Glare Mitigation using a Differentiable Unified Glare Rating

Eurographics Symposium on Rendering 2026 · Symposium Track

Linas Beresna and Eugene Fiume

SFU SIMON FRASER UNIVERSITY

Glare is a core constraint in lighting design

- Glare = perceived excessive brightness causing discomfort or loss of visual performance.
- Quantified in industry by the CIE Unified Glare Rating (UGR); the standard for offices (EN 12464-1).
- Today it is evaluated post-hoc: render, measure, manually tweak materials/lights, repeat.
- This trial-and-error loop is slow for complex multi-bounce scenes or multiple parameters.

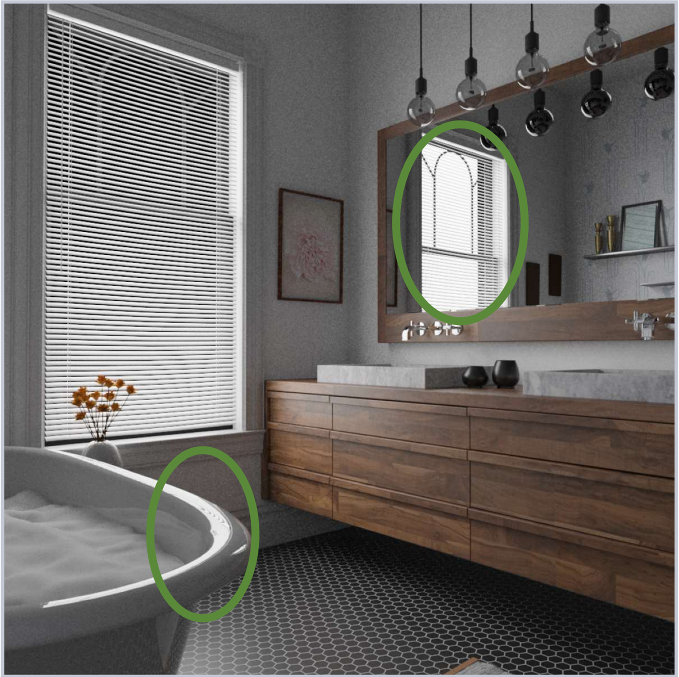

Glare is a core constraint in lighting design

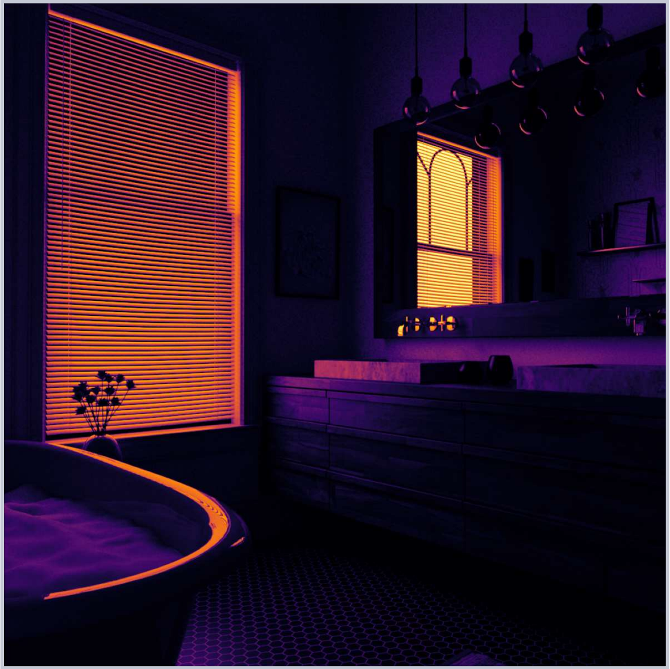

Luminance Map

Glare Mask Map

The Unified Glare Rating

$$UGR = 8 \log_{10} \left(0.25/L^{b} \cdot \Sigma\, L_i^{2}\, \omega_i \,/\, p_i^{2} \right)$$

- L_i = luminance of each glare source
- L^b = background adaptation luminance
- ω_i = solid angle of the source
- p_i = Guth position index (eccentricity)

