

Environment Maps Editing using Inverse Rendering and Adversarial Implicit Functions

Antonio D’Orazio¹ , Davide Sforza¹ , Fabio Pellacini² , Iacopo Masi¹ 

¹Sapienza, University of Rome, Italy

²University of Modena and Reggio Emilia, Italy




Figure 1: We use an adversarial neural environment map combined with an inverse rendering pipeline to generate novel environment maps from a single target rendering. The target can be just edited with a stroke in 2D using image editing software. The figure shows shadow manipulation and adding reflections. The result column displays the rendered results along with the generated environment map.

Abstract

Editing High Dynamic Range (HDR) environment maps using an inverse differentiable rendering architecture is a complex inverse problem due to the sparsity of relevant pixels and the challenges in balancing light sources and background. The pixels illuminating the objects are a small fraction of the total image, leading to noise and convergence issues when the optimization directly involves pixel values. HDR images, with pixel values beyond the typical Standard Dynamic Range (SDR), pose additional challenges. Higher learning rates corrupt the background during optimization, while lower learning rates fail to manipulate light sources. Our work introduces a novel method for editing HDR environment maps using a differentiable rendering, addressing sparsity and variance between values. Instead of introducing strong priors that extract the relevant HDR pixels and separate the light sources, or using tricks such as optimizing the HDR image in the log space, we propose to model the optimized environment map with a new variant of implicit neural representations able to handle HDR images. The neural representation is trained with adversarial perturbations over the weights to ensure smooth changes in the output when it receives gradients from the inverse rendering. In this way, we obtain novel and cheap environment maps without relying on latent spaces of expensive generative models, maintaining the original visual consistency. Experimental results demonstrate the method’s effectiveness in reconstructing the desired lighting effects while preserving the fidelity of the map and reflections on objects in the scene. Our approach can pave the way to interesting tasks, such as estimating a new environment map given a rendering with novel light sources, maintaining the initial perceptual features, and enabling brush stroke-based editing of existing environment maps. Our code is publicly available at github.com/OmnAI-Lab/R-SIREN.

CCS Concepts

• **Computing methodologies** → Artificial intelligence; Computer graphics; Image manipulation;

1. Introduction

High-quality environment maps are essential for achieving photo-realistic renderings. However, they are often treated as static assets, limiting the ability of artists to customize the environment. Instead of editing existing maps, artists typically choose from a pre-curated collection. In light of this, our work aims to present an approach to perform the editing of environment maps by implementing a differentiable rendering in the pipeline, allowing back-propagation from the final rendered image to the initial objects in the 3D scene. This approach can be seen as an inverse problem, which can be solved via optimization techniques such as gradient descent to iteratively update the color values of the environment maps. Instead of optimizing the map directly in the pixel space, which is highly fragile, or resorting to expensive generative models based on diffusion process [LTH*23b], we take a lazy learning approach and parameterize each environment map using a robust implicit function implemented using a standard multi-layer perceptron (MLP). To attain naturally looking images, we train our robust function so that for small perturbations of the parameters—the gradients received by the inverse rendering—the output of the new map remains smooth and stable. Fig. 2 shows an overview of our approach: a target image $\mathbf{I}_{\text{target}}$ is rendered from an *unknown* environment map. We fit an implicit function $f_{\theta}(\mathbf{x})$ to an initial environment map \mathbf{E} . The method renders the \mathbf{E} and updates its representation to match the target image. Instead of using plenty of regularizers to constrain the problem, we make our implicit function smooth and robust so that once we update its weights, the new map looks natural and has fewer artifacts. When designing our implicit representation, we employed SIREN [SMB*20], which has been proven to reconstruct images with very high fidelity, and adapted its training strategies to correctly represent HDR images, which was not possible in the original approach. We call this method HDR SIREN. Fig. 3 shows how it is integrated into the inversion algorithm. Our method is designed to allow computer graphics artists to fine-tune the lighting of 3D scenes by directly manipulating the environment maps, which traditionally has been a complex and time-consuming process. The key objective is to enable precise adjustments to light sources within these maps without corrupting the background, maintaining fidelity in reflections and lighting. Our method can be applied to downstream tasks, as shown in Fig. 1, in which we alter the appearance of the environment map by specifying coarse sketches over the initial rendering.

During our work, we faced the following challenges: \diamond *Sparsity of Relevant Pixels*: The optimization of environment maps is hampered by the sparsity of relevant pixels. In most cases, only a small fraction of the image directly contributes to illuminating the scene, making it difficult to achieve consistent and meaningful updates. This sparsity leads to noise and convergence issues, especially when dealing with HDR images; \diamond *Convergence Difficulties with HDR Images*: High Dynamic Range (HDR) images present significant challenges during optimization. The wide range of pixel intensities in HDR images, which can span from zero to infinity, complicates the optimization process. High learning rates may cause artifacts in the background, while low learning rates fail to effectively manipulate the light sources, leading to suboptimal results.; \diamond *Balancing Background and Light Source Manipulation*: A key challenge lies in balancing the optimization between maintaining

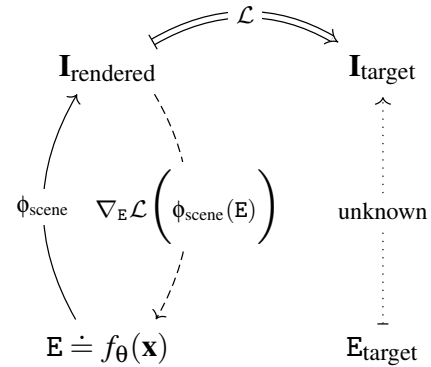


Figure 2: We start from an initial environment map \mathbf{E} that we parameterize using an implicit function via a robust neural network $f_{\theta}(\mathbf{x})$. We optimize the new environment map by receiving gradients through an inverse differentiable rendering that matches the rendered image $\mathbf{I}_{\text{rendered}}$ to a target $\mathbf{I}_{\text{target}}$. The target is generated by an unknown environment map $\mathbf{E}_{\text{target}}$; in practical settings, $\mathbf{I}_{\text{target}}$ could be created by roughly painting possible shades from $\mathbf{I}_{\text{rendered}}$ so to “find” the generating environment map.

the integrity of the background and manipulating the light sources. Techniques such as separating the light sources from the background and applying regularization methods are necessary to prevent the background from being corrupted while still allowing effective light source modification; \diamond *Noise Reduction*: The optimization process often introduces noise into the environment maps, particularly when using plain optimization with no regularization. This noise must be mitigated through regularization techniques like Total Variation regularization, which penalizes rapid changes in pixel values, ensuring a smoother and more coherent output; \diamond *Maintaining Luminance Consistency*: During the optimization process, it is crucial to maintain the overall luminance of the environment map to preserve the correct balance between light sources and shadows. This requires careful management of luminance regularization to ensure that the optimized environment map does not deviate significantly from the original in terms of brightness and contrast.

