

Forget Superresolution, Sample Adaptively (when Path Tracing)

MARTIN BÁLINT, Max Planck Institute for Informatics, Germany
 CORENTIN SALAÜN, Max Planck Institute for Informatics, Germany
 HANS-PETER SEIDEL, Max Planck Institute for Informatics, Germany
 KAROL MYSZKOWSKI, Max Planck Institute for Informatics, Germany




Fig. 1. **Sparse adaptive sampling at 0.25 spp.** Our method learns to predict a sample density map (left, blue means fewer samples, red means more) that concentrates the limited rendering budget on perceptually important regions. Compared to state-of-the-art methods, DLSS 4 Ray Reconstruction [NVIDIA 2025] and NPPD [Bálint et al. 2023], our method recovers many more details. In the first (blue) inset, our method tracks the shadow boundary and achieves a sharp reconstruction, whereas others miss it entirely. In the second (green) inset, we recover the metallic appearance of the character’s suit, even in motion, whereas others produce a flat look. In the third (orange) inset, our method recovers the specular highlight on a disoccluded piece of railing, whereas others offer little lighting detail due to the complex geometry and camera motion.

Real-time path tracing increasingly operates under extremely low sampling budgets, often below one sample per pixel, as rendering complexity, resolution, and frame-rate requirements continue to rise. Superresolution is widely used in production because it reduces path-tracing cost by tracing rays on a coarser image grid and reconstructing missing details. This creates a uniform tradeoff between cost and spatial detail: every image region receives the same reduced ray budget, although path-tracing noise, reconstruction difficulty, and perceptual importance vary strongly across the image. Adaptive sampling offers a compelling alternative, but existing end-to-end approaches rely on approximations that break down in sparse regimes. We introduce an end-to-end adaptive sampling and denoising pipeline explicitly designed for the sub-1-spp regime. Our method uses a stochastic formulation of sample placement that enables gradient estimation despite discrete sampling decisions, allowing stable training of a neural sampler at low sampling budgets. To better align optimization with human perception, we propose a tone-mapping-aware training pipeline that integrates differentiable filmic operators and a state-of-the-art perceptual loss, preventing oversampling of regions with low visual impact. In addition, we introduce a gather-based

pyramidal denoising filter and a learnable generalization of albedo demodulation tailored to sparse sampling. Our results show consistent improvements over uniform sparse sampling, with notably better reconstruction of perceptually critical details such as specular highlights and shadow boundaries, and demonstrate that adaptive sampling remains effective in the sub-1-spp regime.

CCS Concepts: • **Computing methodologies** → **Ray tracing; Neural networks.**

ACM Reference Format:

Martin Bálint, Corentin Salaün, Hans-Peter Seidel, and Karol Myszkowski. 2026. Forget Superresolution, Sample Adaptively (when Path Tracing). *ACM Trans. Graph.* 45, 4, Article 90 (July 2026), 19 pages. <https://doi.org/10.1145/3811377>

1 Introduction

Real-time physically based rendering is fundamentally limited by the number of light paths that can be computed within tight time budgets. At high display resolutions and frame rates, especially on complex scenes, this budget may fall short of even a single sample per pixel (*spp*). Recent hardware advances have made real-time path tracing increasingly practical [Ouyang et al. 2021], yet per-sample complexity will continue to rise as rendering pipelines adopt neural materials [Zeltner et al. 2024], sophisticated light transport, and volumetric effects. As a result, sparse sampling budgets are likely here to stay.

Superresolution has become commonplace in production game engines. Rendering at reduced internal resolution and upscaling the

Authors’ Contact Information: Martin Bálint, Max Planck Institute for Informatics, Saarbrücken, Germany, mbalint@mpi-inf.mpg.de; Corentin Salaün, Max Planck Institute for Informatics, Saarbrücken, Germany, csalaun@mpi-inf.mpg.de; Hans-Peter Seidel, Max Planck Institute for Informatics, Saarbrücken, Germany, hpsidel@mpi-inf.mpg.de; Karol Myszkowski, Max Planck Institute for Informatics, Saarbrücken, Germany, karol@mpi-inf.mpg.de.



This work is licensed under a Creative Commons Attribution 4.0 International License.
 © 2026 Copyright held by the owner/author(s).
 ACM 1557-7368/2026/7-ART90
<https://doi.org/10.1145/3811377>

result is well-suited to rasterization, which relies on a fixed pixel grid. Optionally, learned denoisers are paired with superresolution and used to help resolve simpler Monte Carlo rendered effects. However, as rendering moves toward full path tracing, spatial adaptation is becoming increasingly important. Monte Carlo noise, denoiser reconstruction quality, and the perceptual visibility of artifacts all vary across the image depending on content, yet superresolution sacrifices spatial detail uniformly.

Luckily, with full path tracing, we are free to place samples anywhere in the image. A jointly trained sampler and denoiser can, in theory, learn to navigate the interdependencies between noise, reconstruction, and perception, concentrating samples where they matter most for the final image. Our goal is to build the first end-to-end adaptive sampling and denoising pipeline for the sub-1-spp regime.

The first challenge is building differentiable sparse adaptive sampling: we need to backpropagate gradients from the rendered image to the sampler network. Unfortunately, choosing which pixels to sample is a discrete operation, making gradient computation non-trivial. Prior work overcomes this issue by relying on approximations or surrogate gradients that hold reasonably well at moderate sample counts but break down as the sampling rate decreases. Such approaches limit effective adaptive sampling to budgets above 2 spp [Fan et al. 2025]. We introduce a formulation that chooses sampled pixels stochastically, letting us derive expected gradients from per-pixel statistics. We efficiently approximate these gradients, enabling end-to-end training of the sampler network. Our approach remains robust with as few as one sample per hundred pixels in the sparsest regions.

Second, effectively redistributing a limited sampling budget requires a carefully designed optimization objective. Minimizing per-pixel error alone is not well aligned with human perception of image distortions and often leads to suboptimal visual results. Ideally, the method should learn to prioritize reconstruction in the most perceptually critical regions. Unfortunately, loss formulations used by prior works are not well-aligned with perceptual sensitivity, resulting in apparent over- and undersampling artifacts across the image. The issue is amplified in the sub-1-spp regime, where deviations from reference images are necessarily larger, making it harder for losses to reliably judge which distortions are most noticeable. To this end, we propose a tone-mapping-aware pipeline that includes the state-of-the-art perceptual loss, MILO [Çoğalan et al. 2025].

Third, sparse inputs require architectural changes to the denoiser itself. We introduce a gather-based pyramidal filter that remains stable when most pixels lack samples, and a learnable demodulation mechanism that preserves texture detail without relying on renderer-provided albedo decomposition.

Finally, we demonstrate that stable and consistent adaptive sampling strategies can be learned even at very low budgets, including the sub-1-spp regime highlighted by the teaser example (Fig. 1). Our adaptive sampling provides notably better quality metrics, in addition to resolving small-but-difficult details such as shadow boundaries and specular highlights. We further show how adaptive sampling benefits from the more precise sample placement afforded by higher resolutions, pointing toward a promising direction for real-time path tracing research.

We summarize our contributions as follows:

- A stochastic formulation for differentiable sparse adaptive sampling, enabling stable end-to-end training at sub-1-spp budgets.
- A tone-mapping-aware training pipeline with a family of differentiable filmic curves, allowing the use of state-of-the-art perceptual losses.
- A denoising filter adapted to sparse renderings.
- The dataset and code for our method are available on our project website.¹

2 Related Work

Adaptive sampling. Adaptive sampling plays a crucial role in Monte Carlo rendering, where noise in the produced images is not uniformly distributed across the entire image. Samples from some pixels may exhibit larger variance than others, meaning the number of samples required to achieve satisfactory quality across the image may vary [Mitchell 1987]. Adaptive sampling addresses this challenge by adjusting the sample count for each pixel to equalize per-pixel error, ultimately leading to higher overall image quality within the same rendering time.

We group adaptive sampling techniques into two main categories: iterative and prediction-based sampling. In iterative methods, the image is rendered in multiple passes, typically starting with a uniform pass and then updating the sample distribution in subsequent passes based on gathered statistics [Bolin and Meyer 1998; Firmino et al. 2023]. On the other hand, prediction-based sampling uses guiding information available prior to rendering to estimate the sample budget for each pixel in a single pass. These methods often rely on neural networks for their robust predictive capabilities. In this work, we focus on this second class of adaptive sampling.

When using both neural sampling and denoising, optimizing each component separately using simple heuristics yields limited results. Therefore, adapting the sampler network to a given denoiser is crucial for achieving high-quality results with fewer samples. For instance, Vogels et al. [2018] proposed training an error predictor network to estimate regions with higher error and then using it as a sampling map for adaptive sampling. Kuznetsov et al. [2018] suggested directly applying backpropagation of the denoising error through the rendering process, using a finite-difference gradient approximation of the rounded sample count for low-sample-count adaptive sampling. This method was later extended with temporal processing by Hasselgren et al. [2020], guiding sampling using the previous denoised frame, further reducing rendering costs. Salehi et al. [2022] proposed removing the renderer from the training process by generating path tracing samples from approximate analytic distributions. This method relies on the convergence of the Monte Carlo estimator to a Gaussian distribution as the sample count increases, thanks to the Central Limit Theorem. Unfortunately, Monte Carlo estimators do not follow such an analytic distribution at low sample counts. The method closest to our work, Scardigli et al. [2023], uses reinforcement learning to optimize the sampler network via a critic network trained to predict the denoiser’s expected loss for a given sampling map.

¹<https://balint.io/fssa>




Fig. 2. **Overview of our pipeline.** Given GBuffer features and temporal data from the previous frame, the sampler network predicts a continuous sample density map. After stochastic discretization (Section 4), the renderer traces the requested samples, producing a sparse, noisy image. The denoiser network, which shares latent features with the sampler, predicts weights for our pyramidal gather filter (Section 6). The denoised HDR output is tone-mapped for display and fed back to the sampler in the next frame, closing our perceptual feedback loop (Section 5).

Denoising. Denoising is a crucial step in rendering pipelines utilizing Monte Carlo-based light transport algorithms. Its main goal is to eliminate the noise inherent in Monte Carlo processes, allowing for the generation of high-quality images even with limited sampling budgets. Early denoising methods involved local blurring of noisy input images with non-linear filters [Lee and Redner 1990; Mitchell 1987; Overbeck et al. 2009; Rushmeier and Ward 1994; Xu and Pattanaik 2005]. Subsequent advancements incorporated auxiliary features such as surface normals, albedo, depth, and estimated variance to adapt filters locally, producing more accurate images [Bitterli et al. 2016; Li et al. 2012; McCool 1999; Rousselle et al. 2013]. Recently, the advent of deep learning and corresponding hardware development has spurred the creation of deep learning-based denoisers [Bako et al. 2017; Chaitanya et al. 2017; Kalantari et al. 2015], leading to significant improvements in denoising quality. These approaches range from trainable models for directly predicting the output image from noisy input [Xiao et al. 2020; Zhang and Yuan 2024] to kernel prediction techniques that learn per-pixel kernel parameters [Gharbi et al. 2019; Hofmann et al. 2023; Lee et al. 2024; Zhang et al. 2024]. Most recently, multiscale kernel prediction and filtering have shown significant improvements in interactive denoising [Bálint et al. 2023; Işık et al. 2021].

Most denoising methods assume a uniform distribution of samples across the screen space, often assuming samples in every image pixel. However, in the context of adaptive sampling and real-time rendering, it may not always be feasible or desirable to adhere to such constraints. Using a denoiser optimized for a uniform sample distribution in adaptive sampling methods can result in poor outcomes, especially with deep learning methods. Consequently, recent work has explored denoisers that are aware of adaptive sampling [Firmino et al. 2023; Salehi et al. 2022].

Superresolution. Superresolution is an alternative method to improve performance by rendering at a reduced resolution and upscaling the result to full resolution. This guarantees that all pixels can

be sampled, but on a coarser grid. The approach is widely used in real-time pipelines, notably in production methods such as DLSS, FSR, and XeSS. Early learning-based methods include, for example, the Laplacian pyramid network of Lai et al. [2017]. Later work incorporates temporal information and auxiliary buffers for real-time use [Xiao et al. 2020; Zhong et al. 2023], while generative methods attempt to hallucinate high-frequency details beyond the input signal [Ge et al. 2018].