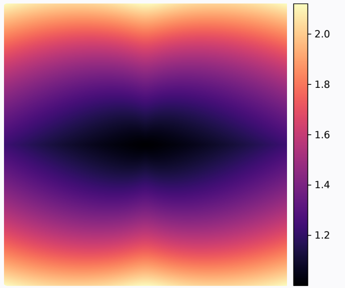

Guth Position

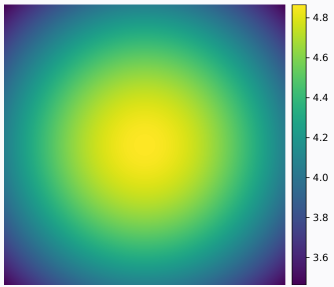

Solid Angle

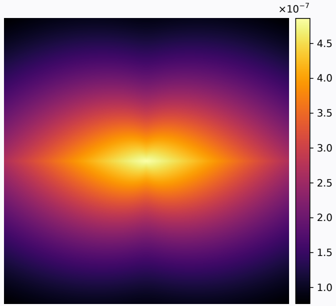

UGR Weight

Why UGR resists gradient-based optimization

- A pixel counts as glare only if $L > 5L^b$:
    - a hard binary threshold.
- Mathematically a Heaviside step:
    - gradients are zero almost everywhere,
- and a Dirac spike exactly at the boundary.
- Standard differentiable rendering may not get a usable signal from this.

Our approach: a differentiable UGR proxy

- Replace the discrete metric with a continuous, fully differentiable proxy.
- Gradients flow from the perceptual rating → through path-traced light transport → into physical scene parameters.
- Turns a passive evaluation into an active loss landscape for inverse design.
- To our knowledge, first integration of a visual-discomfort rating into a DR loss.

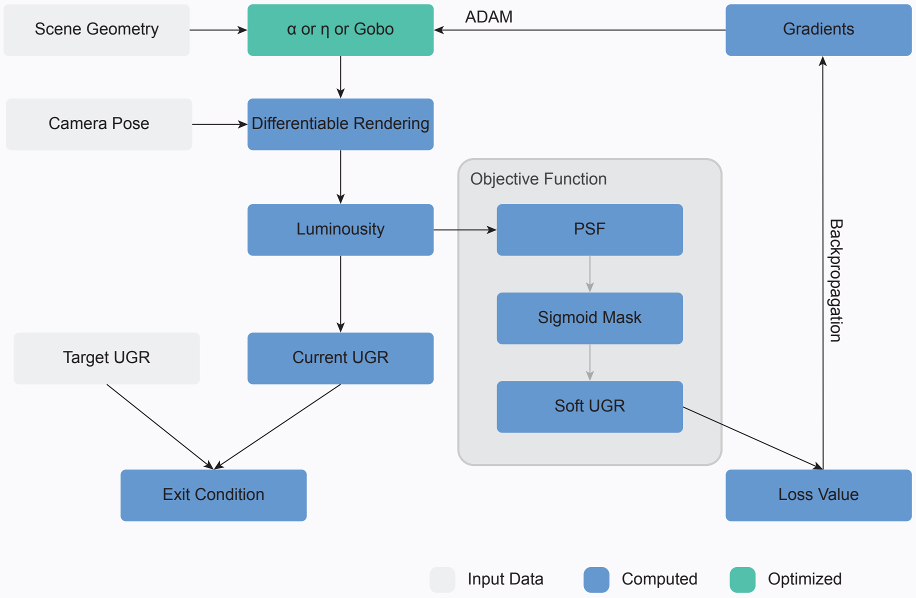

Ingredient 1: the eye's PSF heals the mask

- Inverse rendering runs at low spp (e.g. 16):
  - ’fireflies' get misread as glare.

Raw Glare Mask

Ingredient 1: the eye's PSF heals the mask

- Inverse rendering runs at low spp (e.g. 16):
    - ’fireflies' get misread as glare.
- We convolve luminance with the eye's Point Spread Function (intraocular scattering).

Raw Glare Mask

Ingredient 1: the eye's PSF heals the mask

- Inverse rendering runs at low spp (e.g. 16):
  - ’fireflies' get misread as glare.
- We convolve luminance with the eye's Point Spread Function (intraocular scattering).
- Implemented as p passes of a 3×3 Gaussian.
- Acts as a low-pass filter:
  - noise spikes dissolve
  - the glare mask becomes coherent.