Given the challenges mentioned above, our work aims to provide the following contributions: \diamond *Inverse rendering-based HDR optimization pipeline*: We propose a lazy-learning optimization pipeline that reconstructs a desired environment map given a target rendering. In particular, we explore how to constrain the sparsity and the high variance in values of HDR images to correctly converge to a new HDR image, maintaining the overall consistency of the sky environment. Such exploration involves regularization techniques, e.g. defining a penalty on the total luminance, and image representations that help the Adam optimizer to reach convergence. We will focus on two image representations: 1) pixel-based representations, such as optimizing the environment map in the logarithmic space; 2) implicit representations that map the (x, y) image grid to colors using a neural network, representing HDR images. Our representation is robust to gradient updates received from the inverse rendering, thereby providing smooth editing of environment maps; \diamond *Improvements over the SIREN implicit representation*: The baseline for the SIREN implicit representation does not converge

with HDR images. In this work, we show how to correctly train a neural network based on periodical activation functions such that it can encode high spikes in the image. Finally, we show how to enforce robustness in the weights space to ensure that the result of the inverse rendering-based optimization pipeline will produce natural-looking images.

The next section reviews the related research. Section 3 explains the image representations used in our work, divided between traditional RGB pixel representation and implicit representations. Section 4 describes our proposed robust implicit representation and the process to fit it to an HDR map, while Section 5 describes the inverse rendering process and the optimization pipeline and how it tackles the challenges addressed earlier. Finally, Section 6 covers the training and implementation details, presents the results, and includes a discussion of those results.

2. Related Work

Appearance Editing. *EnvyLight* [Pel10] uses an algorithm to select environment map regions by sketching strokes on lighting features in the rendered image, separating the background and the foreground features in the environment map. The paper proposes the use of wavelets to capture all frequency effects. The separation between background and foreground is proposed again in this work. Instead of being specified manually by the user, the separation is performed automatically during the optimization process. *Illumination Brush: Interactive Design of All-frequency Lighting* [OMSI07] is another traditional approach for environment maps editing for which the authors designed a 3D viewport where the user can paint over the surface of a 3D model desired lighting. Then, it estimates the environment map, creating the lighting style as an inverse problem. The final environment map, however, does not resemble a captured photograph. *Appearance editing with free-viewpoint neural rendering* [GKS*21] proposes a method to estimate a BRDF using a differentiable renderer and a few reference images at different angles, enabling material and shading edits on 3D models, showing the potential of differentiable rendering for appearance editing.

Inverse Rendering and Model Inversion. Differentiable rendering is a broad field, where the work can involve different parts of the pipeline. Mainly, the work can be grouped as solving optimization problems and working on the differentiable renderer itself. *Large Steps in Inverse Rendering* [NJJ21] is a breakthrough method that dramatically improves the reconstruction of 3D meshes from a target image, giving many perspectives on how to tune the optimization function. The idea of using the Laplacian regularizer, which is translated to the Total Variation (TV) regularizer for simpler domains like 2D images, helped with noise and convergence. Developed by the same research group, *Mitsuba 3* [JSRV22] is the differentiable rendered that they used to implement their work, built on top of the just-in-time compiler called Dr.Jit. It is a retargetable renderer, meaning that the data structures can be adapted to perform different tasks. For example, the autograd mode allows the renderer to keep track of the gradients with respect to the input parameters, such as textures, geometries, and material parameters. Given that the environment map is an image texture, the native support for keeping gradients for textures was the reason for choosing Mitsuba over other differentiable renderers. Another example of a differentiable renderer is represented by the *NVIDIA*

Kaolin [FTSL*22] suite, which is a PyTorch library for simplifying the work with 3D assets. However, their differentiable renderer is more suitable for optimizing 3D shapes and light positions. Therefore, Mitsuba 3 resulted in a better choice for our work. Finally, the work in [WLS*21, RMKM23] are one-shot methods for image synthesis that use structured gradients received from the inversion of a pre-trained classifier to control manipulations; [RMKM23] moreover employs a robust model to receive perceptually aligned gradients (PAG) and allows the user to edit the image using a binary mask. Finally, [LTH*23a] proposes a similar pipeline compared to our approach, but, instead, they involve diffusion posterior sampling as a prior for the inverse rendering process, achieving remarkable results but requiring an expensive pre-training of a Denoising Diffusion Probabilistic Model (DDPM) [HJA20] over multiple environment map samples. Finally, *Photorealistic rendering of mixed reality scenes* [KBC*15] provide an extensive survey of differentiable rendering approaches involving HDR illumination.

Adversarial Robustness. Neural network robustness has increasingly become a central topic in deep learning. The work shown in AT [MMS*18], which incorporates adversarial examples into the training process, is still the most effective empirical strategy, with many modifications developed on top of it. Zhang [ZYJ*19] proposed the TRADES loss, which leverages the Kullback-Leibler (KL) divergence to balance the trade-off between standard and robust accuracy. Cui [CTZ*23] improved the KL divergence, addressing its asymmetry property with improved results. Finally, Adversarial Weight Perturbation (AWP) [WXW20] shows how to perturb the training weights, helping the generalization by flattening the loss landscape. Unlike classic adversarial perturbation in the input space, AWP is scaled proportionally to the relative size of the original weights for each layer. Finally, robust classifiers are mainly known for their discriminative power, nevertheless prior art [MBBM24] also linked them to generative models.

3. Image Representations

HDR RGB. HDR imaging is used to accurately represent scenes with a wide range of light intensities. Unlike traditional 8-bit RGB, HDR RGB encodes luminance using higher bit depths (e.g., 16-bit or 32-bit), capturing more detail in both bright and dark areas. This representation is crucial for the environment maps, as it preserves the full spectrum of light in a scene, allowing the representation of the full intensity of the light sources. HDR images are theoretically coded in the range:

$$\mathbf{I}_{\text{hdr}}(\mathbf{x}) \in [0, +\infty) \quad \mathbf{x} = (x, y) \in \mathbb{R}^2$$

Where the practical upper limit is determined by the specific HDR implementation. SDR images are represented in the range:

$$\mathbf{I}_{\text{sdr}}(\mathbf{x}) \in [0, 1], \quad \mathbf{x} = (x, y) \in \mathbb{R}^2$$

This image representation is the most trivial one, but it requires additional image processing and regularization to constrain the inverse optimization problem. Also, the large range in values of HDR images is not well-suited for gradient descent optimization. Higher learning rates will optimize the few light sources correctly at the cost of ruining the remaining background. Lower learning rates will optimize only the background pixels, keeping the light sources substantially unchanged.

