induces a hue shift, a trade-off discussed further in Section 5.4.

- **Source Emission:** Modulating the radiance of the primary area emitter to selectively dim the problematic solid angles.
- **Spatial Blending:** Optimizing a continuous 2D mask to blend between a highly specular base material and a diffuse fallback (simulating localized anti-reflective coating).

Figure 9 tracks the UGR reduction over time for each of these four domains. The results demonstrate that our formulation pro-

vides a stable, continuous loss landscape regardless of the underlying physical parameter being optimized. Whether the target is a scalar material property (Roughness, Index of Refraction) or a high-resolution spatial texture (Radiance, Blend), the framework converges to the target ergonomic threshold.

Beyond geometric and physical versatility, we evaluated the computational scalability of our adjoint framework against stan-

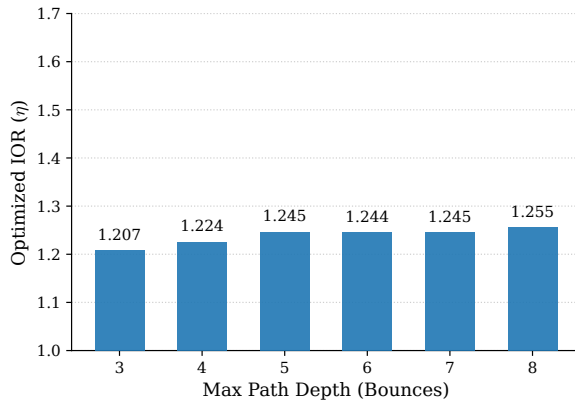


Figure 11: We optimize the Index of Refraction (η) of a dielectric sphere under varying maximum path lengths. At low path lengths, the integrator underestimates the ambient background luminance (L_b), causing the optimizer to overcompensate by aggressively reducing the IOR. Deeper path lengths provide sufficient L_b to naturally mitigate the glare ratio, allowing the optimizer to reach the target UGR while preserving the material’s physical specularity.


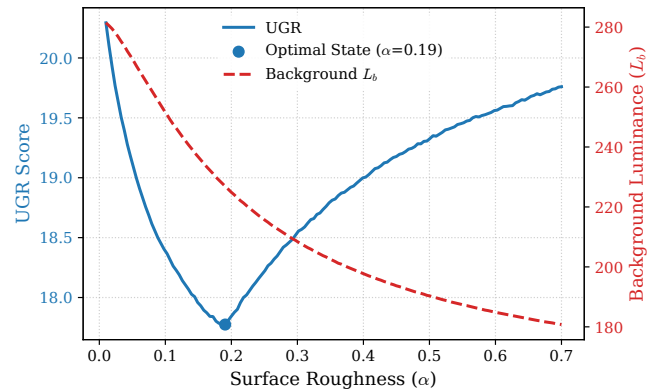
standard derivative-free optimizers, specifically CMA-ES [Han16] and Powell’s method [Pow64]. While these gradient-free baselines can successfully optimize global scalar values, they suffer from exponential scaling constraints. As parameter dimensionality increases, particularly for spatial textures at resolutions of 64×64 ($N = 4096$) and beyond, these solvers exhaust their evaluation budgets and are prone to timing out. In contrast, our adjoint formulation computes gradients for all parameters simultaneously, allowing execution time to remain near-constant regardless of the spatial resolution. We provide a comprehensive quantitative and qualitative breakdown of this dimensionality scaling and baseline comparison in Appendix A.

While the isolated Cornell Box experiments validate the theoretical stability of our framework, practical architectural design potentially requires mitigating glare across multiple surfaces simultaneously. To demonstrate the scalability of our method, we applied it to a highly non-linear, multi-bounce kitchen environment, as shown in Figure 1. Starting from an initial high-glare state, the framework performed a joint optimization across four distinct parameters: the scalar roughness (α) of the metallic extractor hood, the scalar index of refraction (η) of the wooden table top (plastic BSDF), the high-resolution radiance texture of the primary emitter, and the spatially-varying blend BSDF mask of the secondary worktop.