Rendering at reduced resolution, followed by superresolution, offers a simple alternative to adaptive sampling but lacks spatial selectivity. All regions receive the same effective sampling rate, regardless of their rendering complexity. In contrast to adaptive sampling, this prevents allocating additional samples to difficult regions such as reflections, shadows, or high-frequency lighting effects while oversampling simpler areas. In this work, we show that superresolution may be better suited to rasterization pipelines with uniform per-pixel cost than to Monte Carlo path tracing, particularly in scenarios where noise and information content vary significantly across the image, and efficient reconstruction at low sampling budgets relies on highly targeted sampling strategies.

Superresolution has also been successfully combined with denoising for path-traced rendering [Kazmierczyk et al. 2025; Thomas et al. 2022]. The additional denoising stage enables more faithful reconstruction of fine details than superresolution alone, particularly in regions affected by Monte Carlo noise.

Tone-mapping for neural training. Most learning-based denoising and adaptive sampling methods optimize losses defined in linear HDR space, typically using relative-error metrics such as SMAPE or logarithmic $L1$ losses [Bálint et al. 2023; Kuznetsov et al. 2018; Salehi et al. 2022; Vogels et al. 2018]. These losses do not account for perceptual adaptation or downstream tone-mapping and therefore misrepresent which errors are actually visible. As a result, computation is often wasted in very dark or very bright regions where large linear errors are perceptually negligible. Some approaches instead

train directly in tone-mapped (SDR) space [Áfra 2024; Fan et al. 2025; Hasselgren et al. 2020], improving perceptual alignment but introducing bias in kernel-based denoisers and coupling the model to a fixed tone-mapping operator. These limitations highlight the need for training strategies that explicitly account for tone-mapping and perception when jointly optimizing sampling and denoising. To address these limitations, we propose an expressive family of parametric, differentiable tone-mapping operators suitable for end-to-end learning of adaptive sampling and denoising. In addition, we employ a perceptual error metric as the loss function to guide sample allocation toward image regions containing highly visible and objectionable distortions, as perceived by a human observer.

3 Overview

Before detailing our contributions, we outline the pipeline that integrates them. Adaptive sampling and denoising pipelines share a common structure (Fig. 2): a sampler network predicts per-pixel sample densities, the renderer traces paths accordingly, and a denoiser reconstructs the final image from the noisy result. Our pipeline follows this established design, with modifications that enable stable training at sub-1-spp budgets.

The sampler takes GBuffer features—depth, normals, and albedo—alongside the reprojected output and latent features from the previous frame (Green and Blue loops). From these inputs, it predicts a continuous sample density map that distributes the frame’s sampling budget across pixels. This density must then be discretized into integer sample counts before rendering. As we detail in Section 4, this discretization step is the key obstacle to end-to-end training: deterministic rounding [Kuznetsov et al. 2018] produces zero gradients almost everywhere, while existing surrogate approximations become unstable at low sample counts. We propose a stochastic formulation that interprets fractional densities as probabilities, enabling robust gradient estimation even at extreme sparsity.

We obtain the sample density map s by first applying softmax normalization to one of the sampler network’s output channels. We then allocate $1/8$ of the total budget uniformly to ensure numerical stability and distribute the remaining $7/8$ adaptively based on the network’s normalized output. We find that distributing a larger fraction of the budget adaptively slows training convergence without yielding any measurable improvement in quality.

After rendering, the denoiser receives the noisy image together with the same auxiliary features and reprojected data from the previous frame (Blue feedback loop) and the previous frame in HDR (Orange feedback loop). Rather than predicting pixel values directly, it outputs weights for a pyramidal gather filter that reconstructs the final HDR image. The sampler and denoiser share a latent feature space (Blue feedback loop), enabling them to collaborate: the sampler can learn to place samples where the denoiser struggles, while the denoiser can adapt its reconstruction to the sampler’s allocation patterns.

The reconstructed image is then tone-mapped for display. We feed the tone-mapped output back to the sampler in the next frame (Green feedback loop), allowing for alignment with perceptual importance. During training, our differentiable parametric tone-mapper (Section 5) enables gradient flow through this loop. At

inference, any tone-mapper can be substituted: the graphics engine simply feeds its output back to the sampler, making integration straightforward. We describe our gathering pyramidal filter and other architectural choices in Section 6.

Note that, like previous end-to-end approaches, our method does not require a differentiable path tracer. Following prior work by Kuznetsov et al. [2018], we pre-render images at power-of-two sample counts, which can be combined to produce arbitrary budgets during training.

4 Differentiable sparse adaptive sampling

Path tracing is inherently discrete: for each pixel, a sample is either computed or not. Continuous sample densities predicted by the sampler network must therefore be converted to integer counts before rendering. Prior work relies on deterministic rounding with surrogate gradient approximations, but this approach fundamentally limits what adaptive sampling can achieve.

We reframe sample allocation as an exploration-exploitation problem. At each pixel, the sampler must balance *exploiting* regions where it confidently knows samples are needed against *exploring* regions where it is uncertain. A well-tuned strategy samples uncertain regions often enough to gather information while concentrating resources where the payoff is known to be high. Deterministic rounding cannot capture this balance: densities of 0.3 spp and 0.0 spp both yield zero samples, even though the former suggests meaningful uncertainty. We propose stochastic rounding, which interprets the sample density as a probability. A density of 0.3 spp corresponds to a 30% chance of rendering a sample, allowing the network to express confidence continuously and learn to balance exploration and exploitation through gradient-based optimization.

We derive the expected gradient analytically and introduce a relaxed approximation that enables efficient backpropagation while preserving the estimator’s statistical properties. Our approach remains robust at densities as low as 0.01 samples per pixel.

In the following subsections, we compare two ways of discretizing the predicted density map: deterministic rounding and our stochastic formulation. We first introduce common notation. For each pixel i , let $s_i \in \mathbb{R}_{\geq 0}$ denote the continuous sample density predicted by the sampler, $S_{i,j}$ the color of the j -th path-traced sample, L_i^{noisy} the resulting noisy Monte Carlo estimate, and $L_i^{\text{ref}} = \mathbb{E}[S_{i,j}]$ the reference radiance.

4.1 Deterministic rounding

Deterministic rounding for learnable adaptive sampling was first proposed by Kuznetsov et al. [2018]:

$$L_i^{\text{noisy}} = \begin{cases} \frac{\sum_{j=1}^{\text{round}(s_i)} S_{i,j}}{\text{round}(s_i)} & \text{if } s_i \geq 0.5 \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where, for each pixel i , the sample density s_i is rounded to the nearest integer and that many samples $S_{i,j}$ are averaged to produce the noisy estimate L_i^{noisy} .

The gradient of Eq. (1) is zero almost everywhere, reducing to Dirac delta impulses at the transition points. This provides no useful signal for optimization. Instead, Kuznetsov et al. [2018] proposed to




Fig. 3. **Visualization of our relaxed estimator.** On the left, we illustrate the effect of the sampling probability p on the slope of the function at a given temperature parameter λ . On the right, we compare the estimator at a given value of p for different values of λ , along with the Gumbel-Softmax estimator for reference.

use a surrogate gradient:

$$\frac{\partial L_i^{\text{noisy}}}{\partial s_i} = \begin{cases} \frac{L_i^{\text{ref}} - L_i^{\text{noisy}}}{\text{round}(s_i)} & \text{if } s_i \geq 0.5 \\ L_i^{\text{ref}} & \text{otherwise.} \end{cases} \quad (2)$$

Although many later works [Fan et al. 2025; Hasselgren et al. 2020; Salehi et al. 2022] adopt this formulation, the fixed threshold at $s_i = 0.5$ leads to severe failures in sparse regimes. When a spatially extended region is deemed easy to render and $s_i < 0.5$ for all its pixels, deterministic rounding assigns zero samples to the entire region. This leaves large areas with no illumination information, making accurate reconstruction impossible. In principle, even a few samples distributed across the region would suffice to recover a plausible signal, but the bias of the surrogate gradient in Eq. (2) seemingly precludes such partial coverage.

Some approaches mitigate this by imposing mandatory samples in regular patterns [Fan et al. 2025]. While this prevents the most important artifacts, it does not allow the sampler to express fine-grained confidence or receive meaningful gradients. Thus, the learned allocation strategies remain limited.

4.2 Stochastic rounding

We apply the principle of stochastic rounding to the adaptive sampling setting, deriving a formulation that maintains unbiased radiance estimation while enabling gradient-based optimization.

For each pixel i , we always allocate the integer part $\lfloor s_i \rfloor$ samples. Let $p_i = s_i - \lfloor s_i \rfloor$ denote the fractional part. We draw a uniform variate $u_i \sim \mathcal{U}(0, 1)$ and allocate one extra sample whenever $u_i < p_i$. We denote this potential extra sample r_i and define its contribution as:

$$L_i^h = r_i [u_i < p_i]. \quad (3)$$

L_i^h is the contribution of the optional fractional sample: it is either the extra sample color r_i or 0, depending on whether that sample is taken. The final noisy estimate is the normalized sum of all allocated samples:

$$L_i^{\text{noisy}} = \frac{\sum_{j=1}^{\lfloor s_i \rfloor} S_{i,j} + L_i^h}{s_i}. \quad (4)$$




Fig. 4. **Gradient comparison.** We compare the straight-through estimator and our proposed relaxed estimator against the ground-truth gradient estimated via finite differences on a training crop. Our estimator is closest to the finite-difference reference, leading to more stable optimization due to lower gradient bias.

Equation 4 combines the integer samples with the optional hold-out sample and normalizes the sum by the requested density s_i . Rather than dividing by the realized integer sample count, this estimator normalizes by the requested density s_i . This choice has several advantages. First, differentiating this step is trivial for $s_i > 0$. Second, assuming the per-pixel path samples $S_{i,j}$ are identically distributed with mean L_i^{ref} , the estimator remains unbiased across the full density range. Appendix A gives the full derivation, including the correction factor used by the relaxed estimator. Finally, normalization by s_i naturally upweights pixels with the extra sample taken, providing inverse-variance weighting. Intuitively, this formulation remains well-defined even when $s_i \ll 1$: pixels with low densities simply have a correspondingly low probability of being sampled, but when they are, their contribution is appropriately scaled.

Deriving gradients through stochastic processes is a well-established technique for differentiating discrete operations [Li et al. 2018]. The expected gradient of the loss with respect to s_i exists and is well-behaved. However, computing this gradient exactly requires edge sampling of each discontinuity, effectively running the denoiser for each pixel, which is prohibitively expensive for training. We therefore introduce a relaxed formulation that enables efficient gradient computation via standard backpropagation.

4.3 Relaxed stochastic rounding

Following the Gumbel-Softmax relaxation [Jang et al. 2017; Maddison et al. 2017], we replace the discontinuous step function in Eq. (3) with a smooth approximation. While effective for gradient propagation, this relaxation is ill-suited for adaptive Monte Carlo sampling. The resulting sigmoid-like transition has non-zero tails, allowing bright samples to “shine through” in regions that should receive none. Although these tails narrow with increasing temperature parameters, rare outlier draws still occur, resulting in train-test mismatch.

We design our relaxation to avoid this issue. Rather than a symmetric sigmoid, we use a clipped linear ramp whose base sits exactly at the sampling threshold:

$$L_i^h = r_i \cdot h \cdot \min\left(\frac{\lambda}{p_i}(u_i + p_i - 1), 0, 1\right), \quad (5)$$

where the temperature $\lambda \geq 1$ controls the slope of the ramp. With probability $1 - p_i$, the uniform variate u_i falls below the ramp

entirely, yielding $L_i^h = 0$ with zero gradient. In this case, no sample needs to be computed. Only when u_i falls within the linear region (probability p_i/λ) or above (probability $p_i(1 - 1/\lambda)$) do we allocate a sample. This behavior can be replicated exactly at test time: we simply skip sample computation when $u_i < 1 - p_i$. Figure 3 depicts our relaxed stochastic rounding function.

As $\lambda \rightarrow \infty$, the ramp converges to the hard step in Eq. (3). For finite λ , the ramp produces nonzero gradients over the linear region, proportional to p_i/λ . Because the ramp is asymmetric, the expected contribution of the holdout sample would be attenuated without correction. We compensate with a factor:

$$h = \frac{2\lambda}{2\lambda - 1}, \quad (6)$$

which rescales the ramp such that $\mathbb{E}[L_i^h]$ matches the variant in Eq. (3), preserving $\mathbb{E}[L_i^{\text{noisy}}] = \mathbb{E}[S_i]$. Geometrically, we preserve the area under the curve as seen in Fig. 3.

