

Scene Parameter Saliency via Differentiable Light Transport

Linus Beresna
linas_beresna@sfu.ca

Eugene Fiume
eugene_fiume@sfu.ca

Simon Fraser University
Burnaby, Canada

Abstract

Gradient-based saliency methods reveal which input features most influence a neural network’s output, and are a standard tool for model interpretability. We observe that differentiable renderers, which are conventionally used for parameter optimisation, produce an analogous form of saliency: given any scalar metric evaluated on a rendered image, a single reverse-mode differentiation pass yields per-parameter gradients that identify which scene elements most influence the metric. We call these gradient fields metric saliency maps. Unlike neural saliency, which propagates attribution through learned weights, metric saliency propagates through the image formation process itself, including multi-bounce light transport, capturing parameter dependencies that are semi-opaque to manual inspection. We compute metric saliency maps for qualitatively different objectives: psychovisual glare indices, mean scene luminance, and neural perceptual scores. The saliency rankings differ substantially across metrics for the same scene, with parameters that dominate one objective being negligible for another. The saliency map is specific to the metric, not an intrinsic property of the scene. Our results suggest that differentiable renderers produce derivative images that are as informative for scene understanding as the primal images they were designed to generate.

1 Introduction

A central concern of computer vision is understanding how images relate to the scenes that produce them. Classical physics-based vision tackled this relationship directly: shape from shading [14], intrinsic image decomposition [3], and photometric stereo [42] all reason about the physical process of image formation to recover scene properties from pixel observations. These methods succeed because the mapping from scene to image, governed by geometry, reflectance, and illumination, is structured and (in principle) invertible.

Differentiable rendering [16, 20, 21] has revived this physics-based perspective with modern computational tools. By implementing the rendering pipeline as an automatically differentiable program, one obtains gradients of any image-derived quantity with respect to scene parameters: geometry, materials, lighting, and camera. These gradients have been used extensively for inverse rendering [21, 29], pose estimation [38], shape reconstruction [8, 23], and scene generation [43]. In all of these applications, the gradient serves a single purpose: it is fed into an optimiser that iteratively updates scene parameters to minimise a loss function.



Figure 1: **Expanding the diagnostic toolkit for architectural lighting.** While conventional analysis provides (b) glare source identification and (c) luminance maps to evaluate discomfort glare, it leaves the designer to guess the root physical causes. By differentiating the psychovisual unified glare rating through the light transport simulation, we obtain (d) a metric saliency map. This highlights the specific scene materials and geometry driving the metric, explicitly guiding design decisions.

We observe that these gradients carry information that is valuable *independently of any optimisation*. Given a scalar metric M evaluated on a rendered image and a differentiable renderer $\mathcal{R}$ parameterised by scene parameters θ , the gradient $\partial M(\mathcal{R}(\theta))/\partial \theta$ is a vector that assigns to every scene parameter a score reflecting how much it influences the metric. This is a saliency map over the parameter space of the scene.

The analogy to neural network interpretability is direct. Gradient-based saliency methods such as Grad-CAM [33] and vanilla gradient attribution [34] reveal which input features most influence a network’s prediction by inspecting $\partial y/\partial \mathbf{x}$. Our metric saliency maps do the same, but the differentiable function is a physically-based renderer rather than a learned model, and the “input features” are scene parameters rather than pixel values. Where Grad-CAM asks *which pixels matter for this classification?*, we ask *which scene parameters matter for this metric?*

The distinction between neural and physical saliency is more than cosmetic. Neural saliency propagates through learned weights whose structure is opaque by nature. Metric saliency propagates through the image formation process itself, including multi-bounce light transport, where the dependencies have a physical interpretation. A wall surface three bounces removed from a light source may, through interreflection, have a larger effect on a glare metric than the luminaire itself. The gradient encodes this dependency exactly, in a single reverse-mode differentiation pass, without requiring the repeated evaluations of parametric sweeps [36] or finite-difference probing that characterise traditional sensitivity analysis in architectural and lighting design [13, 31].