5.3. Global Illumination and Background Regularization

Because the UGR formulation is fundamentally a ratio between localized source luminance (L_s) and global ambient adaptation (L_b), accurately simulating multiple bounce light transport is paramount. Interestingly, our experiments reveal that whether the background luminance is attached to the computational graph dictates how well the optimization can succeed and profoundly influences the physical plausibility of the resulting materials.

The Influence of Path Length. To investigate the relationship be-



Corresponding Glare Masks

Figure 12: (Top) Sweeping surface roughness α reveals a highly non-convex UGR loss landscape. While initially increasing α diffuses the primary specular reflection and reduces visual discomfort, excessive roughening starves the room of global ambient light (L_b). (Middle/Bottom) As the room progressively darkens ($\alpha = 0.30 \rightarrow 0.40$), the background adaptation luminance L_b drops so severely that secondary, previously comfortable light sources now exceed the relative evaluation boundary. This causes new regions to actively trigger the spatial glare mask, paradoxically spiking the final UGR score.

tween global illumination accuracy and optimization behavior, we conducted an ablation study varying the maximum path length during the inverse rendering process. As demonstrated in Figure 11, optimizing the Index of Refraction (η) of a dielectric surface at short path lengths leads the optimizer to heavily reduce the IOR to remove specular reflections. This occurs because the integrator underestimates the ambient background luminance (L_b).

In order to obtain a better background luminance estimation, we must include longer path lengths. However, simulating higher path lengths increases the computational budget and exacerbates the variance for the gradients which can have adverse effects on the optimization. To balance accuracy and stability, we simulated deeper path lengths during the forward pass to provide accurate ambient lighting, but only record gradients for a small number of path lengths.

Background Gradient Routing and the Alpha Trap. An interesting discovery from our empirical analysis involves the routing of gradients through the L_b term. As established in Figure 12, the UGR objective landscape is nonconvex with respect to spatial roughening (α). Sweeping the roughness parameter reveals that initially increasing α diffuses the glare and lowers the UGR. However,

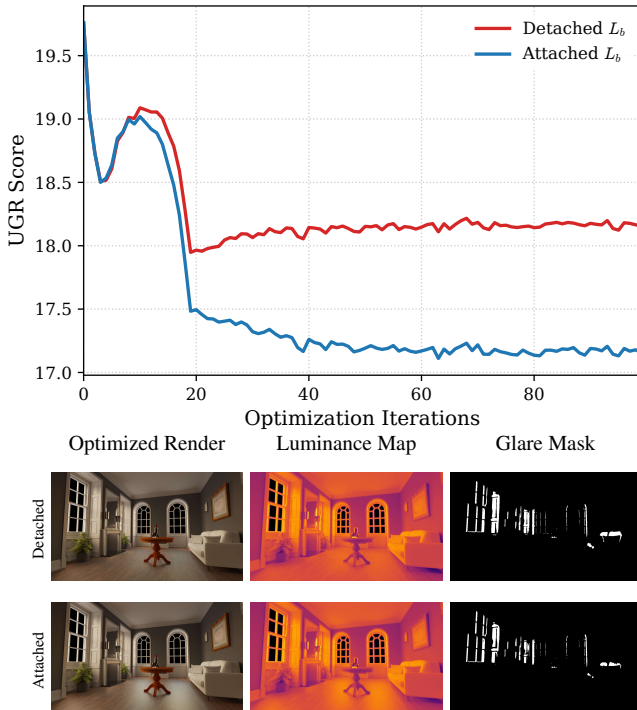


Figure 13: (Top) Convergence plot demonstrating the "over-frosting" trap. When L_b is detached (red), the optimizer greedily minimizes the UGR by darkening the room. When L_b is attached (blue), the loss of ambient light penalizes the UGR score, forcing the optimizer to plateau. (Bottom) Visual results of the optimization. The detached approach results in an unnaturally dark room and an over-blurred highlight. Retaining L_b in the autodiff graph allows the system to autonomously find a "leveled" roughness that diffuses the glare mask while preserving physically plausible global illumination.

if roughening continues, it eventually starves the room of ambient light, causing L_b to drop and the UGR to paradoxically spike, creating a trap for the optimizer.

How the framework handles this nonconvexity depends entirely on whether L_b is attached or detached from the autodiff computational graph. As shown in Figure 13, detaching L_b , treating it as a constant during the backward pass, results in an over frosting trap. The optimizer minimizes the UGR by continually increasing roughness, unnaturally darkening the room and heavily blurring the highlight.

