Appendix A: Baseline Evaluation and Scalability

In this section, we provide an empirical evaluation of our adjoint gradient descent framework against standard derivative-free black-box optimizers (CMA-ES and SciPy's Powell method). The objective is to demonstrate the necessity of gradient-based optimization for high-dimensional spatial mitigation tasks.

Dimensionality Scaling

To evaluate scalability, we optimize the spatial roughness map (α) across exponentially increasing grid resolutions, from a uniform scalar ($N = 1$) to a 64×64 texture ($N = 4096$). To ensure a rigorous comparison, the optimization objective was formulated as a target-seeking task (a symmetric L_2 loss targeting an exact UGR of 17.0) rather than a simple minimization threshold. This prevents derivative-free solvers from trivially satisfying the objective by immediately clamping the parameter to its maximum value.

For our adjoint gradient descent, we utilized a dynamic learning rate that scales proportionally with the loss magnitude, naturally taking smaller steps as the scene approaches the target comfort level. As illustrated in Figure 1, our method maintains a near-constant optimization time regardless of dimensionality. Conversely, derivative-free methods exhibit exponential scaling, ultimately timing out or failing to converge entirely at higher resolutions.

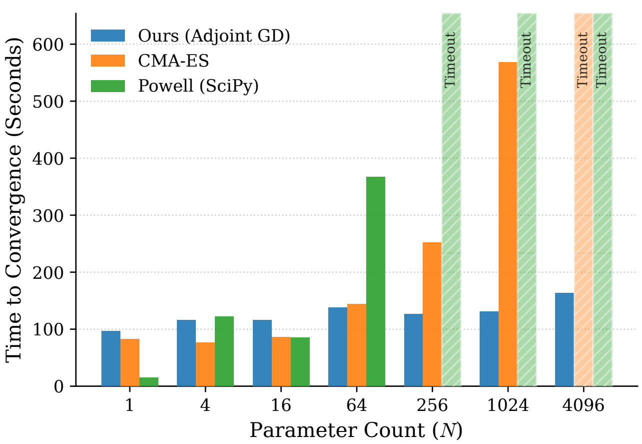

Figure 1: Optimization time vs. parameter count. CMA-ES and Powell methods fail to converge within realistic time constraints at higher resolutions, denoted by the hatched boxes.

It is worth noting that our current gradient-based implementation serves as a baseline proof-of-concept for the adjoint framework. We rely on a relatively simple proportional learning rate adjustment rather than highly optimized, adaptive scheduling algorithms or advanced line-search techniques. Consequently, the performance demonstrated by our method represents a lower bound; integrating more sophisticated solvers (such as L-BFGS or heavily tuned Adam schedulers) into our differentiable pipeline would likely yield even faster convergence, further widening the computational gap between our approach and black-box alternatives.

Domain Generalization

Table 1 evaluates the robustness of our framework across fundamentally different physical material domains. For high-dimensional spatial mitigation tasks ($N = 4096$), derivative-free baselines systematically fail to converge within the allotted time budget, regardless of whether the target is roughness, emitter intensity, or material blend weights.

It is important to note that the absolute convergence time for our adjoint method varies between these spatial domains (ranging from 70.9s to 387.7s). This variance is not a failure of the optimizer, but rather a reflection of the underlying path tracing complexity; for instance, computing gradients for a blend weight requires evaluating multiple BSDF lobes simultaneously, increasing the per-iteration rendering cost compared to a single-lobe roughness evaluation. However, the optimization process itself remains stable and successfully converges in all cases.

To provide a complete baseline, we also include a scalar optimization task ($N = 1$) targeting a global Plastic IOR. At this trivial dimensionality, derivative-free methods are able to succeed and converge slightly faster than our generalized pipeline. This highlights that while black-box methods are highly efficient for tuning isolated, global scalars, our adjoint framework is uniquely capable of scaling to the high-dimensional spatial maps required for complex architectural mitigation.