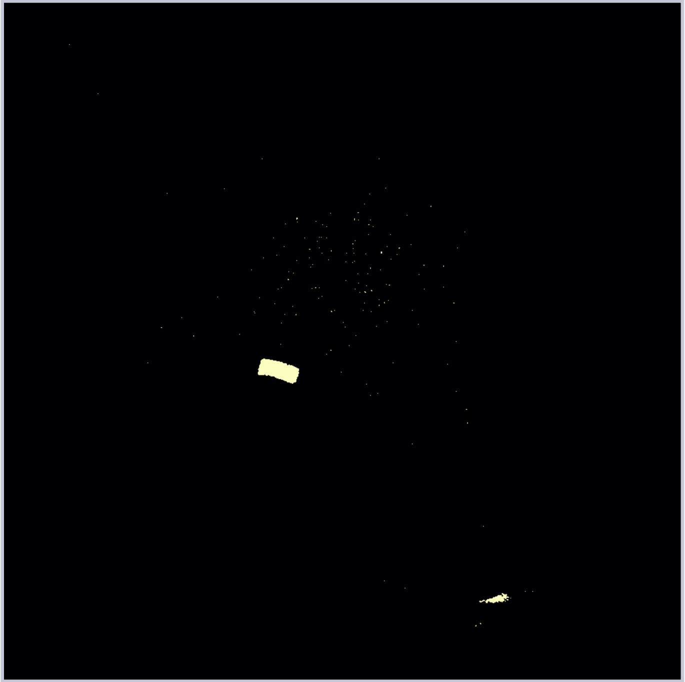

1 Pass

Ingredient 1: the eye's PSF heals the mask

- Inverse rendering runs at low spp (e.g. 16):
    - ’fireflies' get misread as glare.
- We convolve luminance with the eye's Point Spread Function (intraocular scattering).
- Implemented as p passes of a 3×3 Gaussian.
- Acts as a low-pass filter:
    - noise spikes dissolve
    - the glare mask becomes coherent.

2 Pass

Ingredient 2: a soft sigmoid boundary

- Replace the hard $L > 5L^b$ test with a sigmoid mask $m_i \in [0,1]$.
- Sharpness k controls the penumbra (we use $k = 10$).
- Pixels near the threshold now receive a meaningful, proportional gradient.

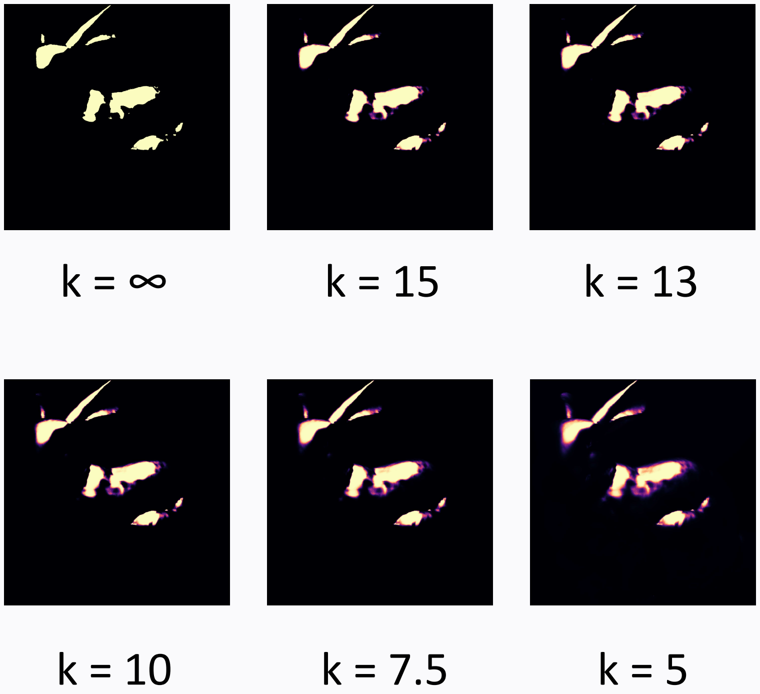

Three radiometric mitigation domains

Source side

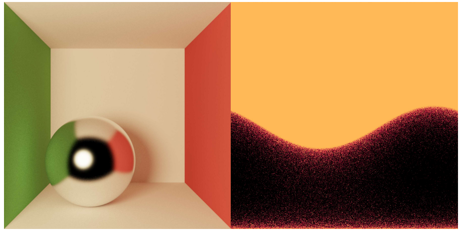

Emitter radiance masking: an automated 'computational gobo' dimming problem angles.