Conversely, keeping L_b attached to the graph introduces a helpful self regulating effect. Attaching the background luminance allows the optimizer to find a leveled area of the roughness where it makes the material rough enough to reduce the highlighted glare but not too rough to starve the room of ambient light. This attached state allows the system to autonomously diffuse the glare mask while preserving physically plausible global illumination.

Contextual Optimization Constraints. While attaching L_b is

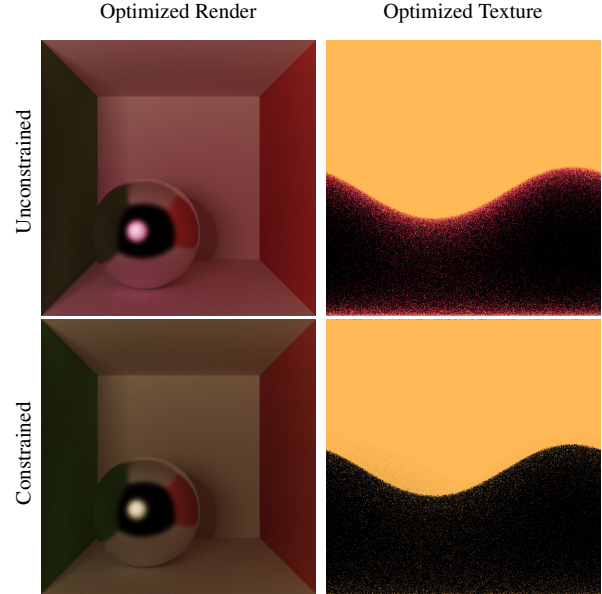


Figure 14: (Top) Unconstrained RGB optimization exploits the psychophysical green-bias of the UGR metric, resulting in unnatural chromaticity shifts. (Bottom) By constraining the latent space to a 1D grayscale mask, the optimizer acts as a neutral density filter, mitigating glare while preserving the original light color.

highly effective for global scene parameters like the living room floor, we observed that care must be taken when applying this regularization to localized texture maps. If L_b is attached while optimizing a spatial roughness texture on a specific object like the sphere in the Cornell Box, the optimizer may exploit the metric by roughening nonglare regions of the texture simply to increase ambient scattering and artificially inflate L_b . Therefore, the decision to attach or detach the background luminance is contextually driven. We detach the background for localized object textures to prevent spatial artifacts, and attach it for broad architectural surfaces to prevent global over darkening.

Relationship to Luminance-Only Optimization. When L_b is detached from the computational graph, the optimizer's effective loss reduces to minimizing $\sum_i L_i m_i$ subject to a fixed threshold, functionally equivalent to a luminance-only objective applied to the glare hotspots. The over-frosting failure demonstrated in Figure 13 is precisely the consequence: without the ambient-preservation signal carried by L_b , the optimizer darkens the room to suppress hotspot luminance with no regard for global illumination. When L_b is attached, the optimization becomes contrast-aware; the optimizer must reduce the source-to-background ratio rather than absolute luminance, which has no analogue in a luminance-only objective.

5.4. Source-Side Constraints and Color Preservation

When optimizing source side emission, we discovered a unique failure mode stemming from the definition of perceptual luminance. Because photometric luminance (L) is derived from radiometric radiance by integrating against the human eye's spectral sensitivity curve, specifically the CIE photopic luminous efficiency

function $V(\lambda)$ [WS00], which heavily favors green wavelengths, an unconstrained optimization will exploit this biological bias.

As shown in Figure 14 (Top), if the optimizer is granted free control over the spectral distribution of an emitter, it will aggressively reduce the wavelengths corresponding to peak human sensitivity to rapidly lower the perceived UGR score, while leaving other wavelengths relatively untouched. This mathematical exploit successfully mitigates the calculated glare, but results in severe unnatural chromaticity shifts in the rendered environment.

To ensure physical and aesthetic plausibility, we found it necessary to restrict the optimizer's latent space when dealing with emission. By mathematically constraining the optimizable parameter to a 1D scalar multiplier applied uniformly across all wavelengths, we force the optimizer to act as a programmable neutral density filter.