Parameter Domain	Ours	CMA-ES	Powell (SciPy)
Roughness (α)	70.9 s	Timeout	Timeout
Emitter Intensity	291.8 s	Timeout	Timeout
Blend Weight	387.7 s	Timeout	Timeout
Plastic IOR	91.5 s	88.7 s	69.6 s

Table 1: Time to convergence (in seconds) across different material properties. All spatial domains (Roughness, Emitter Intensity, and Blend Weight) were optimized at a 64×64 resolution ($N = 4096$), while Plastic IOR was a global scalar optimization ($N = 1$). Derivative-free methods only succeed on the trivial scalar domain, timing out on all high-dimensional spatial maps.

Qualitative Results

Figure 2 provides visual validation of our optimization pipeline compared to derivative-free methods across escalating spatial resolutions. To isolate the effect of dimensionality, we compare our adjoint gradient descent against CMA-ES and Powell's method at grid resolutions of 16×16 ($N = 256$), 32×32 ($N = 1024$), and 64×64 ($N = 4096$).

While our adjoint method smoothly converges to a comfortable visual state across all dimensionalities, the derivative-free baselines exhibit progressive degradation. As the parameter count increases, the black-box optimizers exhaust their evaluation budgets or time out before discovering a viable spatial mitigation map, leaving the scene in an unresolved, highly glaring state.

Proceedings published by Eurographics - The European Association for Computer Graphics.