The choice of λ presents a tradeoff: larger values produce gradients closer to the true expected gradient but with higher variance. Unlike Gumbel-Softmax, which requires a high temperature to avoid significant train-test mismatch, our formulation is robust across a wide range of temperatures. Lower values of λ widen the linear region, adding noise to L_i^h when u_i falls within it, but this effect quickly diminishes with increasing λ . Consequently, we observe that ablations trained with $\lambda = 8, 10$, and 12 all fall within run-to-run variance. Gradient norms remain comparable to the uniform sampling baseline, indicating that these λ values are not large enough to affect training stability either. In practice, we use $\lambda = 10$. We further compare our relaxed gradient estimator to the straight-through estimator [Bengio et al. 2013] in Fig. 4.

5 Tone-mapping-aware training

Adaptive sampling trained end-to-end relies on the loss to guide sample allocation. Under low sample budgets, rendering distortions can be substantial. In this regime, a perceptual loss is needed to weigh distortions against each other and prioritize sampling of those that are more visible and objectionable. Optimizing strict per-pixel fidelity is therefore less effective than minimizing perceptual artifact visibility, which can be further altered by downstream display-referred tone-mapping [Reinhard et al. 2010]. Ignoring these factors can lead to oversampling regions where errors are difficult to perceive, while undersampling regions where additional samples would yield visible improvements.

We address these issues with three modifications to existing approaches. First, we apply a family of differentiable tone-mapping operators (TMOs) to both network outputs and reference images. Differentiability ensures that gradients can flow backward during training, while at inference, any tone-mapper can be used. Second, we provide the sampler network with the previous frame's tone-mapped output. This enables the sampler to automatically adapt to the tone-mapped appearance and integrates trivially into existing rendering pipelines for inference. Finally, instead of computing losses in linear or logarithmic color spaces, we evaluate errors in the tone-mapped domain using the state-of-the-art perceptual loss MILO [Çoğalan et al. 2025]. This leverages the fact that such perceptual metrics are trained and applied in the sRGB domain, where




Fig. 5. **Visualization of our differentiable tone-mapping operator (TMO) for three parameter settings.** The parameters s and h control the compression of the lower (toe) and upper (shoulder) luminance tails, respectively. The curves are plotted as a function of the logarithm of radiance.

large-scale human-labeled image quality datasets are predominantly available.

Differentiable Parametric Filmic Tone-Mapping. Generalization at test time demands training with a diverse set of differentiable TMOs. First, we transform the output of our pyramidal denoising filter to the logarithmic color space and apply basic augmentations:

- **Exposure k :** additive offset in log-space shifts mid-gray.
- **Contrast α :** multiplicative scale about 0 adjusts curve steepness.
- **Saturation β :** per-channel scaling around the log-RGB mean modulates color vividness.

More formally,

$$I^{\text{denoised}} = \text{sRGB}(\tau(\alpha(\log L^{\text{denoised}} \odot \beta + k))), \quad (7)$$

where $\odot$ is a shorthand for the saturation adjustment, and τ is our differentiable filmic curve with further shadow $s \in (0, 1)$ and highlight $h \in (0, 1)$ compression controls:

$$\tau(x; s, h) = \begin{cases} \frac{s}{2} \exp\left(\frac{x+1-s}{s}\right) & \text{if } x < s-1 \\ \frac{1+x}{2} & \text{if } s-1 \leq x < 1-h \\ 1 - \frac{h}{2} \exp\left(\frac{-(x+h-1)}{h}\right) & \text{if } x \geq 1-h. \end{cases} \quad (8)$$

When $s, h \rightarrow 1$, τ approximates a sigmoid (equivalent to Photographic TMO in log-space [Reinhard et al. 2002]). When s or $h \rightarrow 0$, it produces a sharp toe and shoulder. This two-parameter function shape is analogous to common filmic operators [Hable 2010, 2017; SMPTE 2012]. Figure 5 shows the tone-mapping operator for three different sets of parameters s and h .

This tone-mapper is well-suited for gradient-based training. If we omit the constant factor from the contrast and saturation adjustments for brevity, it is possible to bound the derivatives of our TMO as:

$$0 < \frac{\partial \tau}{\partial \log L_t^{\text{denoised}}} < \frac{1}{2}, \quad (9a) \quad 0 < \frac{\partial \tau}{\partial L_t^{\text{denoised}}} < \frac{e}{2}. \quad (9b)$$

Such bounded derivatives avoid gradient explosions or vanishing, ensuring robust training.

Discussion. Although we do not claim state-of-the-art perceptual quality, our family of TMOs is comparable to widely used operators and suffices as a robust training surrogate. Importantly, the

tone-mapping used at test time can differ from the differentiable TMOs employed during training. The sampler learns to respond to tone-mapped appearance rather than to a specific operator, and the proposed TMOs are sufficiently expressive to generalize across different tone-mapping choices. The inference-time TMO does not need to be differentiable, as we do not rely on inference-time back-propagation. For HDR-only workflows, one could instead use the display transfer function.

Perception-based loss. Per-pixel error metrics such as L1, L2, and contrast-normalized variants including the symmetric mean absolute percentage error (SMAPE) are commonly used in image reconstruction tasks such as denoising and superresolution (Section 2). However, when employed to guide adaptive sampling, such losses can be suboptimal, as they assign equal importance to all errors regardless of their visual impact. In dark image regions, Weber-like behavior breaks down, and visual sensitivity is instead governed by the De Vries–Rose law [Van Nes and Bouman 1967], causing these losses to severely overestimate perceptual error and thus oversample such regions. This effect is particularly pronounced for SMAPE, whose Michelson-contrast-like formulation yields small denominators and can indicate near-maximal errors even for barely visible distortions. Conversely, at high luminance levels, all these losses ignore HVS gain control and compressive response [Shapley and Enroth-Cugell 1984], again inflating reported errors. While logarithmic scaling can mitigate oversampling at high luminance, it further exacerbates excessive sample allocation in dark regions. Moreover, per-pixel losses neglect contrast masking and spatial pooling effects, whereby errors become less visible in textured regions [Foley 1994; Legge and Foley 1980]. In particular, visual masking can permit substantial reductions in required sample density in the presence of strong contrast patterns [Myszkowski 1998; Ramasubramanian et al. 1999]. These shortcomings may be further aggravated by tone-mapping, which can significantly alter local luminance and contrast relationships, further decoupling loss predictions from perceived error.

To address these limitations, we instead employ perceptual losses computed on tone-mapped pixel values. Specifically, we adopt the MILO metric [Çoğalan et al. 2025], which implicitly captures masking effects and models distortion visibility across both dark and bright regions. MILO is trained on distortion types relevant to our setting, including residual noise from denoising, over-blurring due to undersampling, and hue shifts similar to those introduced by tone-mapping. As MILO is trained on human mean opinion score (MOS) data, it weights errors according to their perceived visibility and annoyance, while downweighting distortions below perceptual thresholds. This task-informed training places MILO closer to perception-based error visibility prioritization, in contrast to SSIM/MS-SSIM [Wang et al. 2004, 2003], which emphasizes structural fidelity, and LPIPS [Zhang et al. 2018], which focuses on visual similarity. Finally, from a practical standpoint, the visibility maps produced by MILO are more efficient to backpropagate than LPIPS feature activations, resulting in more efficient training. Moreover, prior applications of MILO as a loss function in image restoration tasks demonstrate its ability to reconstruct high-frequency content, which is desirable in our setting.

6 Implementation

6.1 Gathering pyramidal filter

Many recent neural denoisers rely on kernel prediction rather than direct image synthesis. The predicted kernels define how noisy input samples are filtered to produce the final denoised image. As illustrated in Fig. 6, kernel-based denoising can be formulated in two main ways: scattering, where each noisy pixel distributes its contribution to neighboring output pixels, and gathering, where each output pixel aggregates contributions from neighboring noisy input pixels.

With scattering, the network predicts a kernel at each sampled pixel that spreads its value to the surrounding area. When samples are dense, kernels overlap, and every output pixel receives contributions from multiple sources. However, as sampling becomes sparse, the network must predict meaningful kernels only at the sparse set of sampled locations, while the majority of kernel predictions go unused. This creates a complex task for the network: it must learn drastically different behavior depending on whether a pixel is sampled or not.

Gathering transposes this relationship. The network predicts a kernel at each output pixel that reaches into the neighborhood to collect available samples. Even when samples are sparse, every output location has an associated kernel that controls how nearby information is integrated. The prediction task is uniformly distributed across all pixels, and the network learns a consistent strategy for interpolating sparse data. This formulation remains more stable when the average sampling rate drops below one sample per pixel, making it well-suited to the sparsity targeted in our work.

We adopt a multi-scale pyramid structure to provide a sufficient receptive field for reconstruction. Prior work [Bálint et al. 2023] constructs the pyramid through partitioning, a scatter-based operation, and uses 4×4 scatter kernels for coarse-to-fine upsampling. We instead build the pyramid via plain downsampling and use 2×2 gather kernels for upsampling, maintaining consistency with our gathering formulation throughout. At each of the five pyramid levels, the denoiser predicts spatially varying 5×5 gather kernels for denoising. Reconstruction proceeds coarse-to-fine. At full resolution, we additionally predict 5×5 temporal kernels that gather from the warped previous HDR output. Full equations and implementation details are provided in Appendix B.

6.2 Learned demodulation

Kernel-based denoisers produce outputs that are weighted averages of neighboring input pixels. This constrains the output color at each pixel to lie within the convex hull of its neighbors' input colors. For illumination, this is not problematic, but high-frequency textures with saturated colors often require values outside this hull.

Prior albedo demodulation addresses this by dividing out the material's reflectance before filtering, then remodulating afterward. The denoiser operates on the illumination component, while texture detail is preserved through the albedo. However, this approach requires the renderer to provide mathematically correct demodulation maps for diffuse and specular components. Such decomposition is well-defined for simple materials but becomes problematic for volumes, measured BRDFs, and neural materials. It also fails when




Fig. 6. **Gathering versus scattering kernels.** Gathering kernels weight the neighboring sparse samples for every output pixel. Sparsity is distributed *within* the kernels, ensuring an even distribution of difficulty for the network’s prediction task. Scattering kernels scatter the sparse samples to neighboring output pixels. Sparsity is distributed *across* the kernels; kernels at sampled pixels are dense and solely responsible for the output image, while the majority of kernels at unsampled pixels have no effect.

the albedo buffer is noisy or aliased, or when effects like motion blur and depth-of-field blend multiple surface points into a single pixel.

We sidestep these issues by letting the network learn its own per-pixel demodulation. Rather than requiring a precise albedo from the renderer, we provide a rich texture buffer that combines diffuse color, specular tint, and Fresnel information. The denoiser predicts a three-channel demodulation map and uses it to demodulate the noisy input before filtering. After reconstruction, we remodulate using the predicted map. This learned decomposition adapts to arbitrary materials without renderer-side changes and can also handle imperfect GBuffer inputs.

Kernel filtering is fundamentally a neighborhood operation where each output pixel is derived from its surrounding input pixels. The demodulation map also provides per-pixel control: the brightness adjustment at each location depends only on the predicted map value at that location, not on neighboring pixels. We observe that the network learns to encode some geometric shading into this map, producing darkening at corners and crevices that visually resembles ambient occlusion (Fig. 7). Seemingly, such sharp, geometry-dependent variations could be more easily expressed as predicted per-pixel multipliers, rather than through the blending of noisy neighbors with intricately shaped kernels.

6.3 Architecture and training

Both the sampler and denoiser are implemented as U-Net encoder-decoders. Following Salehi et al. [2022], we adopt a global summary module at the sampler’s bottleneck to aggregate image-wide statistics. We chose to use the 15M-parameter architecture proposed by Bálint et al. [2023], a version of which also serves as the highest-quality preset for OIDN [Áfra 2024]. This well-tested architecture lets us evaluate the potential of our contributions under ideal conditions, keeping network limitations out of scope. Nonetheless, we show performance with a real-time network [Thomas et al. 2022] and discuss inference performance in Section 8.1.

We train a single model covering the entire sampling range from 0.11 spp to 4 spp. The discretization of the sample density map described in Section 4 is implemented using blue noise dithering [Ulichney 1993]. We use Falcor 8.0’s default path tracer [Kallweit et al. 2022] with global illumination, a maximum of 10 bounces, and Russian roulette termination. We pre-render samples up to 256 spp to allow flexible budget selection during training. The training set extends Noisebase [Bálint et al. 2023], built from procedural augmentations of 8 outdoor and 8 indoor scenes. The dataset contains 1024 training sequences of 64 frames each. Training runs for 256 epochs on a single H100 GPU, taking approximately 5 days. We use a batch size of 8 and train with 16-bit mixed-precision using the AdamW optimizer [Loshchilov and Hutter 2019] with $\beta = (0.8, 0.985)$ and a weight decay of 0.02. We decay the learning rate following a cosine annealing schedule [Loshchilov and Hutter 2017].