We demonstrate metric saliency maps across qualitatively different objectives, including mean scene luminance, neural perceptual similarity scores, and psychovisual glare ratings. Figure 1 illustrates this approach for architectural lighting by explicitly mapping the root physical causes of discomfort glare. The saliency rankings differ substantially across metrics for the same scene. Parameters that dominate one objective are negligible for another, confirming that the saliency map is specific to the metric rather than being an intrinsic property of the scene. This metric specificity is the practical value of the tool because it tells a designer not just *what to change*, but *what matters for this particular criterion*.

Our contributions are:

1. We identify metric saliency maps, the per-parameter gradients of a scalar objective

through a differentiable renderer, as an interpretability tool analogous to gradient-based attribution in neural networks, and distinct from the optimisation role these gradients conventionally serve.

2. We demonstrate that a single reverse-mode differentiation pass through physically-based rendering exposes non-obvious parameter sensitivities across diverse metrics, including cases where the most influential element is spatially and optically distant from the measured effect.
3. We show that saliency rankings are metric specific: changing the objective function over the same scene produces qualitatively different maps, providing diagnostic information that neither the primal image nor any single metric value can convey alone.

2 Related Work

Our work sits at the intersection of differentiable rendering, gradient-based interpretability, and physics-based scene analysis. We survey each area and identify the gap that metric saliency maps address.

Differentiable rendering. The ability to compute image gradients with respect to scene parameters has been developed across rendering paradigms. For rasterisation, Neural Mesh Renderer [18], SoftRas [23], DIB-R [8], and nvdiffrast [20] introduce differentiable approximations to the discrete visibility function. For Monte Carlo path tracing, Li *et al.* [21] proposed edge sampling for boundary gradients, Loubet *et al.* [25] introduced reparameterisation-based estimators, and Nimier-David *et al.* [30] developed radiative backpropagation within Mitsuba 2. Vicini *et al.* [29] and the Path Replay Backpropagation (PRB) method [39] further reduced the memory and variance costs of differentiating through global illumination. Mitsuba 3 [16] consolidates many of these techniques into a single framework. Neural scene representations, including NeRF [26] and 3D Gaussian Splatting [19], are also inherently differentiable and have been used for view synthesis and reconstruction [2, 10]. DiffCSG [43] extends differentiability to constructive solid geometry and notably includes per-parameter gradient visualisations as part of its pipeline. Across all of these systems, gradients are treated as an intermediate quantity consumed by an optimiser. We propose to treat them as a final output consumed by a human.

Gradient-based saliency and attribution. In the neural network interpretability literature, gradient-based methods are a primary tool for explaining model decisions. Simonyan *et al.* [34] introduced vanilla gradient saliency, computing $\partial y / \partial \mathbf{x}$ to highlight input pixels that most affect a classifier’s output. Grad-CAM [33] and Grad-CAM++ [7] refine this by weighting feature map activations with gradients, producing class-discriminative heatmaps. These methods have been extended to object detection, segmentation, and visual question answering [6]. Integrated Gradients [37] and SmoothGrad [35] address gradient noise and saturation. The general principle, that the derivative of a differentiable function with respect to its inputs constitutes an attribution over those inputs, is well established for neural networks. We apply this same principle to a different differentiable function: a physically-based renderer. The “inputs” become scene parameters (materials, geometry, lighting) rather than pixel values, and the “output” becomes an arbitrary scalar metric rather than a class logit.

Physics-based vision and scene understanding. Our work connects to a tradition in computer vision that reasons about image formation through physical models. Shape from shading [14], intrinsic image decomposition [1, 3], and photometric stereo [42] all invert the rendering process to recover scene properties. More recently, analysis-by-synthesis approaches use differentiable rendering to optimise scene explanations that match observed images [44]. MRD [4] uses physically-based differentiable rendering to probe vision models by finding scene-parameter metamers, physically different scenes that produce identical model activations, offering an interpretability tool grounded in 3D scene descriptions rather than pixel space. Our approach shares MRD’s motivation of using differentiable rendering for understanding rather than reconstruction, but operates through gradient inspection rather than optimisation, and targets scalar scene metrics rather than neural network representations.