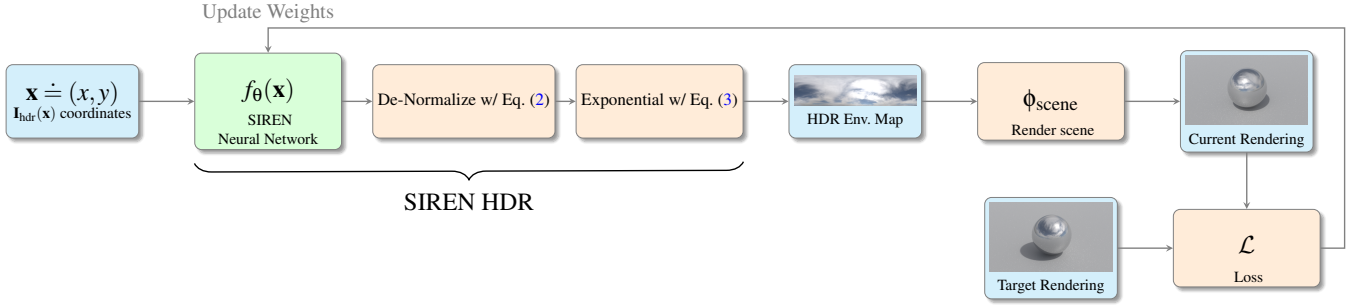


Figure 3: Inverse rendering in the proposed SIREN HDR

SIREN Implicit Representation. Optimizing directly on the environment map pixel will lead to low-quality images due to the sparsity nature of the problem, producing new images that are off the manifold of the natural images. To represent the environment map, we first train a SIREN (Sinusoidal Representation Networks) [SMB*20] model to implicitly represent the image. The SIREN model leverages periodic activation functions to capture the high-frequency details in the image. Specifically, the image is represented as a continuous function $f_{\theta} : \mathbb{R}^2 \rightarrow \mathbb{R}^3$, where θ denotes the network parameters. Given a spatial coordinate $\mathbf{x} = (x, y)$, the network outputs the corresponding RGB color values $f_{\theta}(\mathbf{x}) = (r, g, b)$. The sine activation function is defined as $\sigma(\mathbf{x}) = \sin(\omega_0 \mathbf{x})$, where ω_0 is a frequency parameter controlling the period of the sine function. The first layer processes the input spatial coordinates as follows $\mathbf{h}_1 = \sin(\omega_0 \mathbf{W}_0 \mathbf{x} + \mathbf{b}_0)$. Then, each layer i -th computes:

$$\mathbf{h}_{i+1} = \sin(\mathbf{W}_i \mathbf{h}_i + \mathbf{b}_i)$$

where \mathbf{W}_i and \mathbf{b}_i are the weight matrix and bias vector of the i -th fully connected layer. All the sets of parameters are thus $\theta \doteq \{\mathbf{W}_i, \mathbf{b}_i\}$ and the network is a multi-layer perceptron (MLP). This configuration allows the SIREN model to map spatial coordinates directly to pixel values, although additional care is needed to correctly model HDR images.

4. Representing HDR Images with SIREN

The standard definition of SIREN does not model HDR images directly, as the sinusoidal activation functions can reconstruct only the standard color range before losing its learning capabilities. To address this, we train SIREN to represent the environment map in the logarithmic space. Given a target HDR environment map \mathbf{E} , we train SIREN to represent the normalization in [XPH05]:

$$f_{\theta}(\mathbf{x}) \approx \frac{\log(\mathbf{E}) - \min(\log(\mathbf{E}))}{\max(\log(\mathbf{E})) - \min(\log(\mathbf{E}))}. \quad (1)$$

Then, to obtain the learned environment map \mathbf{E}_{pred} , we de-normalize and transform it back to the HDR space as:

$$\hat{f}_{\theta}(\mathbf{x}) = f_{\theta}(\mathbf{x}) (\max(\log(\mathbf{E})) - \min(\log(\mathbf{E}))) + \min(\log(\mathbf{E})) \quad (2)$$

$$\mathbf{E}_{\text{pred}} = \exp(\hat{f}_{\theta}(\mathbf{x})) \quad (3)$$

Note that, to prevent numerical instability, we add a small value $\epsilon = 0.01$ to the pixels inside the logarithm, which is omitted for readability.

To capture the few high peaks of an HDR image, and thus address the problem of sparsity, we enforce this behavior in the loss function. We define two loss functions as:

$$\mathcal{L} = \underbrace{\lambda_{\log_mse} \|f_{\theta}(\mathbf{x}) - \log(\mathbf{E})\| + \lambda_{\log_ssim} \text{SSIM}(f_{\theta}(\mathbf{x}), \log(\mathbf{E}))}_{\mathcal{L}_{\log}} + \underbrace{\lambda_{hdr} \|\hat{f}_{\theta}(\mathbf{x}) - \mathbf{E}\|}_{\mathcal{L}_{hdr}} \quad (4)$$

The loss \mathcal{L}_{\log} in the log space will focus more on reconstructing the overall features of the environment map. In addition, by defining a loss function in the re-transposed HDR space, e.g. \mathcal{L}_{hdr} , the optimizer can give more importance to the few pixels with high peaks, therefore addressing the sparsity issue. Fig. 4 shows how such strategies result in better coverage of the whole HDR space.

Enabling Smooth Editing via Adversarial Perturbation. We propose to interpret the weights θ of neural implicit representations


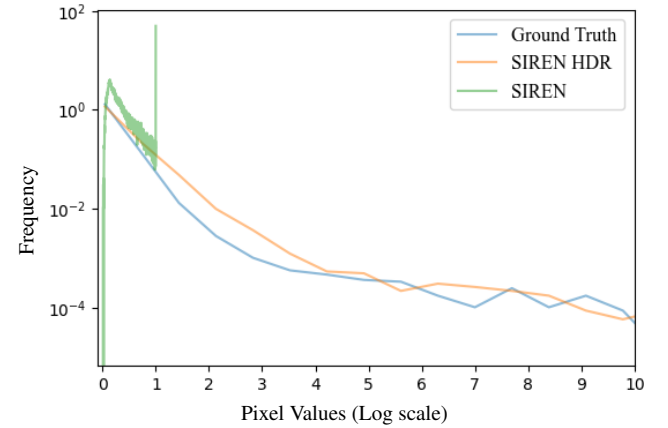


Figure 4: SIREN HDR can represent HDR images with higher fidelity and better pixel intensity distribution, compared to the baseline architecture in [SMB*20].

as a form of “embedding” for the target image. With this interpretation in mind, each specific set of weights corresponds to a unique representation of the data. This raises the question of whether we can adjust these weights such that perturbations do not disrupt the underlying image space. If we do so we could navigate “around the embedding” to synthesize new natural-looking images. Conceptually, this can be viewed as exploring whether the weight set belongs to a smooth implicit latent space, which is not directly observable. Adversarial training *in the weight space* [WXW20] enables us to craft this smooth implicit latent space. We thus seek for a set of weights $f_{\theta}(x)$ such that for small perturbation in the weights space, the implicit representation remains stable and outputs natural environment maps. Once we have fit the MLP and run the optimization, our implicit representation is robust to gradient updates we receive from the inverse rendering.

In practical terms, our approach involves enforcing robustness by incorporating adversarial perturbations during the training process. Specifically, if we consider a specific layer \mathbf{W} , we seek to optimize it to minimize the worst-case loss across potential perturbations $\Delta\mathbf{W} \in \Omega$. This is formalized as the following optimization problem:

$$\min_{\mathbf{W}} \max_{\Delta\mathbf{W} \in \Omega} \rho(\mathbf{W} + \Delta\mathbf{W}) \rightarrow \min_{\mathbf{W}} \max_{\Delta\mathbf{W} \in \Omega} \frac{1}{n} \sum_{i=1}^n \mathcal{L}(f_{\mathbf{W}+\Delta\mathbf{W}}(\mathbf{x}_i, \mathbf{I}_{\text{hdr}}(\mathbf{x}_i))) \quad (5)$$

In this expression, $\rho(\mathbf{W} + \Delta\mathbf{W})$ represents the robust loss function, $f_{\mathbf{W}+\Delta\mathbf{W}}$ denotes the model with perturbed weights, \mathcal{L} is defined in Eq. (4), \mathbf{x}_i and $\mathbf{I}_{\text{hdr}}(\mathbf{x}_i)$ are the input grid and the corresponding target image. Unlike the classic adversarial examples in the input space where each input is perturbed in a fixed neighborhood, in our case, the feasible set of the perturbation Ω is restricted using each relative weight size. Thus $\Omega = \{\forall \Delta : \|\Delta\| \leq \gamma \|\mathbf{W}\|\}$ and γ is a fixed constant used to bound the weight perturbation. Note that, unlike [WXW20], we *just* perturb the weights space and not the input space. Algorithm 1 illustrates the practical approach to perform adversarial weight perturbation, while its role during the training can be seen in Algorithm 2. We will refer to SIREN neural networks trained with the explained robustness strategy as R-SIREN. The effects of training implicit representations with robustness can be seen in Fig. 5.