We train on a sliding window of two consecutive frames to capture temporal dynamics. The loss operates in tone-mapped sRGB space, combining a spatial term with a motion-compensated temporal term:

$$\mathcal{L}_{\text{spatial}} = |I^{\text{denoised}} - I^{\text{ref}}| \cdot m, \quad (10)$$

where m is a visual masking term generated by the MILO perceptual metric [Çoğalan et al. 2025], emphasizing visually more prominent absolute errors. The temporal term, following prior works [Bálint et al. 2023; Fan et al. 2025; Thomas et al. 2022], penalizes frame-to-frame inconsistency:

$$\mathcal{L}_{\text{temporal}} = |\Delta I^{\text{denoised}} - \Delta I^{\text{ref}}|, \quad (11)$$

where Δ denotes the difference between the current frame and the warped previous frame. The final loss takes the maximum of both terms:

$$\mathcal{L} = \max(1.25 \cdot \mathcal{L}_{\text{temporal}}, \mathcal{L}_{\text{spatial}}). \quad (12)$$

We expect such temporal losses to be superseded by video generalizations of metrics like MILO in the near future, as we discuss in Section 8.3.

7 Evaluation

We evaluate our method across five aspects: comparison with prior work, budget scaling behavior, resolution scaling, the advantages of sparse sampling over superresolution, and component ablations. Our experiments demonstrate that adaptive sparse sampling enables high-quality denoising at budgets where previous methods fail, and each of our contributions provides improvements. We also analyze the robustness of adaptive sampling, showing that it reliably improves perceptually important regions without introducing failures.

Setup. We evaluate across a range of rendering budgets from 0.11 spp (equivalent to 33% superresolution scaling) to 4 spp. Our test set comprises 9 animated scenes with 160–300 frames each, rendered at 1440p unless otherwise noted. The test scenes are held out: training scenes use separate assets from different studios. For SDR evaluation and supplementary videos, we use the ACES [Hill 2022] tone-mapper. None of the compared methods are trained for this specific evaluation-time tone-mapper, except RLSNAS, which is trained on tone-mapped images. Reference images are rendered at 6144 spp.




Fig. 7. **Learned demodulation and effect of gather vs scatter pyramids.** Our sampler network receives the renderer’s albedo buffer and the sparse, noisy input. Rather than requiring a mathematically correct albedo decomposition, our method learns a per-pixel demodulation map that adapts to the scene content. The learned map also encodes geometric shading cues resembling ambient occlusion, which are difficult to reconstruct through kernel filtering alone. *Bottom left:* Using learned demodulation (w/ demod.) better preserves texture saturation and helps reduce color bleeding compared to only kernel-based denoising (w/o demod.). *Bottom center and bottom right:* Our gather-based pyramid produces cleaner reconstructions than the scatter-based alternative, which exhibits ringing and oversharpening artifacts around some structures.




Fig. 8. **Main comparison at 0.25 spp.** We compare our method against baselines at a 50% superresolution equivalent budget. Our adaptive method preserves sharp edges and fine texture that superresolution approaches (DLSS, JNDS) lose due to low-resolution rendering. Other baselines also blur details more than our method. The density maps show our learned sample allocation: blue indicates fewer samples than uniform, red indicates more. PSNR for the full-resolution frame is shown in the bottom-right corner of the orange inset.

Baselines. We compare against methods spanning three categories: superresolution, uniform denoising, and adaptive sampling.

DLSS 4 Ray Reconstruction [NVIDIA 2025] denoises and upsamples from low-resolution renderings with low-resolution GBuffers. It supports variable scaling factors, making it suitable for budget comparisons. We use the official implementation in RTXPT [NVIDIA 2023] 1.6.0 adapted to Falcor.

JNDS [Thomas et al. 2022] similarly denoises and upsamples from low resolution but only supports 50% superresolution scaling. As

no official implementation is available, we implement the method in our framework, following the authors’ guidance and using a single filter path for combined diffuse and specular lighting, as they recommended.

NPPD [Bálint et al. 2023] is a denoiser without adaptive sampling. We use the official implementation, retrained on sparse renderings with high-resolution GBuffers to match the inputs of our uniform version.

NTASD [Hasselgren et al. 2020] combines a denoiser with a learned sampler. We implement this method ourselves, as no official code is available, and reproduce the demonstrated performance at 4 spp. Training becomes unstable at lower budgets: at 2 spp, training diverged after 15–20 epochs across 10 attempts, from which we took the best-performing checkpoint. At 1 spp, training failed to complete even a single epoch.

SAUDC [Fan et al. 2025] extends adaptive sampling with mandatory samples for improved stability. Adapting the authors’ implementation to our setup, we found that training diverged after 30–40 epochs at 2 spp despite mandatory samples. We replaced deterministic mandatory samples with our stochastic formulation, which allowed training across our full budget range, though runs still diverged between epochs 80–100. We report results from the best checkpoint at epoch 83.

RLSNAS [Scardigli et al. 2023] uses reinforcement learning for sample allocation. The official implementation is designed for scene-specific training on simple scenes. We trained it on our full dataset until the RL framework’s convergence criterion was met under the official configuration, but generalization to our test scenes remains poor.

Our method, across hundreds of training runs during experimentation, did not diverge once, except for some temporal-loop and gradient-estimator ablations noted below.

Metrics. We report six complementary metrics. PSNR measures pixel-wise error. MS-SSIM [Wang et al. 2004] and HaarPSI [Reisenhofer et al. 2018] assess structural similarity at multiple scales. MILO [Çoğalan et al. 2025] predicts perceptual quality based on human opinion scores. CVVDP [Mantiuk et al. 2024]² models visual differences in videos, accounting for contrast sensitivity and masking, and we report the SDR version unless otherwise noted. CGVQM [Jindal et al. 2025] evaluates video quality across learned features, including temporal aspects. Higher values indicate better quality for all metrics.

7.1 Main comparison

We first compare in the 50% superresolution setting, the most common target in prior work. This corresponds to 0.25 spp sparse sampling at native resolution, or equivalently, upscaling from 720p to 1440p for superresolution methods. Table 1 presents results across all metrics.

Our method consistently outperforms all baselines by a substantial margin. Among superresolution approaches, DLSS achieves 22.88 dB PSNR while JNDS reaches only 21.92 dB, both limited by the loss of high-frequency detail inherent to low-resolution rendering. Sparse uniform sampling with high-resolution GBuffers outperforms superresolution: NPPD achieves 23.92 dB, while our uniform baseline reaches 24.64 dB. Adaptive sampling provides further gains, with our full method achieving 25.41 dB PSNR and 7.056 JOD CVVDP. The HDR CVVDP³ scores generally follow the SDR scores; the largest deviations occur for methods that directly predict display-referred images and then recover HDR radiance

²ColorVideoVDP v0.5.0, 50.4 [pix/deg], Lpeak=200, Lblack=0.2, Lrefl=0.3979 [cd/m²]

³ColorVideoVDP v0.5.0, 50.4 [pix/deg], Lpeak=1500, Lblack=0.0015, Lrefl=0.01592 [cd/m²]

Table 1. **Comparison to baseline methods.** 50% superresolution equivalent (0.25 spp sparse). We report both (SDR/HDR) CVVDP versions. Our method outperforms all baselines across all metrics. Higher is better for all metrics.

Method	PSNR↑	MS-SSIM↑	HaarPSI↑	MILO↑	CVVDP↑	CGVQM↑
<i>Superresolution w/ LR inputs</i>						
DLSS	22.88	0.8950	0.6195	2.506	6.349/6.042	45.01
JNDS	21.92	0.8657	0.5370	2.276	5.621/5.508	28.87
<i>Sparse uniform sampling w/ HR inputs</i>						
NPPD	23.92	0.9062	0.6139	2.534	6.393/6.407	47.02
Ours Uniform	24.64	0.9183	0.6532	2.640	6.753/6.745	55.21
<i>Sparse adaptive sampling w/ HR inputs</i>						
RLSNAS	18.33	0.8077	0.3969	1.986	3.906/–	–9.08
SAUDC	23.09	0.8848	0.5645	2.250	5.965/4.993	38.97
Ours Adaptive	25.41	0.9273	0.6815	2.677	7.056/6.972	57.63

through inverse tone-mapping, because HDR evaluation exposes errors that are compressed by the SDR tone-map. A per-scene breakdown of the same 0.25 spp comparison is provided in Table 10.

Visual comparisons in Fig. 8 reveal that our method recovers fine details that competitors either blur or reconstruct with visible artifacts. The density maps show that our learned allocation strategy concentrates samples in such error-prone regions.

7.2 Budget scaling

We evaluate scaling across two regimes: sub-1-spp (Table 2) and above-1-spp (Table 3). Budget levels correspond to DLSS presets for consistent comparison.

Table 2. **Sub-1-spp regime comparison.** Superresolution equivalent shown in parentheses. We show how our method and baselines scale as we lower the rendering budget. (PSNR↑ / CVVDP↑)

Method	0.44 spp (67%)	0.34 spp (58%)	0.25 spp (50%)	0.11 spp (33%)
DLSS	23.50/6.64	23.19/6.50	22.88/6.35	21.74/5.76
NPPD	24.57/6.76	24.26/6.59	23.92/6.39	22.96/5.85
Ours Uniform	25.46/7.14	25.07/6.96	24.64/6.75	23.40/6.18
SAUDC	23.54/6.23	23.33/6.10	23.09/5.97	22.39/5.56
Ours Adaptive	26.15/7.39	25.80/7.23	25.41/7.06	24.26/6.56

In the sub-1-spp regime, our method maintains stable quality with graceful degradation down to 0.11 spp (one sample per nine pixels on average). At this extreme sparsity, we achieve 24.26 dB PSNR compared to 21.74 dB for DLSS and 22.39 dB for SAUDC.

Table 3. **Above-1-spp regime comparison.** We show how our method and baselines scale to larger budgets. (PSNR↑ / CVVDP↑)

Method	1 spp	2 spp	3 spp	4 spp
NPPD	25.44/7.23	26.08/7.56	26.47/7.75	26.70/7.86
Ours Uniform	26.41/7.57	27.09/7.89	27.45/8.06	27.69/8.17
NTASD	–	18.34/3.98	24.84/6.77	25.91/7.27
SAUDC	24.09/6.57	24.45/6.81	24.61/6.93	24.71/7.00
Ours Adaptive	27.10/7.80	27.82/8.11	28.23/8.28	28.51/8.39




Fig. 9. **Visual comparison across budgets and scenes.** We compare our method against baselines at 4.0 spp (top), 0.444 spp (middle), and 0.111 spp (bottom). Our adaptive method recovers fine details that competitors blur. The sample density maps show our learned sample allocation: blue indicates fewer samples than uniform, red indicates more. PSNR for the full-resolution frame is shown in the bottom-right corner of the orange inset. Our sampler concentrates samples on perceptually important regions such as edges and specular reflections while reducing allocation in smooth areas.

Above 1 spp, adaptive allocation continues to provide consistent gains. At 4 spp, our method achieves 28.51 dB PSNR versus 27.69 dB for our uniform version and 26.70 dB for NPPD.

In addition, Fig. 9 presents visual comparisons at average sampling rates of 4 spp, 0.444 spp, and 0.111 spp. Compared to other methods, both our uniform and adaptive variants consistently produce higher visual fidelity across all budgets. The benefit of adaptive sampling is especially apparent in fine details. For instance, it more accurately reconstructs the grid colors in the top scene, preserves specular highlights in the first crop of that scene, and improves detail recovery in the second crop of the last scene.

7.3 Resolution scaling

A key advantage of sparse sampling over superresolution is favorable scaling with target resolution. We evaluate this effect by rendering four of our test scenes at 1080p, 1440p, and 4K, comparing our method against DLSS at both 912k and 2M total samples (Table 4).

Adaptive sampling maintains its advantage over uniform sampling across all conditions, with both methods outperforming DLSS

by a wide margin. The advantage of adaptive over uniform is approximately 0.7 dB in PSNR, stable across resolutions and budgets.

As resolution increases at a fixed budget, per-pixel metrics decrease, as expected, since each pixel receives fewer samples while the resolution of the reference images increases. This effect is more pronounced for our CNN-based networks than for DLSS, whose transformer architecture appears to scale more favorably with increasing resolution. However, the benefit of additional samples also grows with resolution: at 1080p, doubling the budget from 912k to 2M samples improves our adaptive method by 1.04 dB PSNR, while at 4K the same increase yields 1.28 dB. DLSS shows a similar trend (0.84 dB to 1.09 dB).