Sensitivity analysis in design. In architectural lighting and building performance evaluation, understanding how design parameters affect metrics such as Unified Glare Rating (UGR) [15], Daylight Glare Probability (DGP) [41], spatial daylight autonomy [24], and useful daylight illuminance [28] is essential for informed design. Traditional approaches rely on parametric sweeps [13], variance-based sensitivity analysis (Sobol indices) [36], or Morris screening [27], all of which require many independent simulation evaluations and scale poorly with the number of parameters. Gradient-based lighting optimisation has been explored using differentiable rendering for luminaire placement [22, 32], but the gradients serve an optimiser rather than a human. Our metric saliency maps provide a local sensitivity characterisation at the cost of a single backward pass.

3 Method

The method has three stages: render an image through a differentiable pipeline, evaluate a scalar metric on that image, and backpropagate through both the metric and the renderer to obtain per-parameter gradients. These gradients are then mapped back to image space for visualisation. We describe each stage and the metrics we evaluate.

3.1 Differentiable Image Formation

Let $\mathcal{R}$ be a renderer that computes an image $I = \mathcal{R}(\theta)$ from scene parameters θ . The parameter vector θ may include surface reflectances, roughness values, indices of refraction, emitter intensities, and geometric quantities. We use a physically-based Monte Carlo path tracer with multi-bounce global illumination, so that each pixel is a function of parameters potentially distributed across the entire scene through interreflection. Given a scalar metric M evaluated on the rendered image, the key requirement is that the full composition $M(\mathcal{R}(\theta))$ is differentiable: both the renderer $\mathcal{R}$ and the metric M must admit gradient computation, so that we can obtain $\partial M / \partial \theta$ via reverse-mode automatic differentiation through the metric and the light transport simulation jointly. We use Mitsuba 3 [16] with PRB [39], though any differentiable renderer would serve the same role.

3.2 Metric Saliency Computation

Given a scalar measurement operator $M : \mathbb{R}^D \rightarrow \mathbb{R}$ that evaluates a property of the light transport simulation, let $I = \mathcal{R}(\theta) \in \mathbb{R}^D$ represent a generalised sensor response. By ex-

panding the traditional image formation model, I can encode standard visible images (where $D = H \times W \times C$) as well as non-visible or abstract radiometric properties. The composition $M(\mathcal{R}(\boldsymbol{\theta}))$ thus defines a scalar function of the scene parameters. Applying the chain rule yields the parameter-space gradient:

$$\mathbf{g} = \frac{\partial M(\mathcal{R}(\boldsymbol{\theta}))}{\partial \boldsymbol{\theta}} = \frac{\partial M}{\partial I} \frac{\partial I}{\partial \boldsymbol{\theta}}. \quad (1)$$

The chain rule separates this into two terms: the sensitivity of the metric to the image, which depends on how M is defined, and the sensitivity of the image to scene parameters, which depends on the light transport. Neither term needs to be computed or inspected independently; the product is obtained automatically by backpropagation.

Each component $g_i = \partial M / \partial \theta_i$ has a direct interpretation: it measures how much M would change per unit perturbation of θ_i , accounting for all direct and indirect light paths. The magnitude $|g_i|$ indicates importance, the sign indicates direction. We define the *saliency* of parameter θ_i as $|g_i|$, and for vector-valued parameters (*e.g.*, an RGB reflectance) we take the ℓ_2 norm of the gradient components.

3.3 Image-Space Visualisation

As illustrated in Figure 2, the gradient vector $\mathbf{g}$ lives in parameter space: each component corresponds to a named scene parameter, not a pixel. To produce an image-space visualisation, we need to associate each pixel with the scene parameter(s) visible at that pixel. We achieve this by tracing primary rays from the camera and, at each intersection, identifying which scene object (and therefore which material parameters) the ray hit. This produces a per-pixel material identity map, analogous to an instance segmentation mask in computer vision or a cryptomatte pass in visual effects compositing [9]. Each pixel is then coloured by the saliency of the material visible at that location, producing a heatmap that can be compared side-by-side with the beauty render.

For scalar parameters that apply uniformly to an entire object (*e.g.*, a single roughness value for a surface), every pixel showing that object receives the same saliency value. For spatially-varying parameters such as texture maps, the gradient is itself spatially varying and can be mapped directly onto the surface. In this work we focus on uniform material parameters; spatially-varying saliency fields are left to future work.