Algorithm 1: Computing the Adversarial Weight Perturbation (calc_awp)

Input: Set of weights θ , input coordinates \mathbf{x}

Output: Perturbation θ^*

```

1  $\hat{\theta} \leftarrow \text{clone } \theta$ 
2  $Y \leftarrow f_{\hat{\theta}}(\mathbf{x});$ 
3  $\mathcal{L}_{\text{awp}} \leftarrow -\mathcal{L}_{\text{SSIM}}(Y, \mathbf{E})$ 
4 update weights  $\hat{\theta}$ 
5  $\Delta\theta \leftarrow \theta - \hat{\theta}$ 
6  $\theta^* \leftarrow \frac{\|\theta\|}{\|\Delta\theta\| + \epsilon} \cdot \Delta\theta$ 

```

5. Inverse Rendering Optimization

Given a target image $\mathbf{I}_{\text{target}}$, we want to minimize the loss function between the rendered image $\phi_{\text{scene}}(\mathbf{E})$ and $\mathbf{I}_{\text{target}}$, by iteratively

Algorithm 2: Training SIREN with Adversarial Weights Perturbation

Input: Target environment map \mathbf{E} , model f_{θ} , input

coordinates \mathbf{x} , learning rate η , number of iterations N

Output: Trained model f_{θ}

```

1 for  $i = 1$  to  $N$  do
2    $\theta^* \leftarrow \text{calc\_awp}(\theta_{i-1})$ 
3    $\theta_{\text{adv}} \leftarrow \theta_{i-1} + \theta^*$ 
4    $Y \leftarrow f_{\theta}(\mathbf{x});$ 
5    $\mathcal{L}_{\text{robust}} \leftarrow \mathcal{L}_{\text{log}}(Y, \mathbf{E}) + \mathcal{L}_{\text{hdr}}(Y, \mathbf{E})$ 
6   update weights  $\theta_{\text{adv}}$ 
7    $\theta_i \leftarrow \theta_{\text{adv}} - \theta^*$ 
8 end

```

altering the environment map \mathbf{E} as:

$$\min_{\mathbf{E}} \mathcal{L}(\mathbf{I}_{\text{rendered}}, \mathbf{I}_{\text{target}}) + \mathcal{L}_{\text{reg}}(\mathbf{E}) = \mathcal{L}(\phi_{\text{scene}}(\mathbf{E}), \mathbf{I}_{\text{target}}) + \mathcal{L}_{\text{reg}}(\mathbf{E}) \quad (6)$$

The ill-posed nature of inverse rendering problems often leads to trivial solutions that minimize error without preserving meaningful image features. Unconstrained optimization can exploit these ambiguities, for instance, by manipulating only reflective regions or darkening the brightest pixels. This can result in a loss of smooth transitions, a property found in natural images. To address this, we employ regularization techniques to constrain intensity, luminance, and perceptual characteristics, ensuring that the synthesized image is consistent with the initial sample.

5.1. Regularization

Luminance Regularization. During the optimization process, we want to keep the overall luminance of the environment map constant to maintain the correct balance between light sources and shadows. To achieve this result, we propose a luminance regularization that aims to balance the difference in luminance between the original environment map and the optimized one. We can compute the luminance as:

$$\text{Lum}(\mathbf{E}) = \sum_{i,j} \sum_{k \in \{\text{r,g,b}\}} \alpha_k \mathbf{E}_{i,j}[k]. \quad (7)$$

where $\alpha_r = 0.2126$, $\alpha_g = 0.7152$, $\alpha_b = 0.0722$, per definition of luminance and $\mathbf{E}_{i,j}[k]$ represents the k -th color plane of the environment map at pixel i, j . Subsequently, the proposed luminance regularization is:

$$\mathcal{L}_{\text{lum}} = \lambda_{\text{lum}} |\text{Lum}(\mathbf{E}) - \text{Lum}(\mathbf{E}')| \quad (8)$$

Where \mathbf{E}' is the initial environment map, and λ_{lum} is the trade-off parameter. During the rendering process, usually, only the upper part of the image is used because the lower part is covered by the scene elements. When using regularization techniques such as luminance regularization, some scene configurations, such as the presence of the infinite plane, could cause the optimizer to cheat by simply lowering the brightness of the lower part of the environment map, keeping the amount of luminance constant but burning the usable section. To avoid this from happening, the optimization is run only on the higher part.



Figure 5: Effects of applying the *same* small random perturbation to implicit functions. The random perturbation is $\mathbf{W}^* \sim \alpha \mathcal{N}(0, 1)$ with α set to 1×10^{-3} and is applied in the weights space. From left to right: the original environment map, the rendering of SIREN HDR, the rendering of R-SIREN HDR, the perturbation applied to SIREN HDR, and, finally, the perturbation applied to R-SIREN HDR. In this last case, perturbation of the weights induces naturally looking maps.

Extracting Lights Information. Once the global luminance is constrained, the optimizer can cheat by lighting up a large amount of pixels to avoid penalties from the luminance regularization. However, environment maps usually have a small set of pixels that are very large in value, surrounded by values in the standard range. For example, in the presence of the sun. Finding the pixels belonging to a light source in an entire environment map is a *sparse* problem. For instance, in a traditional 1K environment map, we have a total of $1024 \times 512 = 524,288$ pixels. On the other hand, the light sources, like the sun, usually represent only a few dozen of the whole image. Therefore, we want to encourage sparsity by applying the ℓ_1 regularization over the environment map values. Such regularization, when competing against \mathcal{L}_{lum} , results in environment maps that are faithful to the original color intensity distribution.

$$\mathcal{L}_{L1}(\mathbf{E}) = \lambda_{L1} \sum_{i,j} \sum_{k \in \{r,g,b\}} |\mathbf{E}_{i,j}[k]|. \quad (9)$$

Preserving the Perceptual Features. To ensure that the synthesized environment maps retain the global perceptual qualities of the original, we utilized the Deep Image Structure and Texture Similarity (DISTS) metric as an additional regularization term $\mathcal{L}_{dists} = \lambda_{dists} DISTS(\mathbf{E}, \mathbf{E}_0)$. The DISTS metric [DMWS20] evaluates similarity by focusing on the structural and textural aspects of images, leveraging deep learning to compare these perceptual elements rather than just pixel-wise differences. When used as a regularization term, it guides the optimization towards more natural-looking images.

6. Experiments

We run the inverse rendering optimization described at Eq. (6) with different SIREN training strategies. All the experiments ran on a single RTX 4090 GPU. The environment maps are taken from PolyHaven [pol], resized to 256×128 to speed up the computation time, and cropped to the upper part, which is the only one contributing to the test scenes with the infinite plane. The environment maps are stored in the OpenEXR format, using 32-bit floating-point numbers. The scenes are rendered using the Mitsuba 3 [JSRV22] differentiable renderer.