7.4 Sparse sampling versus superresolution

To understand why sparse sampling outperforms superresolution at equivalent budgets, we run specific ablations at 0.25 spp (Table 5).

We first test our uniform method with low-resolution inputs bilinearly upsampled to full resolution, matching the information available to superresolution methods. This includes path-traced

Table 4. **Comparison across resolutions.** We compare scaling across 3 resolutions, 2 budgets, and 6 metrics. Our adaptive method consistently outperforms our uniform version, which in turn outperforms the baseline DLSS. Comparison across resolutions also depends on the metric’s scaling, since the reference images have different resolutions.

Method	PSNR↑	MS-SSIM↑	HaarPSI↑	MILO↑	CVVDP↑	CGVQM↑
<i>1080p 1 spp (2M samples)</i>						
Ours Adaptive	27.69	0.9565	0.7667	3.003	7.890	69.25
Ours Uniform	27.00	0.9521	0.7427	2.901	7.665	67.28
DLSS	24.56	0.9258	0.6972	2.767	6.909	60.10
<i>1440p 0.44 spp (2M samples)</i>						
Ours Adaptive	26.65	0.9408	0.7135	2.781	7.648	63.81
Ours Uniform	26.01	0.9340	0.6859	2.707	7.363	61.91
DLSS	23.99	0.9154	0.6564	2.594	6.886	54.55
<i>4K 0.25 spp (2M samples)</i>						
Ours Adaptive	26.48	0.9340	0.6861	2.717	7.787	63.08
Ours Uniform	25.80	0.9251	0.6557	2.639	7.445	60.92
DLSS	23.87	0.9120	0.6283	2.524	6.955	52.58
<i>1080p 0.44 spp (912k samples)</i>						
Ours Adaptive	26.65	0.9451	0.7272	2.842	7.486	64.62
Ours Uniform	25.96	0.9369	0.6995	2.754	7.187	62.16
DLSS	23.72	0.9132	0.6544	2.594	6.567	52.83
<i>1440p 0.25 spp (912k samples)</i>						
Ours Adaptive	25.84	0.9299	0.6817	2.676	7.315	60.09
Ours Uniform	25.11	0.9178	0.6506	2.595	6.916	57.48
DLSS	23.41	0.9056	0.6240	2.485	6.637	48.55
<i>4K 0.11 spp (912k samples)</i>						
Ours Adaptive	25.20	0.9179	0.6402	2.580	7.285	57.61
Ours Uniform	24.43	0.9016	0.6048	2.491	6.797	54.29
DLSS	22.78	0.8927	0.5674	2.354	6.522	42.74

samples, normals, albedo, and motion vectors, all at reduced resolution. Quality decreases (from 24.64 dB to 24.16 dB in PSNR), yet remains state-of-the-art compared to all baselines. This suggests that our contributions beyond adaptive sampling—the gather pyramid, learned demodulation, and TMO-aware training—could benefit superresolution pipelines.

Table 5. **Sparse sampling versus superresolution.** We isolate factors contributing to sparse sampling’s advantage at 0.25 spp (50% superresolution equivalent). High-resolution GBuffers recover most of the quality gap, while density map resolution has small impact on aggregate metrics.

Method	PSNR↑	MS-SSIM↑	HaarPSI↑	MILO↑	CVVDP↑	CGVQM↑
Ours Uniform	24.64	0.9183	0.6532	2.640	6.753	55.21
w/ LR input	24.16	0.9088	0.6304	2.498	6.559	50.52
w/ LR inp., HR GB.	24.40	0.9140	0.6413	2.609	6.641	54.00
Ours Adaptive	25.41	0.9273	0.6815	2.677	7.056	57.63
w/ LR density map	25.32	0.9265	0.6772	2.681	7.033	56.90

Restoring high-resolution GBuffers while keeping low-resolution radiance samples and motion recovers much of the quality (24.40 dB PSNR). High-resolution normals, depth, and albedo provide valuable guidance for reconstruction even when the radiance signal is sparse, suggesting that rendering with high-resolution GBuffers is a practical way to boost sharpness with minimal overhead, as we show in Section 8.2.

For adaptive sampling, we test the impact of density map resolution by downsampling predictions to quarter resolution before stochastic sampling. The effect on aggregate metrics is small, but visual inspection reveals more accurate sampling of specular highlights and contact shadows at full resolution. These features occupy

a small fraction of pixels, explaining the minimal metric difference, though they may contribute disproportionately to perceived quality. Running the sampler network at reduced resolution remains a viable trade-off in terms of efficiency.

7.5 Component ablations

We validate each proposed component through ablations (Table 6), evaluating both our adaptive and uniform models to isolate contributions.

Table 6. **Component ablations.** We validate each proposed component by replacing or removing it from our adaptive and uniform models. Loss formulation and recurrent features have the largest impact; our gradient estimator provides the best stability-accuracy balance. (PSNR↑ / CVVDP↑)

Method	0.11 spp	0.25 spp	1 spp	4 spp
Ours Adaptive	24.26/6.56	25.41/7.06	27.10/7.80	28.51/8.39
<i>Gradient variants</i>				
DASR surrogate	23.58/6.23	25.00/6.83	26.86/7.67	28.28/8.31
Straight-through	24.32/6.51	25.40/7.01	27.02/7.76	28.41/8.37
Gumbel-Softmax	24.09/6.42	25.32/6.99	27.10/7.80	28.56/8.41
<i>Component replacements</i>				
SMAPE loss	23.25/6.18	24.57/6.71	26.46/7.51	28.02/8.16
Photographic TMO loss	23.49/6.28	24.68/6.81	26.53/7.62	28.18/8.31
Partitioning pyramid	23.68/6.38	25.15/6.98	27.13/7.81	28.76/8.49
w/o learned demod.	24.33/6.59	25.53/7.10	27.20/7.84	28.62/8.44
Separate rec. features	23.96/6.35	25.03/6.82	26.65/7.62	27.98/8.26
w/o rec. features	23.47/6.10	24.63/6.59	26.30/7.41	27.72/8.11
Ours Uniform	23.40/6.18	24.64/6.75	26.41/7.57	27.69/8.17
<i>Uniform ablations</i>				
SMAPE loss	22.61/5.95	23.91/6.47	25.74/7.31	27.12/7.97
Photographic TMO loss	23.38/6.16	24.68/6.74	26.45/7.55	27.81/8.19
Partitioning pyramid	22.97/6.11	24.37/6.72	26.28/7.58	27.59/8.20
w/o learned demod.	23.49/6.29	24.76/6.86	26.50/7.63	27.79/8.22

Gradient estimation. We compare four gradient estimation strategies for adaptive sampling. With the DASR surrogate [Kuznetsov et al. 2018], training was unstable with only one of five runs succeeding. The resulting quality is also worse due to the biased gradients that guide the optimization toward suboptimal sample distributions. Straight-through estimation also proved unstable (three of five runs succeeded), but surprisingly achieves comparable quality when it converges, suggesting that its gradient bias may act as a regularizer, as observed in other domains [Bengio et al. 2013]. Gumbel-Softmax trains stably but suffers from train-test mismatch as discussed in Section 4.3, slightly degrading quality at low sample counts where the contribution scaling in Eq. (4) amplifies discrepancies. Our relaxed formulation achieves the best balance of stability and accuracy.

Loss formulation. Replacing our tone-mapping-aware pipeline with the SMAPE-based loss from NPPD [Bálint et al. 2023] substantially degrades quality for both adaptive and uniform variants, demonstrating the importance of perceptual alignment for training. Using a fixed photographic TMO instead of our randomized parametric family also hurts the adaptive model, confirming that adaptive sampling adjusts to tone-mapped appearance and requires diverse TMOs for generalization. The uniform model is less affected since it cannot control sampling or observe the tone-mapped image.

Filter architecture. Replacing our gather pyramid with NPPD’s scatter-based partitioning pyramid degrades quality, particularly at low sample counts, confirming our analysis in Section 6.1 that scatter-based approaches are poorly suited to sparse inputs.

Learned demodulation. Removing learned demodulation slightly improves aggregate metrics. However, this component is essential for resolving color artifacts in regions with saturated albedos, as shown in Fig. 7. Deriving more robust formulations holds potential for future work.

Recurrent features. Separating the recurrent feature spaces for the sampler and denoiser degrades quality and increases training variance, suggesting the shared latent space enables beneficial collaboration between networks. Removing recurrent features entirely led to further degradation and needed reruns due to divergence, confirming the importance of temporal context for stability and performance.

7.6 Robustness analysis

A fundamental concern with adaptive sampling is reliability: since the total sample budget is fixed, improving quality in some regions necessarily degrades it in others. Within our exploration-exploitation interpretation, the sampler must reliably identify which regions require more samples and which can afford fewer. We demonstrate that our method addresses regions with the most visible distortions while taking samples only from regions where the impact is perceptually negligible.

Figure 10 presents a balanced comparison. For the given test frame, we identify the patch where adaptive sampling degrades quality the most compared to uniform sampling (measured by mean absolute error, MAE), as well as the patch where adaptive sampling improves quality the most. The figure shows both extremes: regions where adaptive sampling wins, such as fine geometric detail on tree leaves, and regions where it loses to uniform sampling, such as the pavement.

The key observation is that adaptive sampling wins where it matters. The patches showing improvement contain perceptually critical content, while the patches showing degradation are invariably smooth regions, such as diffuse surfaces, where error is difficult to perceive, as contrast is low.

The accompanying error-versus-density plots quantify this relationship. Our method predicts high density where reconstruction error would otherwise be large, thereby substantially reducing it. In regions predicted to require few samples, the error is consistently low as the sampler identifies simpler regions with high confidence. Adaptive sampling primarily loses to uniform sampling in intermediate regions where error is present but modest; here, the density map is only slightly below average, reflecting appropriate caution. We provide similar analyses for additional test scenes in Appendix C.

8 Discussion

8.1 Inference cost

We have not investigated specific network architectures in depth, as inference performance is highly hardware-dependent and lightweight architectures remain a rapidly evolving field. Recent developments in low-precision quantized inference and knowledge distillation from larger to smaller models [Kong et al. 2025] show considerable promise for enabling better real-time neural denoising.

Unfortunately, state-of-the-art networks are not widely available or well-tested across methods. DLSS, for instance, employs a seemingly quite capable transformer architecture, but its architectural details are not public. Developing a custom real-time architecture was beyond the scope of this work.

Nonetheless, we recognize that some readers may question the extent to which real-time networks will match the capacity of the 15M-parameter architecture used in our main evaluation before full real-time path tracing becomes widespread. To address this concern, we also evaluate our method using the network from JNDS [Thomas et al. 2022], a 2.6M-parameter architecture designed for real-time inference on current Intel GPUs.

For a point of reference based on our research prototype, we measured TensorRT INT8 inference at 1440p on an NVIDIA RTX 3090, averaged over 2000 frames. The 15M-parameter architecture used in our main evaluation runs in 48.716 ms/frame, while the 2.6M-parameter JNDS architecture used in the lightweight ablations runs in 4.288 ms/frame.

Our gather-based denoiser has cost comparable to prior multiscale filters, and the gather operations can be fused with the network’s output layers at negligible overhead. Tone-mapping-aware training adds no inference cost: the differentiable TMO and MILO are used only during optimization, and MILO is among the lighter perceptual losses, with roughly 5× faster backpropagation than LPIPS [Çoğalan et al. 2025].

Table 7. **Ablations using real-time network.** 50% superresolution equivalent (0.25 spp sparse). Higher is better for all metrics.

Method	PSNR↑	MS-SSIM↑	HaarPSI↑	MILO↑	CVVDP↑	CGVQM↑
Ours Uniform	24.64	0.9183	0.6532	2.640	6.753	55.21
Mini Uniform	24.16	0.9050	0.6160	2.510	6.443	49.46
NPPD	23.92	0.9062	0.6139	2.534	6.393	47.02
Ours Adaptive	25.41	0.9273	0.6815	2.677	7.056	57.63
Mini Sampler	25.37	0.9243	0.6791	2.735	6.961	57.70
Mini Adaptive	24.97	0.9192	0.6528	2.637	6.820	54.30

As shown in Table 7, our method scales well to reduced network capacity. The Mini Uniform variant, using the 2.6M-parameter JNDS network, outperforms the strongest baseline, NPPD, on most metrics, despite NPPD employing the same 15M-parameter architecture as our main evaluation. Mini Sampler, which substitutes only the sampler network while retaining the large denoiser, incurs only a 0.04 dB PSNR loss. Mini Adaptive, with both a lightweight sampler and denoiser, still surpasses all baselines and performs comparably to our uniform model using the larger network. These results indicate that adaptive sampling continues to provide meaningful improvements, even with architectures that are real-time-capable




Fig. 10. **Best- and worst-case robustness analysis.** This figure compares, for the same scene, the best- and worst-performing regions of our adaptive sampling method relative to uniform sampling. The green inset highlights the crop where uniform sampling performs best, while the orange inset highlights the crop where adaptive sampling performs best. The comparison shows that adaptive sampling better reconstructs fine details, while only slightly underperforming in more uniform regions. The plot on the right illustrates the mean absolute error as a function of relative sampling rate. When adaptive sampling allocates fewer samples than uniform sampling (left side of the plot, sampling ratio < 1), the error increases only marginally. In contrast, when adaptive sampling allocates more samples than uniform sampling (right side, sampling ratio > 1), the error reduction is substantial. This demonstrates that the sampler successfully focuses computation on high-error regions, leading to an overall improvement in image quality.

today. Full-resolution inference is also a common design choice in related work [Kazmierczyk et al. 2025; Xiao et al. 2020]. In practice, this requires keeping the native-resolution stage lightweight, as in the mini network [Thomas et al. 2022].