6. Limitations and Future Work

As detailed in Section 5.3, the decision to attach or detach the background luminance (L_b) from the computational graph is currently a manual, context-driven choice. While attaching L_b acts as a powerful self-regularizer for global parameters, it can introduce spatial artifacts in localized textures. Future work should investigate automated mechanisms for balancing this trade-off. This could involve developing localized background estimators or adaptive gradient weighting schemes that automatically determine the optimal L_b routing based on the specific parameter being optimized, eliminating the need for manual intervention.

Currently, our framework requires the user to manually select which scene parameters (e.g., a specific floor's roughness or a specific window's IOR) are exposed to the optimizer. While our adjoint method can inherently compute gradients for all scene parameters simultaneously, natively optimizing the entire scene would lead to unacceptable, widespread alterations of the original architectural design. A truly autonomous lighting design tool should instead leverage these global gradients to perform an automated sensitivity analysis, identifying and isolating a sparse set of materials or emitters that contribute most significantly to visual discomfort. More broadly, our framework is an inverse design tool rather than a manufacturing pipeline: it identifies which physical parameters to alter and by how much, but does not guarantee that an optimized η or roughness map corresponds to an off-the-shelf material. While we impose basic physical realizability limits (e.g., $\alpha \geq \alpha_{init}$, scalar emitter multipliers), we do not constrain the optimization to pre-manufactured material catalogues. Future research should focus on integrating these discrete constraints, ensuring the optimizer selects a physically manufacturable intervention rather than simply finding the mathematical path of least resistance.

Our current implementation optimizes static scenes from a single, fixed observer viewpoint under static lighting conditions. Discomfort glare, however, is an inherently dynamic phenomenon; occupants move through spaces, and daylighting conditions change continuously throughout the day. Extending our framework to perform joint optimization across multiple viewpoints or temporal daylighting sequences represents a critical next step for holistic architectural design. While our method is theoretically compatible with such extensions, evaluating complex global illumination

across multiple temporal states simultaneously would require significant advances in differentiable rendering frameworks to remain within practical computational and memory budgets.

7. Conclusion

We introduced a framework that bridges psychophysical lighting metrics and inverse rendering, replacing the discrete binary thresholds of the Unified Glare Rating with a continuous, differentiable proxy that turns UGR from a passive post-process evaluation into an active loss landscape for optimizing physical scene parameters. Evaluating mitigation across source-side emission, boundary-side refraction, and surface-side scattering, we show that naive roughening is bottlenecked by ambient light starvation, and that retaining the background adaptation luminance in the computational graph acts as a critical self-regularizer that navigates this non-convex trap without over-frosting; in contrast, refractive and emitter-masking strategies yield inherently stable landscapes. Together, these provide lighting designers, architectural engineers, and material manufacturers with a physically based automated pipeline for mitigating human visual discomfort in complex rendering environments.

8. Acknowledgements

The authors are grateful to the anonymous reviewers for their many constructive suggestions. We particularly commend the insightful comments and suggestions of the primary reviewer. We acknowledge with gratitude the funding provided for this research by Simon Fraser University as well as a Discovery Grant provided by the Natural Sciences and Engineering Research Council of Canada.

Additionally, we would like to thank Benedikt Bitterli for providing the rendering resources [Bit16] (including the Cornell Box, Grey and White Room, Country Kitchen, and Contemporary Bathroom models), and the Stanford Computer Graphics Laboratory for the Stanford Bunny model from the Stanford 3D Scanning Repository [Lab].

References

- [AAL16] AITTALA M., AILA T., LEHTINEN J.: Reflectance modeling by neural texture synthesis. *ACM Trans. Graph.* 35, 4 (July 2016). URL: <https://doi.org/10.1145/2897824.2925917>, doi: [10.1145/2897824.2925917](https://doi.org/10.1145/2897824.2925917). 3
- [Arc24] ARCHITECTURE ET CLIMAT: Taux d'éblouissement d'inconfort (UGR). <https://energieplus-lesite.be/theories/eclairage12/caracteristiques-des-lumineuses/taux-d-ebouissement-d-inconfort-ugr/>, 2024. Accessed: 2026-04-10. 2
- [Bit16] BITTERLI B.: Rendering resources, 2016. <https://benedikt-bitterli.me/resources/>. 11
- [Bla13] BŁASZCZAK U. J.: Method for evaluating discomfort glare based on the analysis of a digital image of an illuminated interior. *Metrology and Measurement Systems* 20, 4 (2013), 623–634. 3
- [CCDP15] CARLUCCI S., CAUSONE F., DE ROSA F., PAGLIANO L.: A review of indices for assessing visual comfort with a view to their use in optimization processes to support building integrated design. *Renewable and Sustainable Energy Reviews* 47 (2015), 1016–1033. URL: <https://www.sciencedirect.com/science/article/pii/S1364032115002154>, doi: <https://doi.org/10.1016/j.rser.2015.03.062>. 3