Boundary side

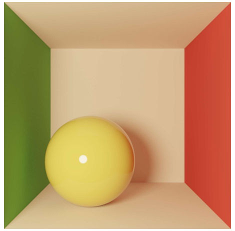

Anti-reflective: optimize index of refraction and spatial AR blend masks.

Surface side

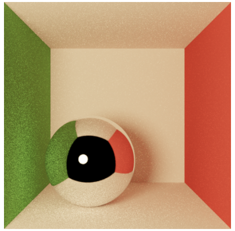

Anti-glare: roughen microfacet α to diffuse the specular highlight.

Validation: all domains converge

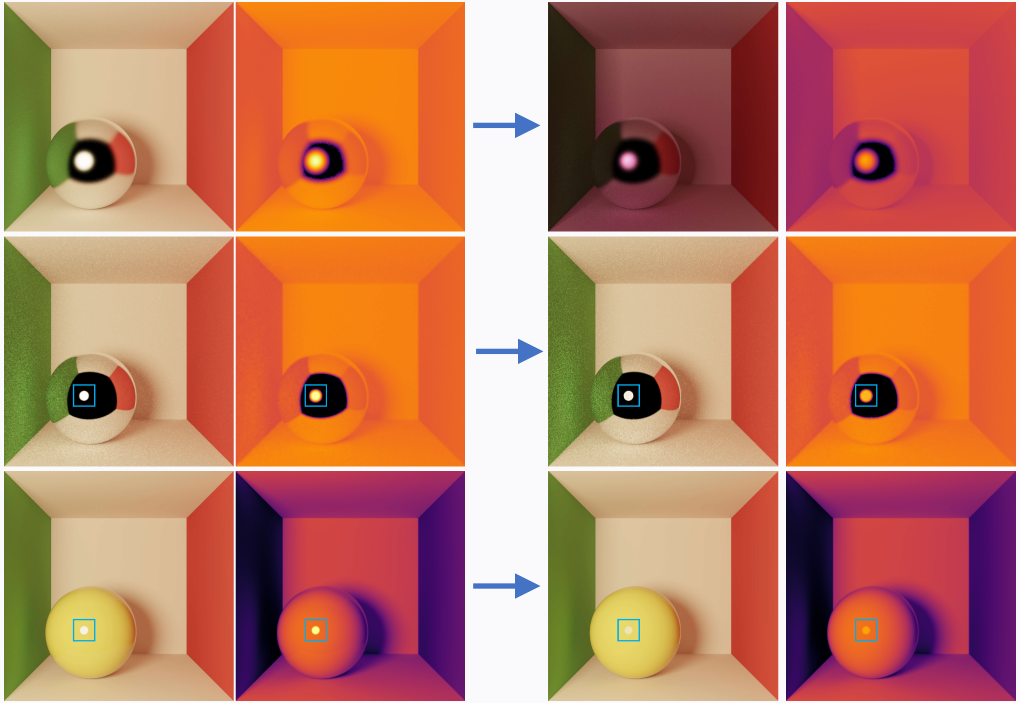

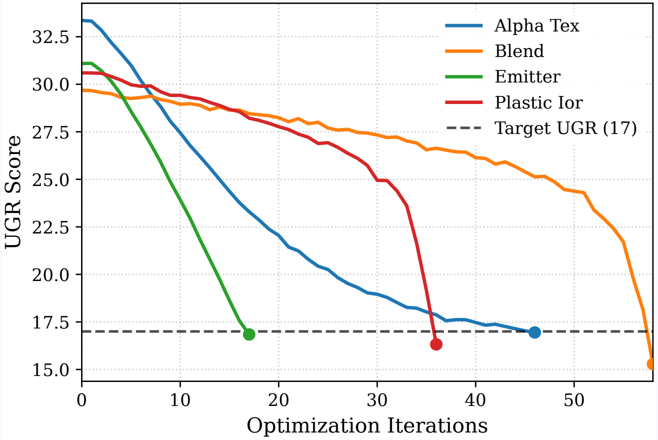

- Stable across scalar & spatial params.
- Also scales past derivative-free baselines (CMA-ES, Powell).