8.2 GBuffer and rendering time

Decreasing reliance on GBuffers is a promising direction of current research. For example, Wu et al. [2024] propose GBuffer-free extrapolation for high-framerate, low-latency rendering. Skipping rendering for the extrapolated frames reduces load on the game engine and saves time on preliminary tasks such as animation, BVH updates, and scheduling. However, scaling the GBuffer resolution only offers modest savings, as the GPU time spent on GBuffer rendering is relatively negligible to begin with. By comparison, tracing the scattered, incoherent rays for global illumination is inherently harder.

In our setup, one spp uniform path tracing is generally 20–30 times the cost of GBuffer rendering at the same resolution. That is despite Falcor using a relatively unoptimized ray-traced GBuffer pass. In a production game engine, we expect GBuffer rendering to be much faster. At 4K, an RGB8 albedo buffer, an RG8_SNORM octahedral normal buffer, and an R16 depth buffer amount to 56 bits per pixel, or 55.4 MiB in total.

Time-wise scaling is further imperfect, as adaptive sampling tends to focus on more challenging parts of the scene, for example, prioritizing sampling areas with complex lighting and geometry over flat regions. These regions sometimes slow down path tracing due to increased BVH and material complexity.

Table 8. **GBuffer timing comparison.** Cost of GBuffer and path tracing on several of our production-quality test scenes, measured on a workstation with an NVIDIA RTX 3090 at 1440p. PT adaptive does not include inference of the sample density.

Scene	Triangles	GBuffer (ms)	PT uniform (ms)	PT 1/4 spp adaptive (ms)	Budget
Bistro	4.2M	0.672	21.80	5.5	≈ 27%
ZeroDay	5.3M	0.892	14.42	3.6	≈ 29%
Sun Temple	600k	0.41	8.47	2.11	≈ 28%

Nonetheless, as shown in Table 8, the overhead of high-resolution GBuffers and adaptive sampling is marginal compared to the improvement over superresolution demonstrated in our work. We focus on sampling budget comparisons throughout most of our paper, as they are more practical and easier to follow. One may assume an approximately 5% overhead on our methods when translating sample count measurements to rendering time.

8.3 Temporal stability

We quantify temporal consistency using two measures computed on motion-compensated frame differences. tPSNR reports the PSNR between the temporal residuals of the reconstruction and the reference, so larger values indicate that frame-to-frame changes match the reference sequence more closely. trMAE reports the mean absolute error of the same residuals, so lower values indicate fewer temporal deviations. We include these metrics because they directly measure agreement with the finite-difference temporal objective used during training.

Table 9. **Temporal stability comparison at 0.25 spp.** Higher is better for tPSNR; lower is better for trMAE.

Metric	DLSS	JNDS	NPPD	Ours Uniform	SAUDC	RLSNAS	Ours Adaptive
tPSNR↑	22.02	20.05	21.83	22.64	21.26	19.69	23.36
trMAE↓	0.240	0.327	0.257	0.239	0.267	0.387	0.227

As shown in Table 9, our adaptive model achieves the highest aggregate tPSNR and lowest trMAE across the evaluation set. Compared with the uniform variant, adaptive sampling improves tPSNR from 22.64 to 23.36 and reduces trMAE from 0.239 to 0.227. This indicates that allocating samples with the learned density map helps the reconstruction match the temporal finite-difference objective, in addition to the spatial and perceptual improvements reported above.

At the same time, temporal stability remains a limitation at these budgets. The density maps show that the sampler allocates extra budget to regions where temporal reuse is least reliable, especially specular highlights and disocclusions, yet this does not eliminate temporal artifacts. On large glossy regions, such as the Cyber Market floor, we observe residual instability, likely because reflected motion is poorly captured by primary-surface motion vectors. In undersampled diffuse regions, we observe a slight lag or loss of temporal detail, but not a clear increase in flicker.

Qualitatively, DLSS appears to apply stronger temporal smoothing in some regions, but its proprietary implementation prevents us from direct comparison or analysis. Nonetheless, a better-aligned temporal objective would fit naturally into our pipeline and is an important direction for future work. Better reprojection or forward projection could reduce the burden on the sampler; the mini-sampler ablation suggests that more accurate temporal transport may enable lighter sample-density prediction, or even a unified sampler-denoiser design.

Future work could also combine adaptive sampling with sample reuse methods such as ReSTIR PT and reservoir splatting [Bitterli et al. 2020; Lin et al. 2022; Liu et al. 2025; Talbot et al. 2005], or explore continuous sample-position formulations connected to differentiable rendering [Li et al. 2018].

9 Conclusion

We presented the first end-to-end adaptive sampling and denoising pipeline designed for the sub-1-spp regime. Our stochastic formulation for sample placement enables stable gradient estimation at extreme sparsity, where prior deterministic approaches fail. Combined with tone-mapping-aware training using perceptual losses and a gather-based pyramidal filter suited to sparse inputs, our method consistently outperforms both superresolution and previous adaptive sampling approaches across all tested budgets.

Our results demonstrate that adaptive sampling remains effective even when the average budget falls below one sample per pixel. The learned sampler reliably concentrates samples on perceptually critical regions while safely reducing allocation in smooth areas where errors are difficult to perceive.

Acknowledgments

We thank Anton Kaplanyan, Anton Sochenov, and Attila Áfra for insightful discussions.

We gratefully acknowledge the support of the Saarland/Intel Joint Program on the Future of Graphics and Media.

We thank the following authors for the scenes used in our experiments: Evermotion (training scenes), Leartes Studios (Cyberpunk Rooftop Market, Village House), Meinel and Kaplanyan [2022] (Sponza), Amazon Lumberyard [2017] and Krzysztof Wolski (Bistro), SłykDrako and Bitterli [2016] (Bedroom), Winkelmann [2019] (ZeroDay).

References

- Attila T. Áfra. 2024. Intel® Open Image Denoise. <https://www.openimagedenoise.org>
- Amazon Lumberyard. 2017. Amazon Lumberyard Bistro, Open Research Content Archive (ORCA). <http://developer.nvidia.com/orca/amazon-lumberyard-bistro>
- Steve Bako, Thijs Vogels, Brian McWilliams, Mark Meyer, Jan Novák, Alex Harvill, Pradeep Sen, Tony DeRose, and Fabrice Rousselle. 2017. Kernel-predicting convolutional networks for denoising Monte Carlo renderings. *ACM Trans. Graph.* 36, 4, Article 97 (July 2017), 14 pages. doi:10.1145/3072959.3073708
- Martin Bálint, Krzysztof Wolski, Karol Myszkowski, Hans-Peter Seidel, and Rafał Mantiuk. 2023. Neural Partitioning Pyramids for Denoising Monte Carlo Renderings. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 60, 11 pages. doi:10.1145/3588432.3591562
- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation. arXiv:1308.3432 [cs.LG] <https://arxiv.org/abs/1308.3432>
- Benedikt Bitterli. 2016. Rendering resources. <https://benedikt-bitterli.me/resources/>
- Benedikt Bitterli, Fabrice Rousselle, Bochang Moon, José A. Iglesias-Guitián, David Adler, Kenny Mitchell, Wojciech Jarosz, and Jan Novák. 2016. Nonlinearly Weighted First-order Regression for Denoising Monte Carlo Renderings. *Computer Graphics Forum* 35, 4 (2016), 107–117. doi:10.1111/cgf.12954
- Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron Lefohn, and Wojciech Jarosz. 2020. Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. *ACM Trans. Graph.* 39, 4, Article 148 (Aug. 2020), 17 pages. doi:10.1145/3386569.3392481
- Mark R. Bolin and Gary W. Meyer. 1998. A perceptually based adaptive sampling algorithm. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. Association for Computing Machinery, New York, NY, USA, 299–309. doi:10.1145/280814.280924
- Chakravarty R. Alla Chaitanya, Anton S. Kaplanyan, Christoph Schied, Marco Salvi, Aaron Lefohn, Derek Nowrouzezahrai, and Timo Aila. 2017. Interactive reconstruction of Monte Carlo image sequences using a recurrent denoising autoencoder. *ACM Trans. Graph.* 36, 4, Article 98 (July 2017), 12 pages. doi:10.1145/3072959.3073601
- Uğur Coğalan, Mojtaba Bemana, Karol Myszkowski, Hans-Peter Seidel, and Colin Groth. 2025. MILO: A Lightweight Perceptual Quality Metric for Image and Latent-Space Optimization. *ACM Trans. Graph.* 44, 6, Article 214 (Dec. 2025), 11 pages. doi:10.1145/3763340
- Hangming Fan, Yuchi Huo, Chuankun Zheng, Chonghao Hu, Yazhen Yuan, and Rui Wang. 2025. Streaming-Aware Neural Monte Carlo Rendering Framework with Unified Denoising-Compression and Client Collaboration. *ACM Trans. Graph.* 44, 4, Article 164 (July 2025), 13 pages. doi:10.1145/3730879
- Arthur Firmino, Jeppe Revall Frisvad, and Henrik Wann Jensen. 2023. Denoising-Aware Adaptive Sampling for Monte Carlo Ray Tracing. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 32, 11 pages. doi:10.1145/3588432.3591537
- John M. Foley. 1994. Human luminance pattern-vision mechanisms: masking experiments require a new model. *J. Opt. Soc. Am.* A 11, 6 (jun 1994), 1710–1719.
- Weifeng Ge, Bingchen Gong, and Yizhou Yu. 2018. Image super-resolution via deterministic-stochastic synthesis and local statistical rectification. *ACM Trans. Graph.* 37, 6, Article 260 (dec 2018), 14 pages. doi:10.1145/3272127.3275060
- Michael Gharbi, Tzu-Mao Li, Miika Aittala, Jaakko Lehtinen, and Frédo Durand. 2019. Sample-based Monte Carlo denoising using a kernel-splatting network. *ACM Trans. Graph.* 38, 4, Article 125 (jul 2019), 12 pages. doi:10.1145/3306346.3322954
- John Hable. 2010. Uncharted 2: HDR Lighting. GDC 2010.
- John Hable. 2017. Filmic Tonemapping with Piecewise Power Curves. <https://filmicworlds.com/blog/filmic-tonemapping-with-piecewise-power-curves/>
- Jon Hasselgren, Jacob Munkberg, Marco Salvi, Anjul Patney, and Aaron Lefohn. 2020. Neural Temporal Adaptive Sampling and Denoising. *Computer Graphics Forum* 39, 2 (2020), 147–155. doi:10.1111/cgf.13919