6.1. R-SIREN HDR Implementation Details

We trained the SIREN models for 1500 epochs, using the Adam optimizer with $lr = 5 \times 10^{-5}$, averaging 30 seconds of training time per image. For the Adversarial Weight Perturbation, we employed a proxy Adam optimizer of 1×10^{-4} , applying perturbations of size




Figure 6: Results over many SIREN training strategies and experiment setups. The initial rendering is obtained with the initial environment map. The goal is to run the optimization to reach the target’s lighting conditions. Note that the target environment map is unknown during the optimization.

$\gamma = 0.01$. Both learning rates are scheduled with a Cosine Annealing scheduler with Warm Restarts, where the number of iterations before the first restart is $T_0 = 100$. We do not enforce a lat-long prior at this step, since it is automatically handled by the inverse rendering process.

For the model architecture, compared to stock SIREN, which terminates the neural network with a linear layer, we appended a final Sigmoid activation function as the last layer to strictly enforce the

results inside the $(0, 1)$ interval to properly undo the normalization, according to Eq. (1). The model is then instantiated with 2 input features, 3 output features, 256 hidden features, and 6 hidden layers. The sinusoidal activation layers have the initial angular frequency set to $\omega_0 = 30$. Our choice of hyperparameters for the loss terms is $\lambda_{\log_mse} = 0.85$, $\lambda_{\log_ssim} = 0.25$, and $\lambda_{hdr} = 1 \times 10^{-2}$.

6.2. Inverse Rendering

We get the HDR prediction from SIREN with Eq. (3). We then use the SSIM loss to enforce the similarity in the shadows, regularized with luminance and L1, according to Eq. (8) and Eq. (9), and the DISTS metric.

$$\mathcal{L} = \lambda_{ssim} \mathcal{L}_{SSIM}(I_{new}, I_{target}) + \underbrace{\lambda_{lum} \mathcal{L}_{lum}(E_{pred}, E_0) + \lambda_{dists} \mathcal{L}_{dists}(E_{pred}, E_0) + \lambda_{L1} \mathcal{L}_{L1}(E_{pred})}_{\mathcal{L}_{reg}} \quad (10)$$

With $\lambda_{ssim} \in [1, 1 \times 10^4]$, $\lambda_{L1} \in [1 \times 10^{-4}, 1 \times 10^{-8}]$, $\lambda_{lum} \in [1 \times 10^{-4}, 1 \times 10^{-8}]$, and $\mathcal{L}_{dists} = 0.6$. Initially, we employ higher λ_{L1} and λ_{lum} to escape the local minima and reduce the original light intensity. Subsequently, we gradually decrease the regularization weights for L1 and luminance, while simultaneously increasing the emphasis on SSIM, to ensure accurate reconstruction of the shadows. We keep the DISTS hyperparameter weight constant to preserve the original perceptual features at each step. We used the Adam optimizer, with $lr = 5 \times 10^{-6}$, running for 400 iterations.

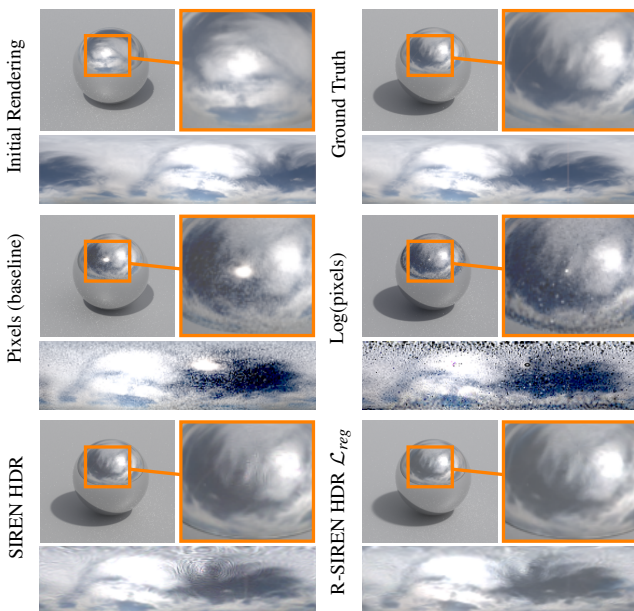


Figure 7: Accuracy in the reproduction of the reflections for a full mirror metallic sphere with roughness set to zero. R-SIREN obtains natural-looking environment maps with fewer artifacts.

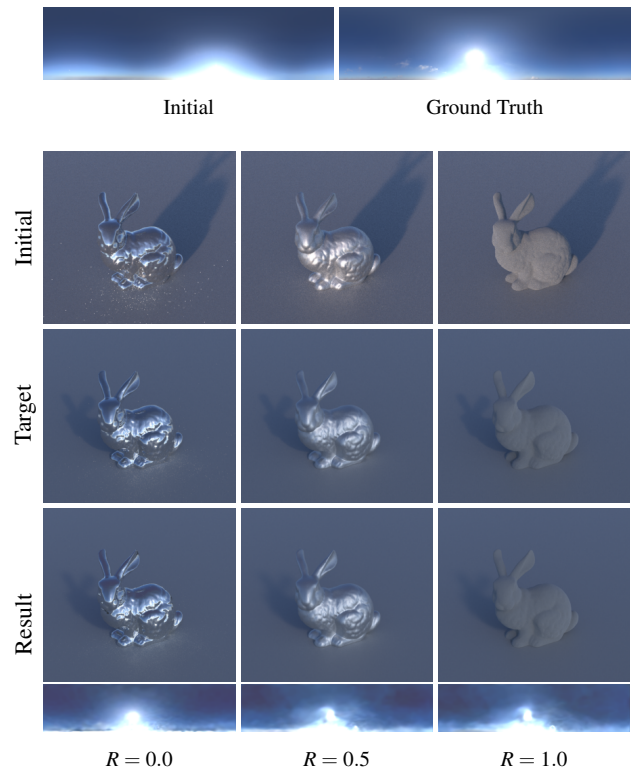


Figure 8: Inverse rendering with roughness R set to 0, 0.5, 1, using R-SIREN HDR \mathcal{L}_{reg} . The ground truth is a different environment map in a different light orientation. Despite the roughness to 1 impairs the inverse rendering process, our method is still able to recover an environment map close to the ground truth.

6.3. Evaluation

Qualitative. To evaluate the representation capabilities of R-SIREN HDR, we started with a physically correct target to avoid biases from the renderer. The first experiment involved reconstructing a new environment map starting from a rendered target in different lighting conditions. In this case, the initial environment map shifted to the left. Fig. 6 shows the results of the preliminary test.

To assess the reflection quality in our sphere scene, we employed a full mirror as the reflective surface. The experiment is conducted with the SIREN HDR and R-SIREN HDR models to evaluate their quality against the baseline optimization, with Adam directly altering the pixel values. As we can see in Fig. 7, both models outperform the baseline optimization in the pixel space, including when the optimizer is altering the values in the logarithmic space. Furthermore, R-SIREN HDR provides finer details in the reconstruction quality.

We further investigated the impact of surface roughness on the inversion process. By varying the roughness values of the reflective surface, we observed that surfaces with minor roughness generally yielded the best reconstructions, as shown in Fig. 8. This is attributed to the fact that minor roughness facilitates the inverse rendering process. The smoother surface provides more predictable

Table 1: Evaluation of rendering and environment map against target and ground truth, for the experiment presented in Fig. 6.