- [CIE95] CIE: *Discomfort Glare in Interior Lighting*. Tech. Rep. CIE 117-1995, International Commission on Illumination, 1995. doi:10.25039/TR.117.1995. 2, 3
- [DAD*18] DESCHAINTE V., AITTALA M., DURAND F., DRETAKIS G., BOUSSEAU A.: Single-image svbrdf capture with a rendering-aware deep network. *ACM Trans. Graph.* 37, 4 (July 2018). URL: <https://doi.org/10.1145/3197517.3201378>, doi:10.1145/3197517.3201378. 3
- [DIA26] DIAL GMBH: DIALux evo: Lighting Design Software. <https://www.dialux.com>, 2026. Accessed: 2026-04-14. 2
- [DLW*22] DENG X., LUAN F., WALTER B., BALA K., MARSCHNER S.: Reconstructing translucent objects using differentiable rendering. In *ACM SIGGRAPH 2022 Conference Proceedings* (New York, NY, USA, 2022), SIGGRAPH '22, Association for Computing Machinery. URL: <https://doi.org/10.1145/3528233.3530714>, doi:10.1145/3528233.3530714. 3
- [DLZ*22] DAI Y., LI C., ZHOU S., FENG R., LOY C. C.: Flare7k: A phenomenological nighttime flare removal dataset. *Advances in Neural Information Processing Systems* 35 (2022), 3926–3937. 3
- [EGS*26] EKIM OLSSON Z., GOVÉN T., SELVANDER M., LAIKE T., SJÖBERG K.: Perception of discomfort glare in relation to the cie unified glare rating scale (ugr) and the impact of ambient light on both ugr and subjective glare indexes (sgis). *Indoor Environments* 3, 1 (2026), 100149. URL: <https://www.sciencedirect.com/science/article/pii/S2950362026000019>, doi:<https://doi.org/10.1016/j.indenv.2026.100149>. 3
- [Eur21] EUROPEAN COMMITTEE FOR STANDARDIZATION (CEN): *EN 12464-1:2021 Light and lighting - Lighting of work places - Part 1: Indoor work places*. Tech. rep., CEN, Brussels, Belgium, 2021. 2, 3
- [GHZ18] GUO Y., HAŠAN M., ZHAO S.: Position-free monte carlo simulation for arbitrary layered bsdfs. *ACM Trans. Graph.* 37, 6 (Dec. 2018). URL: <https://doi.org/10.1145/3272127.3275053>, doi:10.1145/3272127.3275053. 3
- [GKM20] GOEL S., KANAZAWA A., MALIK J.: Shape and viewpoint without keypoints. In *European Conference on Computer Vision* (2020), Springer, pp. 88–104. 3
- [GLD*19] GAO D., LI X., DONG Y., PEERS P., XU K., TONG X.: Deep inverse rendering for high-resolution svbrdf estimation from an arbitrary number of images. *ACM Trans. Graph.* 38, 4 (July 2019). URL: <https://doi.org/10.1145/3306346.3323042>, doi:10.1145/3306346.3323042. 3
- [Han16] HANSEN N.: The CMA evolution strategy: A tutorial. *CoRR abs/1604.00772* (2016). URL: <http://arxiv.org/abs/1604.00772>, arXiv:1604.00772. 9
- [JdJM14] JAKOB W., D'EON E., JAKOB O., MARSCHNER S.: A comprehensive framework for rendering layered materials. *ACM Trans. Graph.* 33, 4 (July 2014). URL: <https://doi.org/10.1145/2601097.2601139>, doi:10.1145/2601097.2601139. 3
- [Jon17] JONES N. L.: *Validated interactive daylighting analysis for architectural design*. PhD thesis, Massachusetts Institute of Technology, 2017. 3
- [JSR*22] JAKOB W., SPEIERER S., ROUSSEL N., NIMIER-DAVID M., VICINI D., ZELTNER T., NICOLET B., CRESPO M., LEROY V., ZHANG Z.: Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>. 5
- [JSRV22] JAKOB W., SPEIERER S., ROUSSEL N., VICINI D.: Dr.jit: A just-in-time compiler for differentiable rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (July 2022). doi:10.1145/3528223.3530099. 5
- [KB14] KINGMA D. P., BA J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014). 5