- Stephen Hill. 2022. ACES Tone Mapping Operator. <https://github.com/TheRealMJP/BakingLab/blob/master/BakingLab/ACES.hsl>
- Nikolai Hofmann, Jon Hasselgren, and Jacob Munkberg. 2023. Joint Neural Denoising of Surfaces and Volumes. *Proc. ACM Comput. Graph. Interact. Tech.* 6, 1, Article 10 (May 2023), 16 pages. doi:10.1145/3585497
- Mustafa Işık, Krishna Mullia, Matthew Fisher, Jonathan Eisenmann, and Michael Gharbi. 2021. Interactive Monte Carlo Denoising Using Affinity of Neural Features. *ACM Trans. Graph.* 40, 4, Article 37 (jul 2021), 13 pages. doi:10.1145/3450626.3459793
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparameterization with Gumbel-Softmax. arXiv:1611.01144 [stat.ML] <https://arxiv.org/abs/1611.01144>
- Akshay Jindal, Nabil Sadaka, Manu Mathew Thomas, Anton Sochenov, and Anton Kaplanyan. 2025. CGVQM+D: Computer Graphics Video Quality Metric and Dataset. *Computer Graphics Forum* (2025). doi:10.1111/cgf.70221
- Nima Khademi Kalantari, Steve Bako, and Pradeep Sen. 2015. A machine learning approach for filtering Monte Carlo noise. *ACM Trans. Graph.* 34, 4, Article 122 (July 2015), 12 pages. doi:10.1145/2766977
- Simon Kallweit, Petrik Clarberg, Craig Kolb, Tomáš Davidovič, Kai-Hwa Yao, Theresa Foley, Yong He, Lifan Wu, Lucy Chen, Tomas Akenine-Möller, Chris Wyman, Cyril Crassin, and Nir Benty. 2022. The Falcor Rendering Framework. <https://github.com/NVIDIAGameWorks/Falcor>
- Pawel Kazmierczyk, Sungye Kim, Wojciech Uss, Wojciech Kalinski, Tomasz Galaj, Mateusz Maciejewski, and Rama Harihar. 2025. Joint Denoising and Upscaling via Multi-branch and Multi-scale Feature Network. *Proc. ACM Comput. Graph. Interact. Tech.* 8, 1, Article 16 (May 2025), 18 pages. doi:10.1145/3728297
- Mengxun Kong, Jie Guo, Chen Wang, Ye Yuan, and Yanwen Guo. 2025. Real-Time Neural Denoising with Render-Aware Knowledge Distillation. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 4 (apr 2025), 4392–4400. doi:10.1609/aaai.v39i4.32462
- Alexandr Kuznetsov, Nima Khademi Kalantari, and Ravi Ramamoorthi. 2018. Deep Adaptive Sampling for Low Sample Count Rendering. *Computer Graphics Forum* 37, 4 (2018), 35–44. doi:10.1111/cgf.13473
- Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. 2017. Deep Laplacian Pyramid Networks for Fast and Accurate Super-Resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5835–5843. doi:10.1109/CVPR.2017.618
- Junmin Lee, Seunghyun Lee, Min Yoon, and Byung Cheol Song. 2024. Real-Time Monte Carlo Denoising With Adaptive Fusion Network. *IEEE Access* 12 (2024), 29154–29165. doi:10.1109/ACCESS.2024.3369588
- Mark E. Lee and Richard A. Redner. 1990. A note on the use of nonlinear filtering in computer graphics. *IEEE Computer Graphics and Applications* 10, 3 (1990), 23–29. doi:10.1109/38.55149
- Gordon E. Legge and John M. Foley. 1980. Contrast masking in human vision. *J. Opt. Soc. Am.* 70, 12 (dec 1980), 1458–1471.
- Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. 2018. Differentiable Monte Carlo ray tracing through edge sampling. *ACM Trans. Graph.* 37, 6, Article 222 (Dec. 2018), 11 pages. doi:10.1145/3272127.3275109
- Tzu-Mao Li, Yu-Ting Wu, and Yung-Yu Chuang. 2012. SURE-based optimization for adaptive sampling and reconstruction. *ACM Trans. Graph.* 31, 6, Article 194 (nov 2012), 9 pages. doi:10.1145/2366145.2366213
- Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized resampled importance sampling: foundations of ReSTIR. *ACM Trans. Graph.* 41, 4, Article 75 (July 2022), 23 pages. doi:10.1145/3528223.3530158
- Jeffrey Liu, Daqi Lin, Markus Kettunen, Chris Wyman, and Ravi Ramamoorthi. 2025. Reservoir Splatting for Temporal Path Resampling and Motion Blur. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Papers (SIGGRAPH Conference Papers '25)*. Association for Computing Machinery, New York, NY, USA, Article 35, 11 pages. doi:10.1145/3721238.3730646
- Ilya Loshchilov and Frank Hutter. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. arXiv:1608.03983 [cs.LG] <https://arxiv.org/abs/1608.03983>
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. arXiv:1711.05101 [cs.LG] <https://arxiv.org/abs/1711.05101>
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. arXiv:1611.00712 [cs.LG] <https://arxiv.org/abs/1611.00712>
- Rafal K. Mantiuk, Param Hanji, Maliha Ashraf, Yuta Asano, and Alexandre Chapiro. 2024. ColorVideoVDP: A visual difference predictor for image, video and display distortions. *ACM Trans. Graph.* 43, 4, Article 129 (July 2024), 20 pages. doi:10.1145/3658144
- Michael D. McCool. 1999. Anisotropic diffusion for Monte Carlo noise reduction. *ACM Trans. Graph.* 18, 2 (apr 1999), 171–194. doi:10.1145/318009.318015
- Frank Meinel and Anton Kaplanyan. 2022. Intel Sample Library. <https://www.intel.com/content/www/us/en/developer/topic-technology/graphics-processing-research/samples.html>
- Don P. Mitchell. 1987. Generating antialiased images at low sampling densities. *SIGGRAPH Comput. Graph.* 21, 4 (aug 1987), 65–72. doi:10.1145/37402.37410
- Karol Myszkowski. 1998. The Visible Differences Predictor: applications to global illumination problems. In *Proceedings of the 9th Eurographics Workshop on Rendering (EGWR)*. Eurographics Association, Vienna, 223–236.
- NVIDIA. 2023. NVIDIA® RTX Path Tracing. <https://github.com/NVIDIA-RTX/RTXPT>
- NVIDIA. 2025. DLSS 4: Transforming Real-Time Graphics with AI. <https://research.nvidia.com/labs/adlr/DLSS4>
- Yaobin Ouyang, Shiqiu Liu, Markus Kettunen, Matt Pharr, and Jacopo Pantaleoni. 2021. ReSTIR GI: Path Resampling for Real-Time Path Tracing. *Computer Graphics Forum* 40, 8, 17–29. doi:10.1111/cgf.14378
- Ryan S. Overbeck, Craig Donner, and Ravi Ramamoorthi. 2009. Adaptive wavelet rendering. *ACM Trans. Graph.* 28, 5 (Dec. 2009), 1–12. doi:10.1145/1618452.1618486
- Mahesh Ramasubramanian, Sumanta N. Pattanaik, and Donald P. Greenberg. 1999. A perceptually based physical error metric for realistic image synthesis. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., USA, 73–82. doi:10.1145/311535.311543
- Erik Reinhard, Wolfgang Heidrich, Paul Debevec, Sumanta Pattanaik, Greg Ward, and Karol Myszkowski. 2010. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann, San Francisco, CA, USA.
- Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. 2002. Photographic tone reproduction for digital images. *ACM Trans. Graph.* 21, 3 (July 2002), 267–276. doi:10.1145/566654.566675
- Rafael Reisenhofer, Sebastian Bosse, Gitta Kutyniok, and Thomas Wiegand. 2018. A Haar wavelet-based perceptual similarity index for image quality assessment. *Signal Processing: Image Communication* 61 (2018), 33–43. doi:10.1016/j.image.2017.11.001
- Fabrice Rousselle, Marco Manzi, and Matthias Zwicker. 2013. Robust Denoising using Feature and Color Information. *Computer Graphics Forum* 32, 7 (2013), 121–130. doi:10.1111/cgf.12219
- Holly E. Rushmeier and Gregory J. Ward. 1994. Energy preserving non-linear filters. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. Association for Computing Machinery, New York, NY, USA, 131–138. doi:10.1145/192161.192189
- Farnood Salehi, Marco Manzi, Gerhard Roethlin, Romann Weber, Christopher Schroers, and Marios Papas. 2022. Deep Adaptive Sampling and Reconstruction Using Analytic Distributions. *ACM Trans. Graph.* 41, 6, Article 259 (nov 2022), 16 pages. doi:10.1145/3550454.3555515
- Antoine Scardigli, Lukas Cavigelli, and Lorenz K. Müller. 2023. RL-based Stateful Neural Adaptive Sampling and Denoising for Real-Time Path Tracing. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 66390–66407. https://proceedings.neurips.cc/paper_files/paper/2023/file/d1422213c9f2b5d578b77d166fba86a-Paper-Conference.pdf
- Robert Shapley and Christina Enroth-Cugell. 1984. Visual Adaptation and Retinal Gain Controls. In *Visual Perception: The Neurophysiological Foundations*, L. Spillman and J. S. Werner (Eds.). Academic Press, 53–101.
- SMPTE. 2012. SMPTE ST 2065-1, Academy Color Encoding Specification (ACES). (2012).
- Justin F. Talbot, David Cline, and Parris Egbert. 2005. Importance resampling for global illumination. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques (Konstanz, Germany) (EGSR '05)*. Eurographics Association, Goslar, DEU, 139–146.
- Manu Mathew Thomas, Gabor Liktó, Christoph Peters, Sungye Kim, Karthik Vaidyanathan, and Angus G. Forbes. 2022. Temporally Stable Real-Time Joint Neural Denoising and Supersampling. *Proc. ACM Comput. Graph. Interact. Tech.* 5, 3, Article 21 (jul 2022), 22 pages. doi:10.1145/3543870
- Robert A. Ulichney. 1993. Void-and-cluster method for dither array generation. In *Human Vision, Visual Processing, and Digital Display IV*, Jan P. Allebach and Bernice E. Rogowitz (Eds.), Vol. 1913. International Society for Optics and Photonics, SPIE, 332–343. doi:10.1117/12.152707
- Floris L. Van Nes and Maarten A. Bouman. 1967. Spatial Modulation Transfer in the Human Eye. *J. Opt. Soc. Am.* 57, 3 (1967), 401–406.
- Thijs Vogels, Fabrice Rousselle, Brian McWilliams, Gerhard Rothlin, Alex Harvill, David Adler, Mark Meyer, and Jan Novák. 2018. Denoising with kernel prediction and asymmetric loss functions. *ACM Trans. Graph.* 37, 4, Article 124 (jul 2018), 15 pages. doi:10.1145/3197517.3201388
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- Zhou Wang, Eero P. Simoncelli, and Alan C. Bovik. 2003. Multiscale Structural Similarity for Image Quality Assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers*, 2003, Vol. 2. 1398–1402. doi:10.1109/ACSSC.2003.1292216
- Mike Winkelman. 2019. Zero-Day, Open Research Content Archive (ORCA). <https://developer.nvidia.com/orca/beeples-zero-day>
- Songyin Wu, Deepak Vembar, Anton Sochenov, Selvakumar Panneer, Sungye Kim, Anton Kaplanyan, and Ling-Qi Yan. 2024. GFFE: G-buffer Free Frame Extrapolation for Low-latency Real-time Rendering. *ACM Trans. Graph.* 43, 6, Article 248 (Nov. 2024), 15 pages. doi:10.1145/3687923

- Lei Xiao, Salah Nouri, Matt Chapman, Alexander Fix, Douglas Lanman, and Anton Kaplanyan. 2020. Neural supersampling for real-time rendering. *ACM Trans. Graph.* 39, 4, Article 142 (aug 2020), 12 pages. doi:10.1145/3386569.3392376
- Ruifeng Xu and Sumanta N. Pattanaik. 2005. A novel Monte Carlo noise reduction operator. *IEEE Computer Graphics and Applications* 25, 2 (2005), 31–35. doi:10.1109/MCG.2005.31
- Tizian Zeltner, Fabrice Rousselle, Andrea Weidlich, Petrik Clarberg, Jan Novák, Benedikt Bitterli, Alex Evans, Tomáš Davidovič, Simon Kallweit, and Aaron Lefohn. 2024. Real-time Neural Appearance Models. *ACM Trans. Graph.* 43, 3, Article 33 (June 2024), 17 pages. doi:10.1145/3659577
- Boyuan Zhang and Hongliang Yuan. 2024. High-Quality Real-Time Rendering Using Subpixel Sampling Reconstruction. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 7 (mar 2024), 7006–7014. doi:10.1609/aaai.v38i7.28527
- Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 586–595. doi:10.1109/CVPR.2018.00068
- Xian Yao Zhang, Gerhard Röhlin, Shilin Zhu, Tunç Ozan Aydin, Farnood Salehi, Markus Gross, and Marios Papas. 2024. Neural Denoising for Deep-Z Monte Carlo Renderings. *Computer Graphics Forum* 43, 2 (2024), 18 pages. doi:10.1111/cgf.15050
- Zhihua Zhong, Jingsen Zhu, Yuxin Dai, Chuankun Zheng, Guanlin Chen, Yuchi Huo, Hujun Bao, and Rui Wang. 2023. FuseSR: Super Resolution for Real-time Rendering through Efficient Multi-resolution Fusion. In *SIGGRAPH Asia 2023 Conference Papers* (Sydney, NSW, Australia) (SA '23). Association for Computing Machinery, New York, NY, USA, Article 8, 10 pages. doi:10.1145/3610548.3618209

A Unbiasedness of stochastic rounding and the relaxed correction factor

For a fixed pixel i , write the requested density as $s_i = k_i + p_i$ with $k_i = \lfloor s_i \rfloor$ and $p_i \in [0, 1)$. Let the path-traced samples $S_{i,j}$ be i.i.d. with mean $\mu_i = \mathbb{E}[S_{i,j}] = L_i^{\text{ref}}$. The holdout contribution is $L_i^h = r_i [\mathbf{u}_i < p_i]$, where $\mathbf{u}_i \sim \mathcal{U}(0, 1)$ and r_i is distributed as an additional path sample.

$$\mathbb{E}[L_i^h] = \mathbb{E}[r_i] \mathbb{P}[\mathbf{u}_i < p_i] = \mu_i p_i. \quad (13)$$

Substituting into the noisy estimator gives

$$\mathbb{E}[L_i^{\text{noisy}}] = \frac{\sum_{j=1}^{k_i} \mathbb{E}[S_{i,j}] + \mathbb{E}[L_i^h]}{k_i + p_i} = \frac{k_i \mu_i + p_i \mu_i}{k_i + p_i} = \mu_i = L_i^{\text{ref}}. \quad (14)$$

For the relaxed estimator with $h = 1$, the unclipped ramp has support length p_i/λ and average height $(2\lambda - 1)/(2\lambda)$ over the region where a sample is evaluated. Its expected contribution is therefore

$$\mu_i p_i \frac{2\lambda - 1}{2\lambda}. \quad (15)$$

Choosing

$$h = \frac{2\lambda}{2\lambda - 1} \quad (16)$$

restores the expected holdout contribution to $\mu_i p_i$, which in turn preserves $\mathbb{E}[L_i^{\text{noisy}}] = L_i^{\text{ref}}$.