We normalise saliency values across all parameters to $[0, 1]$ for visualisation, so that the map highlights *relative* importance within a scene configuration. The resulting image is

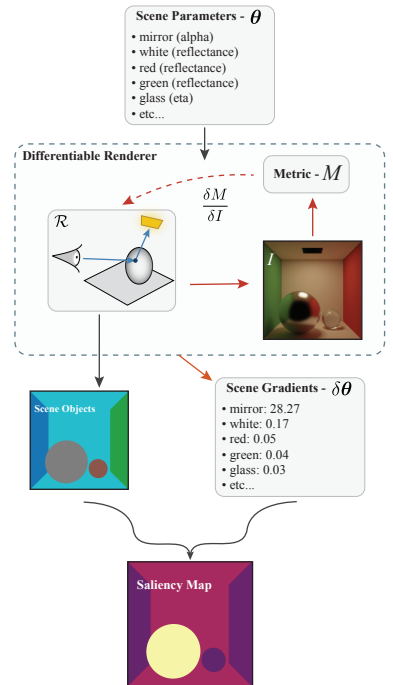


Figure 2: **Image-Space Saliency Mapping.** Per-pixel material identities (Scene Objects) are coloured by their corresponding parameter sensitivities (Scene Gradients) to produce the final metric-specific heatmap.

metric specific: the same scene rendered under the same conditions will produce different saliency maps for different metrics.

3.4 Choice of Metrics

To demonstrate generality, we evaluate metrics from three qualitatively different families.

Mean luminance. This is the simplest metric: the average luminance of the rendered image, computed as a weighted sum of RGB channels using standard photometric coefficients. This serves as a baseline. The saliency map should highlight surfaces with large visible area and high reflectance, which are the parameters that most affect overall image brightness.

Unified Glare Rating (UGR). A psychovisual metric standardised by the CIE [15] for evaluating discomfort glare in architectural lighting. UGR aggregates the luminance, solid angle, and positional weight of glare sources against a background luminance. It is non-linear, involves a logarithm and a ratio, and its dependence on scene parameters is mediated by the glare/background decomposition and the Guth position index [11].

Neural perceptual score. To demonstrate that the approach extends beyond hand-crafted metrics, we evaluate a metric defined by a neural network. Specifically, we backpropagate from the pre-softmax class logit of a pre-trained ResNet-50 image classifier [12]. This case is notable because the gradient must flow through both the neural network latent space and the physical light transport simulation, and the resulting saliency map reflects the joint sensitivity of both systems.

4 Results

Experimental Setup. We render standard interior scenes including a dining room, kitchen, and living room using PRB [39] with Mitsuba 3 [16] and DrJit [17]. We trace paths up to a maximum depth of six bounces with 1024 samples per pixel. For each scene, we track all differentiable material parameters and evaluate metric saliency for mean luminance, discomfort glare, and neural classification logits.

4.1 Metric Specificity

Psychovisual metrics like UGR depend on complex spatial relationships between the observer and the environment, whereas simple metrics like mean luminance measure global scene energy. Figure 3 demonstrates how our pipeline adapts to these fundamentally different objectives across both a living room and a kitchen environment.

By default, the differentiable engine produces a raw mathematical list of parameter sensitivities. While plotting these gradients as bar charts (top row) yields precise statistical rankings, isolating this data from its spatial context makes it difficult for a human designer to parse. By aggregating these parameter gradients and mapping them directly back onto the 3D scene geometry, we produce an intuitive visual heatmap of metric sensitivity (bottom row).

The UGR saliency maps explicitly isolate the specific reflective surfaces and light sources that maximise discomfort glare for the given camera pose, such as the mirror in the living room or the metallic utensils in the kitchen. This provides highly localised, actionable targets for a designer to mitigate visual discomfort. To verify that this is not simply a generic scene




Figure 3: Visualizing Metric Specificity. While extracting gradients provides precise ranking data for scene parameters (top row), mapping this raw data back into the 3D scene (bottom row) allows for intuitive spatial visualisation. Across multiple scenes, the UGR evaluation isolates highly localised reflective surfaces driving discomfort glare (e.g., the mirror or metallic utensils). In contrast, the mean luminance evaluation shifts focus entirely to large surfaces driving global illumination.

feature map, we compare it against the mean luminance objective. As shown in Figure 3, the parameters dictating overall scene brightness differ from those causing localised glare. The mean luminance maps correctly shift emphasis to large geometric areas, such as the diffuse walls, ceilings, and floors, confirming that our spatial projection accurately visualises the underlying physical drivers of the chosen metric regardless of the scene’s architecture.