The ambient-starvation trap

- Roughening diffuses glare, but over-roughening starves the room of ambient light.
- Because UGR ∝ $1/L^b$, that darkening makes the score paradoxically **spike**.

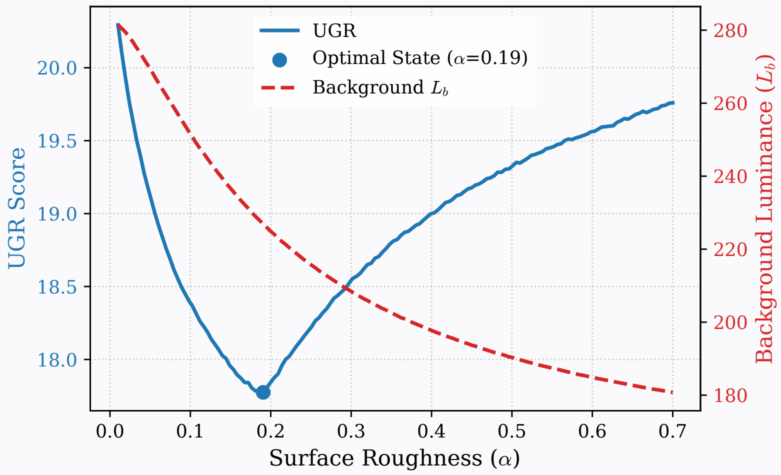

The ambient-starvation trap

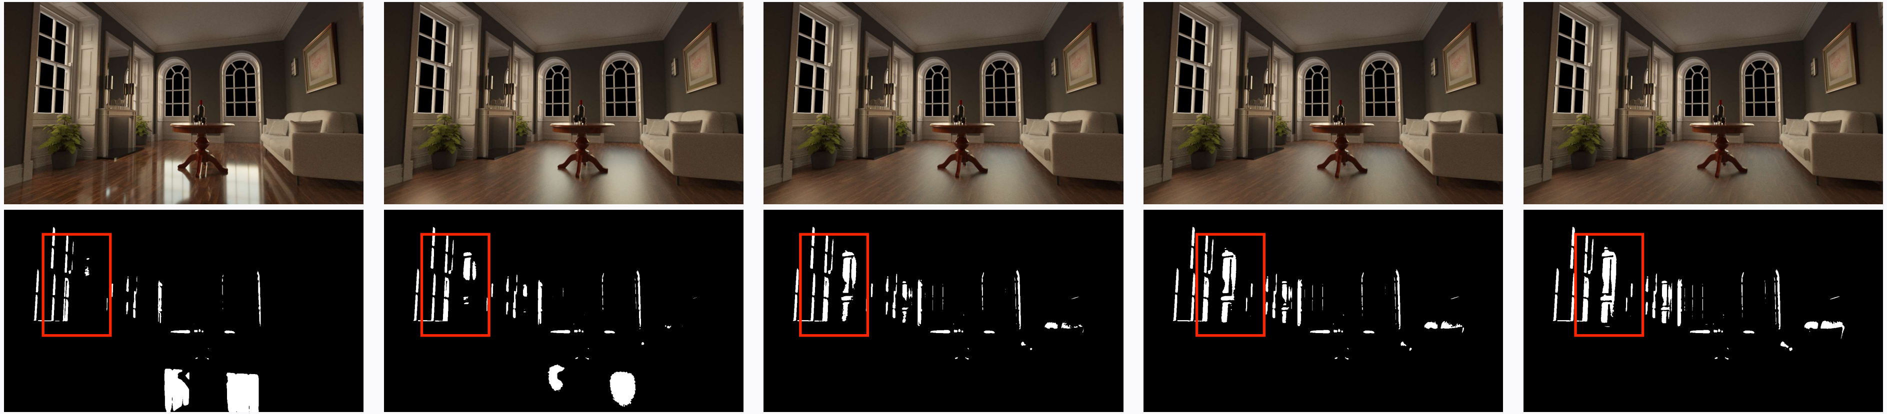

α = 0.01 | α = 0.15 | α = 0.25 | α = 0.3 | α = 0.4

The ambient-starvation trap

- Detaching $L^b \approx$ luminance-only optimization:
    - 'over-frosting', a dark room.
- Keeping L^b attached makes the loss contrast-aware:
    - a built-in self-regularizer.