Experiment	Envmap Metrics				Rendering Metrics			
	↑ PSNR	↓ LPIPS	↓ MSE	↑ SSIM	↑ PSNR	↓ LPIPS	↓ MSE	↑ SSIM
Pixels	8.91931 dB	0.63838	0.12825	0.31881	29.91687 dB	0.17168	0.00102	0.94152
Log(Pixels)	8.32278 dB	0.76582	0.14714	0.02603	46.27205 dB	0.06250	0.00002	0.99040
SIREN HDR w/o \mathcal{L}_{hdr}	14.89818 dB	0.60814	0.03237	0.47149	38.95401 dB	0.08466	0.00013	0.98292
SIREN HDR	14.56471 dB	0.62082	0.03496	0.34456	46.37604 dB	0.05282	0.00002	0.99316
R-SIREN HDR	14.58689 dB	0.55724	0.03478	0.55372	39.95935 dB	0.05812	0.00010	0.99125
R-SIREN HDR \mathcal{L}_{reg}	14.85838 dB	0.53229	0.03267	0.52657	41.23402 dB	0.09429	0.00008	0.98908

Table 2: Evaluation of rendering and environment map against target and ground truth, for the experiment presented in Fig. 7.

Experiment	Envmap Metrics				Rendering Metrics			
	↑ PSNR	↓ LPIPS	↓ MSE	↑ SSIM	↑ PSNR	↓ LPIPS	↓ MSE	↑ SSIM
Pixels	18.35015 dB	0.48241	0.01462	0.54777	26.01747 dB	0.25495	0.00250	0.88970
Log(Pixels)	14.60996 dB	0.63089	0.03459	0.28612	30.00922 dB	0.05187	0.00100	0.97107
SIREN HDR \mathcal{L}_{reg}	29.19136 dB	0.32992	0.00120	0.85988	40.78696 dB	0.09896	0.00008	0.98072
R-SIREN HDR \mathcal{L}_{reg}	29.77839 dB	0.29141	0.00105	0.91918	35.46278 dB	0.17063	0.00028	0.95863

and consistent reflections, making it easier for the inverse rendering process to discover the surrounding environment texture.

Quantitative. To provide a more objective assessment of R-SIREN HDR’s performance, we employed quantitative metrics such as Peak Signal-to-Noise Ratio (PSNR), Learned Perceptual Image Patch Similarity (LPIPS) [ZIE*18], Mean Squared Error (MSE), and Structural Similarity Index (SSIM) [WBSS04]. These metrics measure noise ratio, perceptual similarity, and reconstruction similarity. We measured the quality of both rendering and environment maps against the ground truth for the Mirror Sphere and the Horse experiments, with Table 1 and Table 2, showing the computed scores. The results display a significant gap in the environment map metrics in favor of the SIREN-based approaches. The discrepancy in the results between the environment map and the rendering metrics proves that the unconstrained optimizations lead to overfitting the target scene without focusing on the overall image quality, resulting in final rendering that is very close to the target but degrading the visual features of the environment map. In contrast, LPIPS is a perceptual metric that prioritizes high-level visual features and global similarity, making it less sensitive to local noise and pixel-level discrepancies. Our approach with SIREN is not designed to reconstruct the unknown ground truth by overfitting it, but rather to synthesize a new one while preserving the aesthetic similarity with the original environment map, minimizing noise, and maintaining a continuous structure. This objective is rewarded by the LPIPS, as it shows that R-SIREN HDR \mathcal{L}_{reg} has the best capabilities in synthesizing more natural-looking images compared to the pixel-based baselines.

Additionally, Table 3 presents the evaluation results for all 44 PureSky maps from PolyHaven [pol] on the bunny scene. These results assess the impact of varying roughness on reconstruction quality. When the roughness is set to zero, the baseline approach rapidly overfits the desired sample, guided by the scene’s reflection on the object. However, the low LPIPS score indicates that the synthesized environment map lacks the perceptual fidelity of the orig-

inal, revealing the method’s tendency to cheat by using shortcuts in order to fit the target. This is further supported by the low Peak Signal-To-Noise Ratio, demonstrating the baseline’s production of noisy images.

Furthermore, removing the reflection information quickly compromises the baseline process, leading to optimization failure. In contrast, our approach achieves superior LPIPS and PSNR scores at lower roughness, resulting in the best performance for generating natural images that accurately capture the structure of the final rendering, as evidenced by the LPIPS and SSIM scores in the rendering metrics.

Limitations. While R-SIREN HDR demonstrates promising performance without relying on strong and computationally expensive priors, it shows certain limitations posed by both the single-view inverse rendering and the absence of a latent space. The lack of a strong prior, such as a latent space, makes the model more prone to visual hallucinations, especially when the environment map is not a sky. The model’s sensitivity to the training process of SIREN introduces a trade-off between robustness capabilities and sharpness. Additionally, the method’s reliance on reflective materials for effective inversion can be challenging in downstream tasks where more reflections require the renderer to establish physical consistency between existing reflections and new sketched lighting conditions. Furthermore, when applying R-SIREN HDR to a downstream task, the user’s input specification is crucial. Strokes that do not align with the underlying geometry can lead to physically incorrect and difficult-to-interpret results for the renderer.

6.4. Downstream Task: Appearance Editing

Moving the focus to downstream applications, appearance editing is the task of synthesizing a new environment map based on some user’s input over the original rendering. We used a photo editing program to sketch target images over the initial renderings, capturing our desired lighting conditions. Compared to the previous

Table 3: Comprehensive Evaluation of All Puresky Maps from PolyHaven in the Bunny Scene [pol]. The colors indicate the best result for each level of roughness employed. The proposed R-SIREN HDR \mathcal{L}_{reg} achieves the best performance among all the metrics for the environment map while still fitting well the target (rendering metrics) across different roughness values.

Experiment		Envmap Metrics				Rendering Metrics			
Name	Roughness	↑ PSNR	↓ LPIPS	↓ MSE	↑ SSIM	↑ PSNR	↓ LPIPS	↓ MSE	↑ SSIM
Pixels (baseline)	0.0	11.47963 dB	0.63746	0.07877	0.33030	35.00420	0.18879	0.00465	0.94913
	0.5	9.50746 dB	0.68670	0.12130	0.23377	34.86679	0.20606	0.00507	0.95656
	1.0	9.80608 dB	0.66478	0.11211	0.28438	36.98002	0.20773	0.00337	0.95282
Log(pixels)	0.0	12.65512 dB	0.66607	0.09697	0.36745	31.23284	0.32984	0.04663	0.79454
	0.5	9.39572 dB	0.72085	0.16545	0.28348	24.51719	0.54301	0.08933	0.69031
	1.0	11.01039 dB	0.69084	0.13392	0.33982	29.61368	0.45936	0.05712	0.74746
SIREN HDR \mathcal{L}_{reg}	0.0	11.52815 dB	0.48079	0.08656	0.35959	30.14610	0.24235	0.03281	0.93911
	0.5	11.18299 dB	0.51395	0.09276	0.29486	36.25536	0.21043	0.03165	0.95327
	1.0	10.91892 dB	0.52043	0.09769	0.28518	38.52983	0.21901	0.03354	0.95579
R-SIREN HDR \mathcal{L}_{reg}	0.0	11.77427 dB	0.47657	0.08324	0.37889	31.73236	0.20716	0.01330	0.95850
	0.5	11.37157 dB	0.50795	0.08931	0.31936	38.55692	0.16639	0.01403	0.97359
	1.0	11.09265 dB	0.51780	0.09469	0.30140	40.64959	0.17940	0.01541	0.97294