- [KMN*04] KAKIMOTO M., MATSUOKA K., NISHITA T., NAEMURA T., HARASHIMA H.: Glare generation based on wave optics. In *12th Pacific Conference on Computer Graphics and Applications, 2004. PG 2004. Proceedings.* (2004), pp. 133–140. doi:10.1109/PCCGA.2004.1348343. 4
- [KUH18] KATO H., USHIKU Y., HARADA T.: Neural 3d mesh renderer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018). 3
- [Lab] LABORATORY S. C. G.: The stanford 3d scanning repository. <http://graphics.stanford.edu/data/3Dscanrep/>. 11
- [LADL18] LI T.-M., AITTALA M., DURAND F., LEHTINEN J.: Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph.* 37, 6 (Dec. 2018). URL: <https://doi.org/10.1145/3272127.3275109>, doi:10.1145/3272127.3275109. 3
- [LB14] LOPER M. M., BLACK M. J.: Opendr: An approximate differentiable renderer. In *European Conference on Computer Vision* (2014), Springer, pp. 154–169. 3
- [LHEN*24] LIPP L., HAHN D., ECORMIER-NOCCA P., RIST F., WIMMER M.: View-independent adjoint light tracing for lighting design optimization. *ACM Trans. Graph.* 43, 3 (May 2024). URL: <https://doi.org/10.1145/3662180>, doi:10.1145/3662180. 3
- [LHJ19] LOUBET G., HOLZSCHUCH N., JAKOB W.: Reparameterizing discontinuous integrands for differentiable rendering. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (Dec. 2019). doi:10.1145/3355089.3356510. 3
- [LLCL19] LIU S., LI T., CHEN W., LI H.: Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2019). 3
- [MBGJ22] MISSO Z., BITTERLI B., GEORGIEV I., JAROSZ W.: Unbiased and consistent rendering using biased estimators. *ACM Transactions on Graphics (ToG)* 41, 4 (July 2022), 1–13. doi:10/gqjn66. 6
- [NDDJK21] NIMIER-DAVID M., DONG Z., JAKOB W., KAPLANYAN A.: Material and Lighting Reconstruction for Complex Indoor Scenes with Texture-space Differentiable Rendering. In *Eurographics Symposium on Rendering - DL-only Track* (2021), Bousseau A., McGuire M., (Eds.), The Eurographics Association. doi:10.2312/sr.20211292. 3
- [NDMKJ22] NIMIER-DAVID M., MÜLLER T., KELLER A., JAKOB W.: Unbiased inverse volume rendering with differential trackers. *ACM Trans. Graph.* 41, 4 (July 2022), 44:1–44:20. URL: <https://doi.org/10.1145/3528223.3530073>, doi:10.1145/3528223.3530073. 3
- [NRN*23] NICOLET B., ROUSSELLE F., NOVÁK J., KELLER A., JAKOB W., MÜLLER T.: Recursive control variates for inverse rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)* 42, 4 (Aug. 2023). doi:10.1145/3592139. 6
- [Pel90] PELI E.: Contrast in complex images. *J. Opt. Soc. Am. A* 7, 10 (Oct 1990), 2032–2040. URL: <https://opg.optica.org/josaa/abstract.cfm?URI=josaa-7-10-2032>, doi:10.1364/JOSAA.7.002032. 4
- [Pow64] POWELL M. J. D.: An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The Computer Journal* 7, 2 (01 1964), 155–162. URL: <https://doi.org/10.1093/comjnl/7.2.155>, arXiv:<https://academic.oup.com/comjnl/article-pdf/7/2/155/959784/070155.pdf>, doi:10.1093/comjnl/7.2.155. 9
- [PWB18] PIERSON C., WIENOLD J., BODART M.: Review of factors influencing discomfort glare perception from daylight. *LEUKOS* 14, 3 (2018), 111–148. URL: <https://doi.org/10.1080/15502724.2018.1428617>, arXiv:<https://doi.org/10.1080/15502724.2018.1428617>, doi:10.1080/15502724.2018.1428617. 3
- [Rei02] REINHARD E.: Parameter estimation for photographic tone reproduction. *Journal of graphics tools* 7, 1 (2002), 45–51. 5
- [SSZG95] SPENCER G., SHIRLEY P., ZIMMERMAN K., GREENBERG D. P.: Physically-based glare effects for digital images. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive*