4.2 Neural and Perceptual Scores

We demonstrate that this pipeline handles high dimensional learned spaces by evaluating a semantic classifier. Figure 4 illustrates the metric saliency for the dining room scene when back-propagating from the ResNet-50 logit for the folding chair class (class 559).

The gradient flows from the neural network latent space directly into the physical light transport simulation, capturing complex optical dependencies. Notably, the left wall of this scene consists of a highly reflective chrome material. The resulting saliency map explicitly visualises how this specular surface influences the classification score. It highlights that the neural network relies not only on direct geometric edges, but also on the reflected features and specific material roughness values to conceptually recognise the objects in the scene.



(a) Beauty Render



(b) Saliency Map

Figure 4: Semantic attribution mapped into 3D parameter space. The saliency map shows which materials act as the cues for the classification logit.

4.3 Validation via Independent Parameter Ablation

To verify that our computed gradients accurately predict the physical sensitivities of the objective functions, we perform an independent ablation test on both the non-linear UGR metric and the global mean luminance metric (Figure 5). Because these metrics represent complex optical dependencies, a computed gradient provides a strictly local first-order approximation. To rigorously validate this without colliding with physical material bounds (e.g., surface roughness clamping at 1.0), we evaluate the metric’s actual response to a small, independent gradient descent step (learning rate = 0.01) for each of the top-ranked parameters.




Figure 5: **Independent Ablation Validation.** Parameters are ranked strictly by computed gradient magnitude (y-axis). Taking a small gradient descent step yields an absolute change in the objective metric (x-axis). The UGR evaluation (a) perfectly mirrors the computed saliency rank, whilst the mean luminance evaluation (b) highlights the noise floor challenges of global stochastic metrics.

As shown in Figure 5a, the absolute measured impact on the localised UGR metric perfectly follows our computed saliency ranking in a monotonically decreasing curve. This confirms that a single reverse-mode evaluation in our differentiable framework successfully isolates and ranks the true physical drivers of discomfort glare, effectively replacing the need for dozens of expensive finite-difference simulations.

When evaluating global metrics like mean luminance (Figure 5b), the highest-ranked parameter correctly dominates the physical impact. However, at lower ranks, we observe minor deviations from strict monotonicity. This highlights an inherent challenge in differentiable path tracing: as absolute metric changes fall below 10^{-4} units, the evaluation plunges into the stochastic noise floor of the renderer. Furthermore, applying a uniform gradient step across fundamentally different physical domains, such as diffuse reflectance versus the index of refraction (η), yields disparate macroscopic behaviors due to the non-linear nature of Fresnel equations. We leave the exploration of parameter-specific adaptive step scaling and second-order Hessian approximations for future work.

5 Discussion

What the maps reveal. The central finding of this work is that metric saliency maps expose parameter-metric relationships that are not directly accessible through inspection of the rendered image or the metric value alone. The primal image shows what the scene looks like; the metric returns a single number summarising some property of that image; but neither reveals *why* the number is what it is, or *what* would change it most effectively. The

saliency map fills this gap. The fact that saliency rankings differ substantially across metrics for the same scene confirms that this information is genuinely metric-specific and cannot be recovered from any single rendering or evaluation.

Limitations. Several limitations should be noted. First, Monte Carlo gradient estimators inherit the variance of the underlying path tracer. At low sample counts, saliency values for individual parameters may be noisy. In our experiments, moderate sample counts (64–256 spp) produced stable rankings, but fine-grained spatially-varying saliency fields would require higher sample budgets or variance reduction techniques. Second, the approach requires the metric M to be differentiable. Metrics with hard thresholds or discrete decisions (*e.g.*, binary glare source detection) must be relaxed into smooth approximations. For instance, evaluating UGR within our framework required us to reformulate standard binary visibility checks into smooth sigmoid relaxations. The choice of relaxation can affect the resulting saliency map, and the relationship between the relaxed and the original metric deserves further investigation. Third, gradient magnitude measures local, first-order sensitivity. It does not capture interactions between parameters (second-order effects), nor does it account for parameter constraints that may prevent a suggested change in practice. A parameter with high saliency may not be actionable if it is physically fixed. Higher-order methods, such as Hessian-based analysis [40], could provide richer sensitivity information at greater computational cost.