Figure 9: Effects of the regularization terms over the downstream task. *R-SIREN HDR* already achieves remarkable results without regularization. In addition, employing $\mathcal{L}_{lum} + \mathcal{L}_{l_1}$ results in sharper shadows due to constraining the brightest pixel in a small region. On the other hand, $\mathcal{L}_{lum} + \mathcal{L}_{l_1} + \mathcal{L}_{dists}$ produces slightly blurrier shadows but preserves more of the perceptual features in the environment map.

inverse rendering tasks, we are not computing the loss over a physically correct image because older features, such as reflections over the objects, are still present, creating conflicts that will be blended by the inversion process. The renderer will have to find the best compromise to a physically correct sample that covers all the light conditions expressed in the target. We reduced the luminance regularization weights to give the optimization more freedom, allowing new lights without completely removing older ones. The effects of the regularization terms over the downstream task are illustrated in Fig. 9. We tested various appearance editing methods: shadow translation, light addition, shadow addition, color temperature adjustment, and shadow enhancement/blurring. For shadow translation, we removed the existing shadow by painting over it with the background color and then drew a new shadow in the desired location. To add a new shadow, we applied a dark brush stroke to the desired area. For light addition, we introduced a lighter brush stroke to simulate the desired light reflection on the model. To adjust the color temperature, we applied a color temperature filter. Finally, we enhanced or blurred the shadows by increasing the contrast and

applying Gaussian blur, respectively. Fig. 10 illustrates the synthesized environment maps and the corresponding final renderings.

7. Conclusions and Future Work

In this work, we proposed a method to modify existing environment maps using a lazy-learning optimization technique. This approach leverages the output of a differentiable renderer as a target within the loss function. To preserve the overall brightness characteristics, we incorporated a luminance regularization. Additionally, we used L1 regularization to restrict the HDR light features to small, meaningful sectors, imitating the natural behavior of light distribution in environment maps. Lastly, we applied the DISTS-based regularizer to keep the original perceptual features. We then demonstrated how to transition the image representation from RGB to neural implicit functions, leading to significant improvements in the consistency of the newly generated images. To achieve this, we trained a neural network to represent HDR images and ensured its robustness to small changes in the weights, facilitating optimization guided by the inverse rendering loss. Enforcing robustness during the training

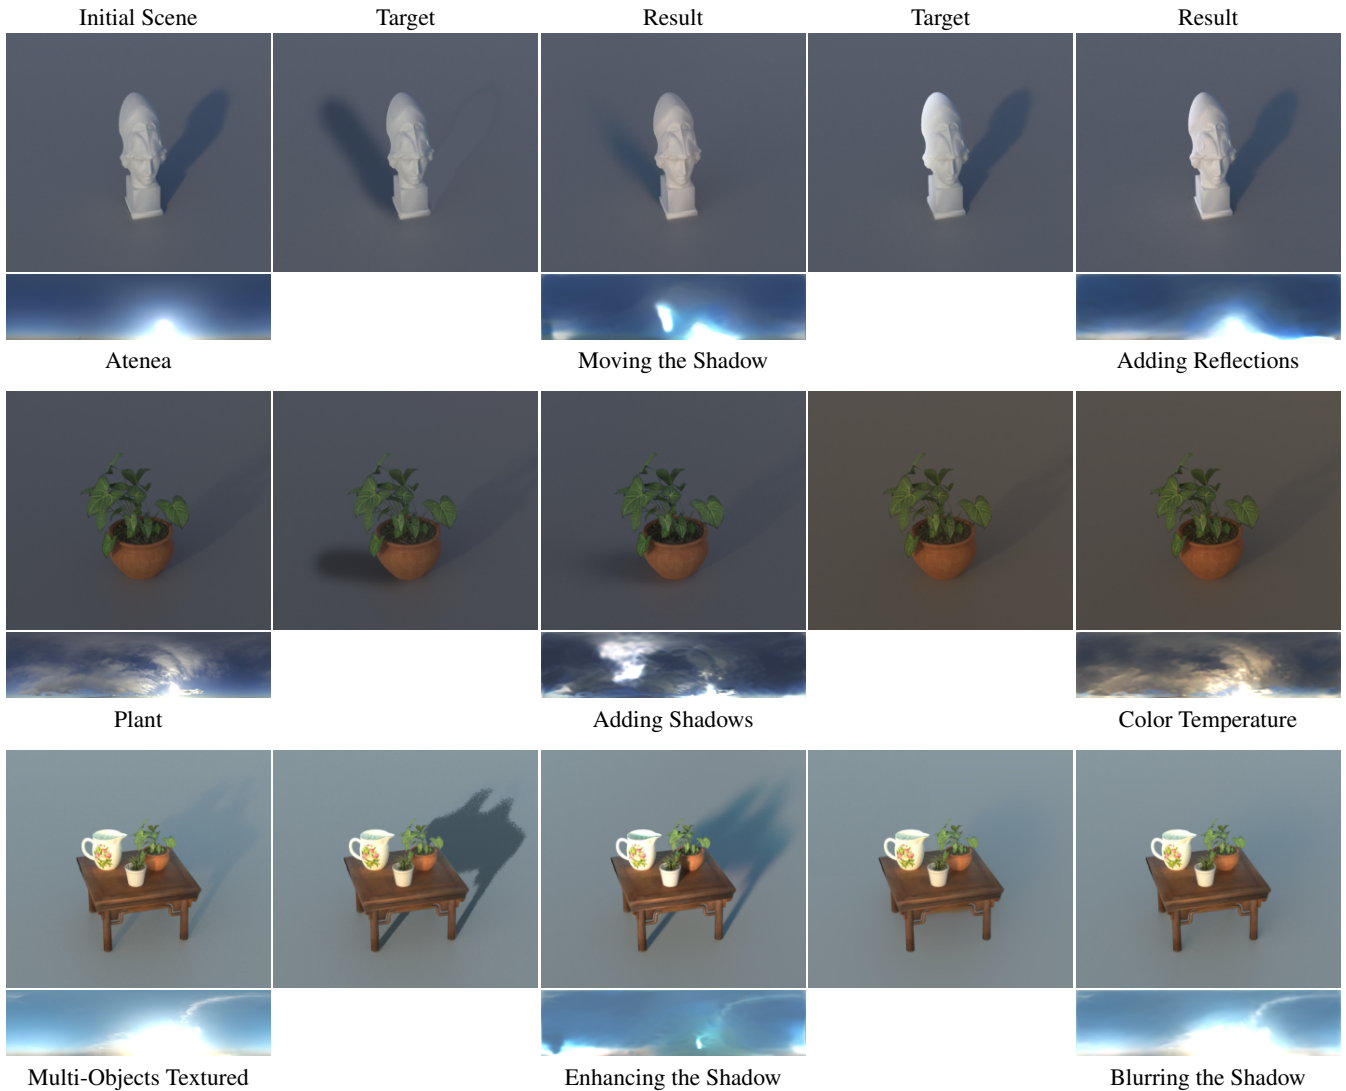


Figure 10: Downstream task. Synthesized environment maps and renderings with various appearance edits: \diamond *Moving the Shadow*: the existing shadow is painted over with the background color, and then a new one is sketched in the desired location. \diamond *Adding Reflections*: the new reflection coming from the light is painted over the 3D model. \diamond *Adding Shadows*: the new shadow is sketched without removing the original ones. \diamond *Color temperature*: the color temperature of the image is altered to a warmer appearance. \diamond *Enhancing the Shadow*: the shadow is coarsely selected and then increased in contrast. \diamond *Blurring the Shadow*: a Gaussian blur filter is applied within the desired region.