- Techniques* (New York, NY, USA, 1995), SIGGRAPH '95, Association for Computing Machinery, p. 325–334. URL: <https://doi.org/10.1145/218380.218466>, doi:10.1145/218380.218466. 4
- [SWF*21] SUN Q., WANG C., FU Q., DUN X., HEIDRICH W.: End-to-end complex lens design with differentiable ray tracing. *ACM Trans. Graph.* 40, 4 (July 2021). URL: <https://doi.org/10.1145/3450626.3459674>, doi:10.1145/3450626.3459674. 3
- [TMM*21] TSENG E., MOSLEH A., MANNAN F., ST-ARNAUD K., SHARMA A., PENG Y., BRAUN A., NOWROUZEZAHRAI D., LALONDE J.-F., HEIDE F.: Differentiable compound optics and processing pipeline optimization for end-to-end camera design. *ACM Trans. Graph.* 40, 2 (June 2021). URL: <https://doi.org/10.1145/3446791>, doi:10.1145/3446791. 3
- [TVB*25] TEH A., VICINI D., BICKEL B., GKIOULEKAS I., O'TOOLE M.: Automated design of compound lenses with discrete-continuous optimization. SA Conference Papers '25, Association for Computing Machinery. doi:10.1145/3757377.3763850. 3
- [Vea98] VEACH E.: *Robust Monte Carlo methods for light transport simulation*. Stanford University, 1998. 4
- [VJS21] VICINI D., SPEIERER S., JAKOB W.: Path replay backpropagation: Differentiating light paths using constant memory and linear time. *Transactions on Graphics (Proceedings of SIGGRAPH)* 40, 4 (Aug. 2021), 108:1–108:14. doi:10.1145/3450626.3459804. 3, 5
- [War94] WARD G. J.: The radiance lighting simulation and rendering system. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1994), SIGGRAPH '94, Association for Computing Machinery, p. 459–472. URL: <https://doi.org/10.1145/192161.192286>, doi:10.1145/192161.192286. 3
- [WC06] WIENOLD J., CHRISTOFFERSEN J.: Evaluation methods and development of a new glare prediction model for daylight environments with the use of ccd cameras. *Energy and Buildings* 38, 7 (2006), 743–757. Special Issue on Daylighting Buildings. URL: <https://www.sciencedirect.com/science/article/pii/S0378778806000715>, doi:https://doi.org/10.1016/j.enbuild.2006.03.017. 3
- [WHX*21] WU Y., HE Q., XUE T., GARG R., CHEN J., VEERARAGHAVAN A., BARRON J. T.: How to train neural networks for flare removal. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (October 2021), pp. 2239–2247. 3
- [WRG*25] WEIER P., RIVIERE J., GUSEINOV R., GARBIN S., SLUSALLEK P., BICKEL B., BEELER T., VICINI D.: Practical inverse rendering of textured and translucent appearance. *Transactions on Graphics (Proceedings of SIGGRAPH)* 44, 4 (Aug. 2025). doi:10.1145/3730855. 3
- [WS00] WYSZECKI G., STILES W. S.: *Color science: concepts and methods, quantitative data and formulae*. John wiley & sons, 2000. 11
- [ZRJ23] ZHANG Z., ROUSSEL N., JAKOB W.: Projective sampling for differentiable rendering of geometry. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 42, 6 (Dec. 2023). doi:10.1145/3618385. 3