Scope of the rendering model. Our demonstration uses a physically-based Monte Carlo path tracer, but the principle applies to any differentiable image formation process. Rasterisation-based differentiable renderers [20, 23], neural radiance fields [26], and Gaussian splatting [19] all produce parameter gradients that could be interpreted as saliency maps. The multi-bounce setting we emphasise is where the approach is most valuable, because it is precisely where human intuition about parameter influence breaks down, but simpler rendering models would work equally well for simpler scenes.

Future work. This paper establishes the concept of metric saliency maps and demonstrates their utility. Several directions remain open. The visualisation of gradient information in scene and image space is itself a research question: signed maps, vector-field representations, and interactive exploration interfaces could provide richer and more actionable information than the magnitude heatmaps we present here. Real-time or interactive computation of saliency maps would enable designers to explore parameter sensitivities dynamically as they modify a scene, motivating work on efficient differentiable rendering. Finally, coupling saliency maps with optimisation, using the map to guide which parameters to optimise and which to hold fixed, could combine the interpretability of our approach with the automation of existing inverse rendering pipelines.

6 Conclusion

We have shown that differentiable renderers produce gradient information that is valuable as a direct output, not only as an intermediate quantity for optimisation. By backpropagating a scalar metric through the rendering pipeline, we obtain metric saliency maps that reveal which scene parameters most influence the metric. These maps are inexpensive (a single

backward pass), general (applicable to any differentiable metric and renderer), and informative (exposing non-obvious dependencies mediated by multi-bounce light transport). The saliency rankings are metric-specific: the same scene produces qualitatively different maps under different objectives. This metric-specificity is the practical contribution, as it provides diagnostic information that neither the rendered image nor the metric value can convey alone. Rendering has always served a single purpose: making the invisible visible, turning the mathematics of light transport into images that humans can see and reason about. Differentiable rendering computes derivatives of that same process, but to date these derivatives have been consumed exclusively by optimisation algorithms, never treated as a primary interpretability object. We suggest that derivative images deserve the same status as the primal images they accompany: they are structured, interpretable, and made for human eyes.

7 Acknowledgements

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 [5] for providing the rendering resources (including the Cornell Box, Grey and White Room, Country Kitchen, and The Breakfast Room models).

References

- [1] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence*, 37(8):1670–1687, 2014.
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5855–5864, 2021.
- [3] Harry Barrow, J Tenenbaum, A Hanson, and E Riseman. Recovering intrinsic scene characteristics. *Comput. vis. syst*, 2(3-26):2, 1978.
- [4] Benjamin Beilharz and Thomas SA Wallis. Mrd: Using physically based differentiable rendering to probe vision models for 3d scene understanding. *arXiv preprint arXiv:2512.12307*, 2025.
- [5] Benedikt Bitterli. Rendering resources, 2016. <https://benedikt-bitterli.me/resources/>.
- [6] Angie Boggust, Harini Suresh, Hendrik Strobelt, John Gutttag, and Arvind Satyanarayan. Saliency cards: A framework to characterize and compare saliency methods. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 285–296, 2023.

-
- [7] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE winter conference on applications of computer vision (WACV)*, pages 839–847. IEEE, 2018.
 - [8] Wenzheng Chen, Huan Ling, Jun Gao, Edward Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. *Advances in neural information processing systems*, 32, 2019.
 - [9] Jonah Friedman and Andrew C Jones. Fully automatic id mattes with support for motion blur and transparency. In *ACM SIGGRAPH 2015 Posters*, pages 1–1. 2015.
 - [10] Jian Gao, Chun Gu, Youtian Lin, Zhihao Li, Hao Zhu, Xun Cao, Li Zhang, and Yao Yao. Relightable 3d gaussians: Realistic point cloud relighting with brdf decomposition and ray tracing. In *European Conference on Computer Vision*, pages 73–89. Springer, 2024.
 - [11] Sylvester K Guth. Computing visual comfort ratings for a specific interior lighting installation. *Illuminating Engineering*, 61(10):634, 1966.
 - [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
 - [13] Lisa Heschong. Daylighting and human performance. *ASHRAE journal*, 44(6):65–67, 2002.
 - [14] Berthold K. P. Horn and Michael J. Brooks, editors. *Shape from shading*. MIT Press, Cambridge, MA, USA, 1989. ISBN 0262081830.
 - [15] International Commission on Illumination. Discomfort glare in interior lighting. Technical Report CIE 117-1995, CIE, Vienna, Austria, 1995. ISBN 978-3-900734-70-1.
 - [16] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, Merlin Nimier-David, Delio Vicini, Tizian Zeltner, Baptiste Nicolet, Miguel Crespo, Vincent Leroy, and Ziyi Zhang. Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>.
 - [17] Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, and Delio Vicini. 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.
 - [18] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3907–3916, 2018.
 - [19] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis, et al. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4): 139–1, 2023.
 - [20] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics (ToG)*, 39(6):1–14, 2020.