of the neural implicit function helped maintain the stability of the overall perceptual features and reduce noise. Finally, we applied the proposed R-SIREN representation to the downstream task of editing the appearance of environment maps by specifying a coarse target crafted with photo editing software. In this way, we were able to obtain novel environment maps such that the final rendering resembles said target, significantly improving the perceptual fidelity and reducing the overall noise. We did so without relying on expensive pre-trained generative models such as Generative Adversarial Networks (GANs) or Denoising Diffusion Probabilistic Models (DDPMs).

In our future research, we aim to explore adversarial implicit functions for 3D model manipulation via inverse rendering, potentially extending to other representations like SDFs. Furthermore, we plan to refine the way we enforce the "local" latent space of neural implicit functions by incorporating diverse samples during perturbations. This includes augmentations of the same sample or introducing new references to bring their representation closer inside the weights space of the trained sample, enabling the appearance editing process to properly blend the different content and perceptual features. By doing so, we expect to enhance the model's ability to generate more realistic and diverse outputs.

Acknowledgments

This work was supported by projects PNRR MUR PE0000013-FAIR under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU and PRIN 2022 project 20227YET9B “AdVVent” CUP code B53D23012830006.

References

- [CTZ*23] CUI J., TIAN Z., ZHONG Z., QI X., YU B., ZHANG H.: Decoupled kullback-leibler divergence loss. *arXiv preprint arXiv:2305.13948* (2023). 3
- [DMWS20] DING K., MA K., WANG S., SIMONCELLI E. P.: Image quality assessment: Unifying structure and texture similarity. *CoRR abs/2004.07728* (2020). URL: <https://arxiv.org/abs/2004.07728>. 6
- [FTSL*22] FUJI TSANG C., SHUGRINA M., LAFLECHE J. F., TAKIKAWA T., WANG J., LOOP C., CHEN W., JATAVALLABHULA K. M., SMITH E., ROZANTSEV A., PEREL O., SHEN T., GAO J., FIDLER S., STATE G., GORSKI J., XIANG T., LI J., LI M., LEBAREDIAN R.: Kaolin: A pytorch library for accelerating 3d deep learning research. <https://github.com/NVIDIAGameWorks/kaolin>, 2022. 3
- [GKS*21] GERA P., KT A., SIRIKONDA D., SAKURIKAR P., NARAYANAN P. J.: Appearance editing with free-viewpoint neural rendering. 2021. *arXiv:2110.07674*. 3
- [HJA20] HO J., JAIN A., ABBEEL P.: Denoising diffusion probabilistic models. In *Advances in neural information processing systems* (2020). 3
- [JSRV22] JAKOB W., SPEIERER S., ROUSSEL N., VICINI D.: Dr.jit: A just-in-time compiler for differentiable rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (July 2022). doi:10.1145/3528223.3530099. 3, 6
- [KBG*15] KRONANDER J., BANTERLE F., GARDNER A., MIANDJI E., UNGER J.: Photorealistic rendering of mixed reality scenes. *Computer Graphics Forum* 34 (2015). URL: <https://api.semanticscholar.org/CorpusID:631737>. 3
- [LTH*23a] LYU L., TEWARI A., HABERMANN M., SAITO S., ZOLLHÖFER M., LEIMKÜHLER T., THEOBALT C.: Diffusion posterior illumination for ambiguity-aware inverse rendering. *ACM Transactions on Graphics* 42, 6 (2023). 3
- [LTH*23b] LYU L., TEWARI A., HABERMANN M., SAITO S., ZOLLHÖFER M., LEIMKÜHLER T., THEOBALT C.: Diffusion posterior illumination for ambiguity-aware inverse rendering. *ACM TOG* 42, 6 (2023), 1–14. 2
- [MBBM24] MIRZA M. H., BRIGLIA M. R., BEADINI S., MASI I.: Shedding more light on robust classifiers under the lens of energy-based models. In *European Conference on Computer Vision* (2024). 3
- [MMS*18] MADRY A., MAKELOV A., SCHMIDT L., TSIPRAS D., VLADU A.: Towards deep learning models resistant to adversarial attacks. 3
- [NIJ21] NICOLET B., JACOBSON A., JAKOB W.: Large steps in inverse rendering of geometry. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 40, 6 (Dec. 2021). doi:10.1145/3478513.3480501. 3
- [OMSI07] OKABE M., MATSUSHITA Y., SHEN L., IGARASHI T.: Illumination brush: Interactive design of all-frequency lighting. In *15th Pacific Conference on Computer Graphics and Applications (PG'07)* (2007), pp. 171–180. doi:10.1109/PG.2007.9. 3
- [Pel10] PELLACINI F.: envylight: an interface for editing natural illumination. *ACM SIGGRAPH 2010 papers* (2010). URL: <https://api.semanticscholar.org/CorpusID:7019773>. 3
- [pol] Polyhaven dataset of hdri environment maps. <https://polyhaven.com/>. 6, 8, 9
- [RMKM23] ROUHSEDAGHAT M., MONAJATIPOOR M., KUO C.-C. J., MASI I.: MAGIC: Mask-Guided Image Synthesis by Inverting a Quasi-Robust Classifier. In *In Proceedings of AAAI Conference on Artificial Intelligence* (2023). 3
- [SMB*20] SITZMANN V., MARTEL J. N., BERGMAN A. W., LINDELL D. B., WETZSTEIN G.: Implicit neural representations with periodic activation functions. In *Proc. NeurIPS* (2020). 2, 4
- [WBSS04] WANG Z., BOVIK A., SHEIKH H., SIMONCELLI E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. doi:10.1109/TIP.2003.819861. 8
- [WLS*21] WANG P., LI Y., SINGH K. K., LU J., VASCONCELOS N.: Imagine: Image synthesis by image-guided model inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 3681–3690. 3
- [WXW20] WU D., XIA S., WANG Y.: Adversarial weight perturbation helps robust generalization. *arXiv: Learning* (2020). URL: <https://api.semanticscholar.org/CorpusID:222322720>. 3, 5
- [XPH05] XU R., PATTANAİK S. N., HUGHES C.: High-dynamic-range still-image encoding in jpeg 2000. *IEEE Computer Graphics and Applications* 25 (2005), 57–64. URL: <https://api.semanticscholar.org/CorpusID:1366624>. 4
- [ZIE*18] ZHANG R., ISOLA P., EFROS A. A., SHECHTMAN E., WANG O.: The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2018), pp. 586–595. doi:10.1109/CVPR.2018.00068. 8
- [ZYZ*19] ZHANG H., YU Y., JIAO J., XING E. P., GHAOU L. E., JORDAN M. I.: Theoretically principled trade-off between robustness and accuracy. 3