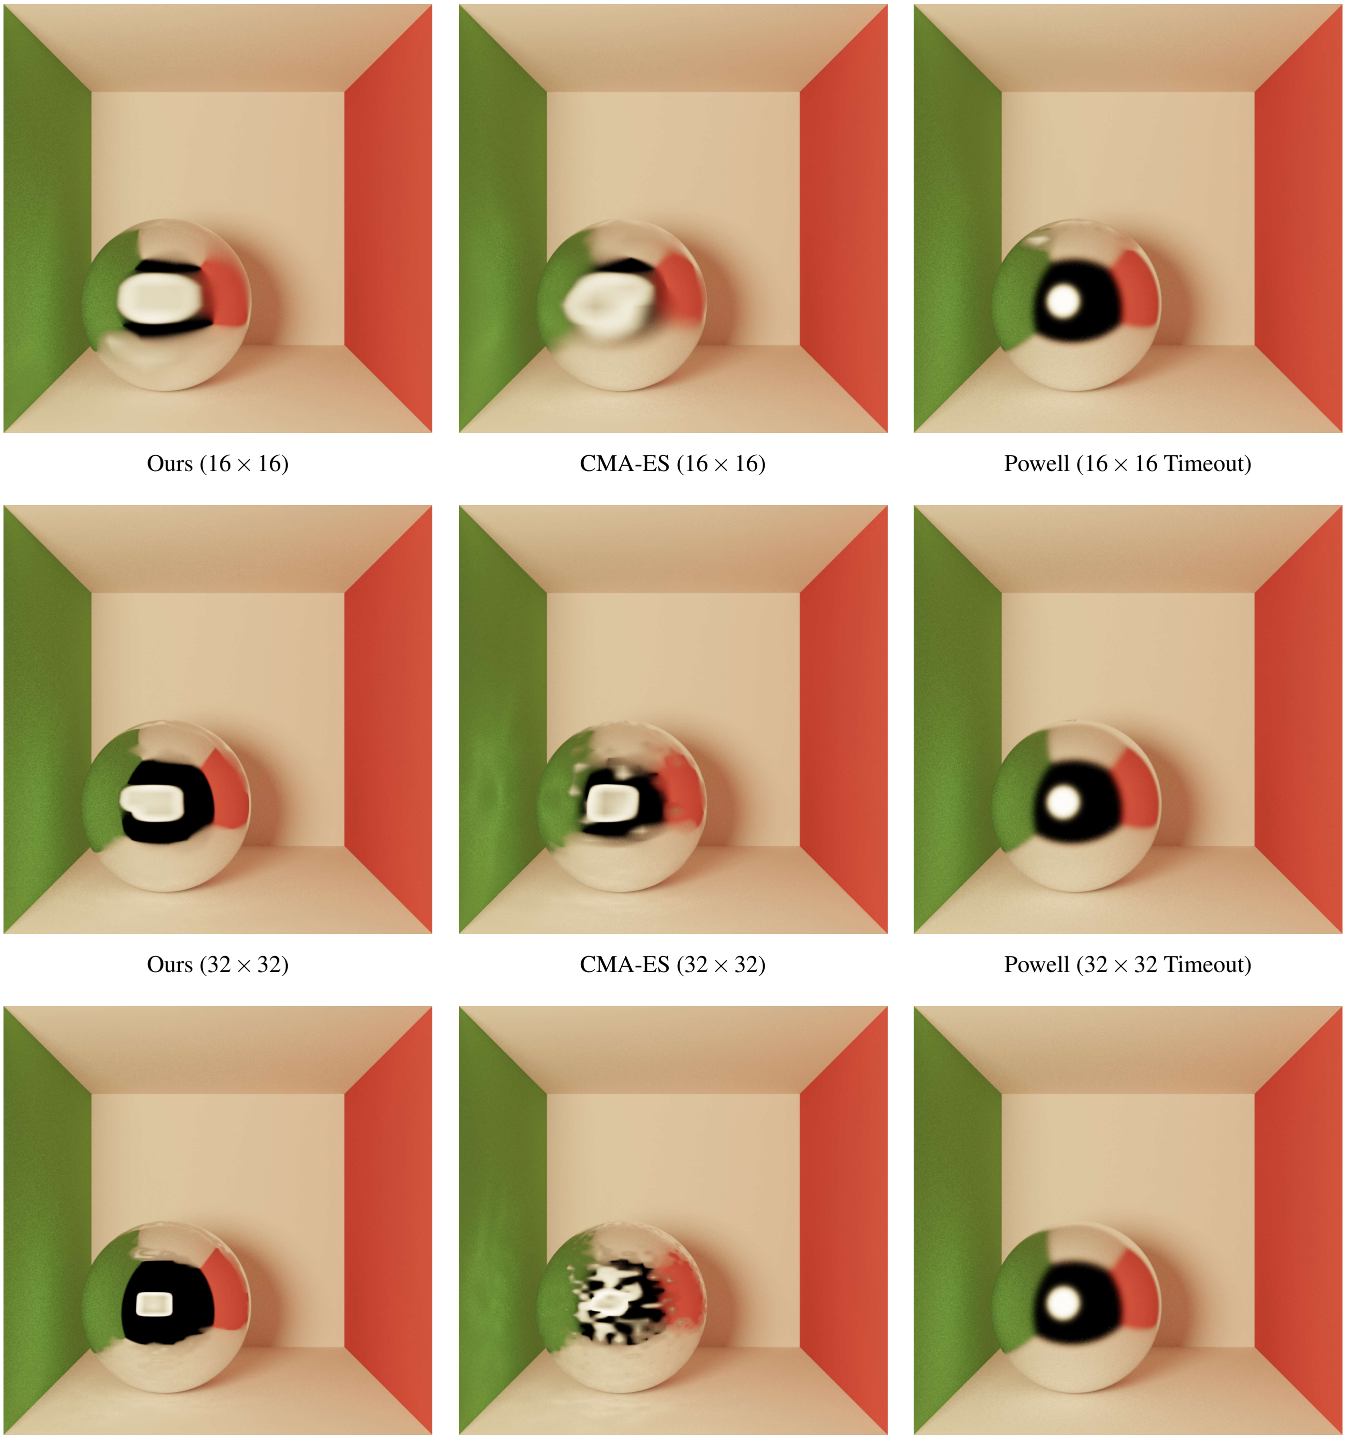

Ours (16 × 16) | CMA-ES (16 × 16) | Powell (16 × 16 Timeout)

Ours (32 × 32) | CMA-ES (32 × 32) | Powell (32 × 32 Timeout)

Ours (64 × 64) | CMA-ES (64 × 64 Timeout) | Powell (64 × 64 Timeout)

Figure 2: Visual states after the allocated optimization budget. The initial, unoptimized scene exhibits extreme glare (UGR $\approx$ 32.3). Across all spatial resolutions, our gradient-based approach successfully suppresses targeted glare without destroying global aesthetics. Notably, this baseline implementation operates without explicit spatial regularization (e.g., Total Variation or sparsity constraints); consequently, minor, non-essential roughness alterations (blurring) can be observed in some non-glaring regions of the sphere. In contrast, as dimensionality increases (top to bottom), the derivative-free optimizers exhaust their evaluation budgets (indicated by *Timeout*), failing to converge and leaving the severe glare largely unmitigated.