- [21] Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. *ACM Transactions on Graphics (TOG)*, 37(6):1–11, 2018.
- [22] Lukas Lipp, David Hahn, Pierre Ecornier-Nocca, Florian Rist, and Michael Wimmer. View-independent adjoint light tracing for lighting design optimization. *ACM Trans. Graph.*, 43(3), May 2024. ISSN 0730-0301. doi: 10.1145/3662180.
- [23] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7708–7717, 2019.
- [24] I Lm. Approved method: Ies spatial daylight autonomy (sda) and annual sunlight exposure (ase). *Illuminating Engineering Society*. <https://www.ies.org/product/ies-spatial-daylight-autonomy-sda-and-annual-sunlight-exposure-ase>, 2013.
- [25] Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. Reparameterizing discontinuous integrands for differentiable rendering. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 38(6), December 2019. doi: 10.1145/3355089.3356510.
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [27] Max D Morris. Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33(2):161–174, 1991.
- [28] Azza Nabil and John Mardaljevic. Useful daylight illuminances: A replacement for daylight factors. *Energy and buildings*, 38(7):905–913, 2006.
- [29] Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. Mitsuba 2: A retargetable forward and inverse renderer. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 38(6), December 2019. doi: 10.1145/3355089.3356498.
- [30] Merlin Nimier-David, Sébastien Speierer, Benoît Ruiz, and Wenzel Jakob. Radiative backpropagation: An adjoint method for lightning-fast differentiable rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 39(4), July 2020. doi: 10.1145/3386569.3392406.
- [31] Christoph Reinhart. Daylight performance predictions. In *Building performance simulation for design and operation*, pages 221–269. Routledge, 2019.
- [32] Anson Savage. Differentiable objectives for 3d scene relighting via gradient descent on olat basis coefficients. 2026.
- [33] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [34] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.

- [35] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- [36] Ilya M Sobol. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and computers in simulation*, 55(1-3):271–280, 2001.
- [37] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International conference on machine learning*, pages 3319–3328. PMLR, 2017.
- [38] Jonathan Tremblay, Bowen Wen, Valts Blukis, Balakumar Sundaralingam, Stephen Tyree, and Stan Birchfield. Diff-dope: Differentiable deep object pose estimation. *arXiv preprint arXiv:2310.00463*, 2023.
- [39] Delio Vicini, Sébastien Speierer, and Wenzel Jakob. Path replay backpropagation: Differentiating light paths using constant memory and linear time. *Transactions on Graphics (Proceedings of SIGGRAPH)*, 40(4):108:1–108:14, August 2021. doi: 10.1145/3450626.3459804.
- [40] Zican Wang, Michael Fischer, and Tobias Ritschel. Stochastic gradient estimation for higher-order differentiable rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 28198–28206, 2025.
- [41] Jan Wienold and Jens Christoffersen. Evaluation methods and development of a new glare prediction model for daylight environments with the use of ccd cameras. *Energy and buildings*, 38(7):743–757, 2006.
- [42] Robert J Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):139–144, 1980.
- [43] Haocheng Yuan, Adrien Bousseau, Hao Pan, Quancheng Zhang, Niloy J Mitra, and Changjian Li. Diffcsg: Differentiable csg via rasterization. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–10, 2024.
- [44] Alan Yuille and Daniel Kersten. Vision as bayesian inference: Analysis by synthesis? *Trends in cognitive sciences*, 10(7):301–308, 2006.