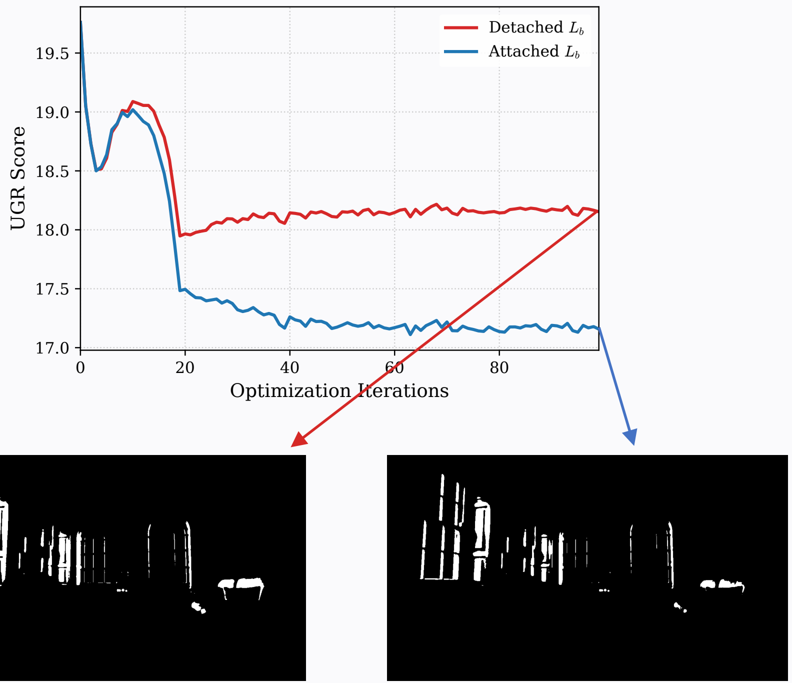

Limitations & future work

- Attach/detach L^b is currently a manual rule:
  - *automate via localized estimators*.
- User selects which parameters to expose:
  - *automated sensitivity analysis*.
- An inverse-design tool, not a manufacturing pipeline:
  - *integrate material catalogues*.
- Static, single viewpoint:
  - *extend to multi-view*
  - *temporal daylighting*.

Differentiable UGR: a perceptual rating that drives inverse design

- Continuous proxy (PSF + sigmoid) makes UGR optimizable.
- Stable across source, boundary, and surface domains.
- Attaching L^b self-regularizes against ambient starvation.

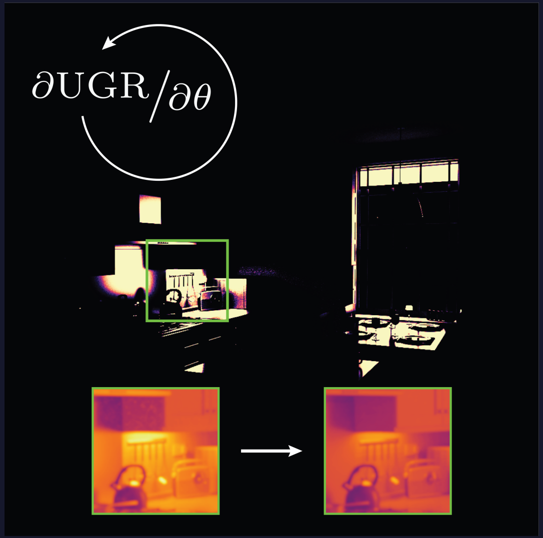

Thank you · Code available online · Questions welcome

SFU SIMON FRASER UNIVERSITY

Color exploits

- Perceptual luminance weights green heavily (CIE V(λ)).
- Unconstrained emitter optimization games this:
    - kills green, shifts the hue.
- Fix: constrain to a 1D scalar multiplier:
    - acts as a neutral-density filter.
- Mitigates glare while preserving the original light color.