B Gathering pyramid definition

Our pyramidal filter builds on NPPD [Bálint et al. 2023] but replaces scatter-based operations with gather-based equivalents. With scatter kernels, the network must predict meaningful weights only at sampled pixels while most predictions go unused; with gather kernels, every output pixel predicts weights to collect from its neighborhood, distributing the task uniformly regardless of input sparsity.

We also simplify the pyramid by removing learned partitioning. NPPD predicts per-pixel weights that partition radiance across levels before downsampling. We instead use plain average pooling and rely on joint softmax normalization across kernel types to implicitly balance contributions.

Pyramid construction. We build a low-pass pyramid with $L = 5$ levels by average pooling:

$$L_{xyt}^l = \frac{1}{4^l} \sum_{u,v \in \mathcal{N}_{xy}^l} L_{uvt}^{\text{noisy}}, \quad (17)$$

where $\mathcal{N}_{xy}^l$ denotes the $2^l \times 2^l$ block of full-resolution pixels mapping to coarse pixel (x, y) at level l . Level $l = 0$ is full resolution; $l = L - 1$ is coarsest.

Kernel prediction. At each level, we predict spatially varying gather kernels:

- 5×5 denoising kernels ${}^d\kappa^l$ at all levels (25 weights),
- 2×2 upsampling kernels ${}^u\kappa^l$ at levels $l < L - 1$ (4 weights),
- 5×5 temporal kernels ${}^t\kappa$ at $l = 0$ only (25 weights).

We apply joint softmax normalization across all kernels at each level. At the coarsest level, only denoising weights are normalized. At intermediate levels:

$$\sum_{ij} d_{ij,xyt}^l + \sum_{ij} u_{ij,xyt}^l = 1. \quad (18)$$

At full resolution, the softmax spans all 54 weights:

$$\sum_{ij} d_{ij,xyt}^0 + \sum_{ij} u_{ij,xyt}^0 + \sum_{ij} t_{ij,xyt} = 1. \quad (19)$$

This joint normalization eliminates separate blending coefficients: the network learns to balance denoising, upsampling, and temporal contributions through relative kernel magnitudes.

Gather filtering. Each output pixel collects from its 5×5 neighborhood:

$$G(L^l, d_{ij,xyt}^l) = \sum_{i,j=-2}^2 d_{ij,xyt}^l \cdot L_{(x+i)(y+j)t}^l. \quad (20)$$

Learnable bilinear upsampling. Standard bilinear interpolation computes each output pixel as a fixed weighted average of four neighbors. We replace the fixed weights with learned per-pixel weights, allowing edge-aware upsampling. For upsampling from level $l+1$ to level l :

$$U(\hat{L}^{l+1}, u_{ij,xyt}^l) = \sum_{i,j=0}^1 u_{ij,xyt}^l \cdot \hat{L}_{\lfloor \frac{x+i}{2} \rfloor \lfloor \frac{y+j}{2} \rfloor t}^{l+1}. \quad (21)$$

The four kernel weights replace the position-dependent bilinear coefficients, enabling the network to sharpen edges or smooth across them as needed.

Coarse-to-fine reconstruction. Starting from the coarsest level:

$$\hat{L}_{xyt}^{L-1} = G(L^{L-1}, d_{ij,xyt}^{L-1}), \quad (22)$$

$$\hat{L}_{xyt}^l = G(L^l, d_{ij,xyt}^l) + U(\hat{L}^{l+1}, u_{ij,xyt}^l), \quad (23)$$

for $l = L-2, \dots, 1$. At full resolution, we add temporal filtering:

$$L_{xyt}^{\text{denoised}} = G(L^{\text{noisy}}, d_{ij,xyt}^0) + U(\hat{L}^1, u_{ij,xyt}^0) + G(\mathcal{W}_t L_{t-1}^{\text{denoised}}, t_{ij,xyt}), \quad (24)$$

where $\mathcal{W}_t$ warps the previous frame using motion vectors. The joint softmax ensures that in disoccluded regions or with unreliable motion, the network suppresses temporal contribution by reducing t_{ij} weights.

C Additional results

Table 10 reports the 0.25 spp comparison separately for each test scene. In addition to the spatial and perceptual metrics used in the main comparison, the final two columns report tPSNR and trMAE, computed from motion-compensated frame differences as described in Section 8.3. The per-scene results show that our method maintains strong performance across the full test set. Adaptive sampling also improves over the uniform variant for most metrics and scenes.

Figure 11 presents additional comparisons extending Fig. 10 to three further scenes. For each scene, we show the best- and worst-performing crops of our adaptive method relative to the uniform

Table 10. **Comparison with the top three baselines and our variants on each test scene.** 50% superresolution equivalent (0.25 spp sparse). We report both (SDR/HDR) CVVDP versions.

Method	PSNR↑	MS-SSIM↑	HaarPSI↑	MILO↑	CVVDP↑	CGVQM↑	tPSNR↑	trMAE↓
<i>Cyberpunk Market Soldier</i>								
DLSS	23.32	0.9143	0.6100	2.567	5.948/6.056	43.27	21.40	0.245
JNDS	21.23	0.8837	0.5179	2.282	4.971/5.223	21.23	19.28	0.343
NPPD	23.12	0.9143	0.5872	2.518	5.675/5.947	39.29	20.90	0.265
Ours Uniform	23.78	0.9280	0.6262	2.595	6.078/6.307	48.79	21.80	0.248
Ours Adaptive	24.80	0.9409	0.6603	2.665	6.480/6.674	54.09	22.72	0.236
<i>Cyberpunk Market Flythrough</i>								
DLSS	23.92	0.9321	0.6815	2.639	6.559/6.801	45.05	21.63	0.236
JNDS	21.10	0.8849	0.5580	2.270	5.521/5.657	15.77	18.95	0.329
NPPD	23.55	0.9233	0.6477	2.593	6.298/6.479	39.38	21.09	0.258
Ours Uniform	24.19	0.9337	0.6786	2.688	6.717/6.812	49.07	21.88	0.232
Ours Adaptive	24.93	0.9407	0.7055	2.737	6.964/7.079	52.91	22.54	0.222
<i>Bistro Exterior</i>								
DLSS	22.85	0.9219	0.6504	2.664	7.384/6.551	37.86	23.03	0.273
JNDS	22.65	0.8885	0.5658	2.393	6.568/6.595	21.07	20.79	0.358
NPPD	23.56	0.8977	0.5985	2.571	6.277/6.550	33.39	21.54	0.328
Ours Uniform	24.99	0.9248	0.6571	2.808	7.165/7.481	49.56	22.89	0.279
Ours Adaptive	25.12	0.9218	0.6633	2.804	7.192/7.414	49.00	23.15	0.269
<i>Bistro Day</i>								
DLSS	25.92	0.9407	0.6910	2.769	7.388/7.591	64.27	24.06	0.210
JNDS	23.99	0.9099	0.5941	2.481	6.568/6.774	52.42	21.84	0.274
NPPD	25.95	0.9452	0.6835	2.863	7.406/7.576	67.52	23.66	0.210
Ours Uniform	26.26	0.9262	0.6987	2.925	7.175/7.478	69.54	23.90	0.226
Ours Adaptive	27.77	0.9586	0.7553	3.130	7.999/8.175	75.11	25.33	0.190
<i>Bistro Interior</i>								
DLSS	24.37	0.9139	0.5806	2.366	6.219/6.691	46.75	22.84	0.246
JNDS	23.06	0.8887	0.5104	2.223	5.620/5.582	31.99	21.04	0.342
NPPD	24.82	0.9199	0.5793	2.385	6.361/6.302	48.41	22.80	0.262
Ours Uniform	25.06	0.9246	0.6115	2.417	6.548/6.460	53.22	23.26	0.251
Ours Adaptive	25.73	0.9361	0.6550	2.501	7.017/6.945	56.37	24.18	0.240
<i>ZeroDay Measure 7</i>								
DLSS	20.48	0.8460	0.5741	2.141	5.557/4.997	45.31	19.62	0.261
JNDS	20.11	0.8167	0.4905	1.963	5.363/3.997	32.11	17.93	0.336
NPPD	23.19	0.8750	0.5893	2.151	6.362/5.999	50.38	20.74	0.255
Ours Uniform	24.15	0.8954	0.6352	2.229	6.777/6.359	57.60	21.62	0.227
Ours Adaptive	24.73	0.9030	0.6533	2.268	7.050/6.678	59.86	22.15	0.221
<i>Village</i>								
DLSS	19.04	0.7412	0.5035	2.208	4.585/4.819	12.79	16.61	0.303
JNDS	17.84	0.7192	0.4734	2.079	3.828/3.937	-10.05	15.28	0.398
NPPD	19.45	0.8038	0.5241	2.285	4.835/4.915	21.32	16.70	0.326
Ours Uniform	20.31	0.8412	0.5727	2.460	5.433/5.449	34.24	17.62	0.294
Ours Adaptive	20.58	0.8416	0.5746	2.397	5.490/5.497	36.07	17.83	0.283
<i>Sponza</i>								
DLSS	24.09	0.9317	0.6183	2.612	6.500/6.624	48.08	22.42	0.219
JNDS	22.45	0.9126	0.5554	2.399	5.937/5.915	34.48	20.46	0.316
NPPD	24.65	0.9458	0.6351	2.797	6.813/6.770	57.70	22.42	0.237
Ours Uniform	25.31	0.9506	0.6710	2.939	7.083/7.055	64.57	23.21	0.218
Ours Adaptive	25.60	0.9516	0.6741	2.830	7.058/6.960	62.42	23.40	0.222
<i>Bedroom</i>								
DLSS	21.91	0.9136	0.6664	2.584	7.004/6.245	61.66	26.55	0.169
JNDS	24.83	0.8875	0.5675	2.390	6.478/5.891	60.78	24.91	0.246
NPPD	26.99	0.9309	0.6808	2.645	7.505/7.123	65.79	26.64	0.173
Ours Uniform	27.72	0.9405	0.7280	2.700	7.801/7.307	70.28	27.57	0.172
Ours Adaptive	29.39	0.9516	0.7922	2.764	8.248/7.331	72.88	28.90	0.164

baseline. These results demonstrate that the adaptive sampler consistently reallocates the sampling budget toward visually complex regions with fine details, leading to noticeable improvements in quality, while preserving high reconstruction quality in simpler regions.

In addition to these figures, we provide an interactive image viewer in the supplement to facilitate detailed visual comparisons between methods, as well as a video that further illustrates temporal behavior and perceived quality.




Fig. 11. **Best- and worst-case robustness analysis.** This figure presents additional results analogous to Figure 10 on different scenes. It highlights the best- and worst-performing regions of our adaptive method compared to the uniform baseline. The plot on the right shows the mean absolute error as a function of the relative sampling rate: regions undersampled by our method appear on the left, while oversampled regions appear on the right.