
Latent Intrinsic Properties Emerge from Training to Relight

Xiao Zhang¹ William Gao¹ Seemandhar Jain² Michael Maire¹
D.A. Forsyth² Anand Bhattad³

¹University of Chicago ²University of Illinois Urbana Champaign

³Toyota Technological Institute at Chicago

zhang7@uchicago.edu bhattad@ttic.edu

<https://latent-intrinsics.github.io/>

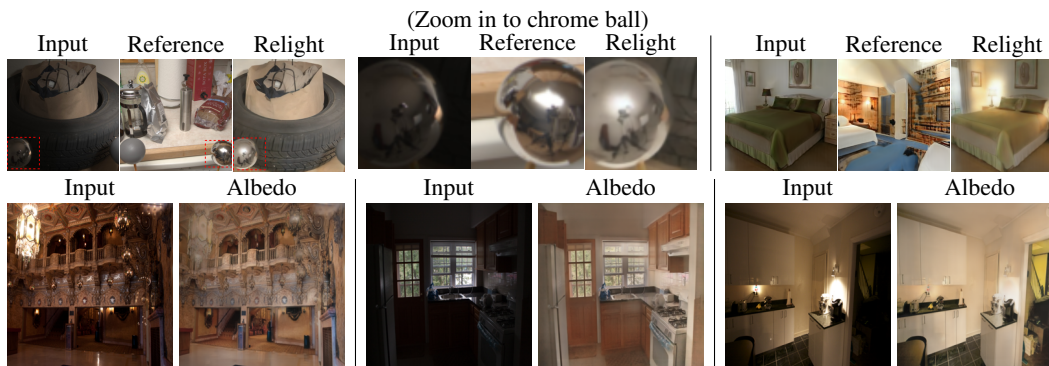


Figure 1: We describe a purely data-driven image relighting model. Our model recovers latent variables representing scene intrinsic properties from one image, latent variables representing lighting from another, then applies the lighting to the intrinsics to produce a relighted scene (**top row**). There is no physical model of intrinsics, extrinsics or their interaction. Our model relights images of real scenes with SOTA accuracy and is more accurate than current supervised methods. Note how, for the chrome ball detail in **top center**, the specular reflections on the chrome ball (which give an approximate environment map) change when the extrinsics are changed. Note how our model ascribes lighting to visible luminaires when it can (**top right**), despite the absence of any physical model. A physical model accounts only for effects in that model, and most physical models of surfaces are approximate; in contrast, a latent intrinsic model accounts for whatever produces substantial effects in training data. Latent intrinsics yield albedo in a natural fashion (light the scene with an appropriate illuminant). **Bottom row** shows SOTA albedo estimates recovered from our latent intrinsics.

Abstract

Image relighting is the task of showing what a scene from a source image would look like if illuminated differently. Inverse graphics schemes recover an explicit representation of geometry and a set of chosen intrinsics, then relight with some form of renderer. However error control for inverse graphics is difficult, and inverse graphics methods can represent only the effects of the chosen intrinsics. This paper describes a relighting method that is entirely data-driven, where intrinsics and lighting are each represented as latent variables. Our approach produces SOTA relightings of real scenes, as measured by standard metrics. We show that albedo can be recovered from our latent intrinsics without using any example albedos, and that the albedos recovered are competitive with SOTA methods.

1 Introduction

Relighting – taking an image of a scene, then adjusting it so it looks as though it had been taken under another light – has a range of applications, including commercial art (e.g., photo enhancement) and data augmentation (e.g., making vision models robust to varying illumination). As a technical problem, relighting is very hard indeed, likely because how a scene changes in appearance when the light is changed can depend on complex surface details (grooves in screws; bark on trees; wood grain) that are hard to capture either in geometric or surface models.

One common approach to relighting a scene is to infer scene characteristics (geometry, surface properties) using inverse graphics methods, then render the scene with a new light source. This approach is fraught with difficulties, including the challenge of selecting which material properties to infer and managing error propagation. These methods perform best in outdoor scenes with significant shadow movements but struggle with indoor scenes where interreflections create complex effects (Section 4.2).

As this paper demonstrates, a purely data-driven method offers an attractive alternative. A source scene, represented by an image, is encoded to produce a latent representation of intrinsic scene properties. A source illumination, represented by another image, is encoded to produce a latent representation of illumination properties. These intrinsic and extrinsic properties are combined and then decoded to produce the relighted image. As a byproduct of this training, we find that the latent representation of intrinsic scene properties behaves like an albedo, while another latent representation acts as a lighting controller.

Our model can capture complex scene characteristics without explicit supervision by capturing intrinsic properties as latent phenomena, making it particularly appealing. In contrast to a physical model, we are not required to choose which effects to capture. This latent approach reduces the need for detailed geometric and surface models, simplifies the learning process, and enhances the model’s ability to generalize to diverse and unseen scenes. This makes it highly applicable to a wide range of real-world scenarios.

Contributions: We present the first fully data-driven relighting method applicable to images of real complex scenes. Our approach requires no explicit lighting supervision, learning to relight using paired images alone. We demonstrate that this method effectively trains and generalizes, producing highly accurate relightings. Furthermore, we demonstrate that albedo-like maps can be generated from the model without supervision or prior knowledge of albedo-like images. These intrinsic properties emerge naturally within the model. We validate our model on a held-out dataset, applying target lighting conditions from various scenes to assess its generalization capability and precision in real-world scenarios (Section 4.2).

2 Related Work

Intrinsic Images. Humans have been known to perceive scene properties independent of lighting since at least 1867 [46, 21, 4, 20]. In computer vision, the idea dates to Barrow and Tenenbaum [3] and comprises at least depth, normal, albedo, and surface material maps. Depth and normal estimation are now well established (eg [24]). There is a rich literature on albedo estimation (dating to 1959 [30, 31]!). A detailed review appears in [16], which breaks out methods as to what kinds of training data they see. Early methods do not see any form of training data, but more recently both CGI data and manual annotations of relative lightness (labels) have become available. Early efforts, such as SIRFS [1], focused on using shading information to recover shape, illumination, and reflectance, highlighting the importance of modeling these factors for intrinsic image analysis. Recent strategies include: deep networks trained on synthetic data [33, 23, 15]; and conditional generative models [29].

The weighted human disagreement ratio (WHDR) evaluation framework was introduced by [5] using the IIW dataset. This is a dataset of human judgments that compare the absolute lightness at pairs of points in real images. Each pair is labeled with one of three cases (first lighter; second lighter; indistinguishable) and a weight, which captures the certainty of labelers. One evaluates by computing a weighted comparison of algorithm predictions with human predictions; WHDR scores can be improved by postprocessing because most methods produce albedo fields with very slow gradients, rather than piecewise constant albedos. [10] demonstrate the value of “flattening” albedo (see also [39]); [11] employ a fast bilateral filter [2] to obtain significant improvements in WHDR.

Using Intrinsic Images for Relighting. Bhattad and Forsyth [6] demonstrated that intrinsic images could be used for reshading inserted objects. This approach can be extended by adjusting the shading in both the foreground and background to eliminate discrepancies [12]. Intrinsic images and geometry-aware networks have been used for multi-view relighting [41]. StyLitGAN [9] introduced a method to relight images by identifying directional vectors in the latent space of StyleGAN, but can only relight StyleGAN generated images and requires explicit albedo and shading to guide relighting. It can be extended to real images using a GAN inversion, but does not generalize [7]. LightIt [28] controls lighting changes in image generation using diffusion models, by conditioning on shading and normal maps to achieve consistent and controllable lighting. Like these methods, we use intrinsics and extrinsics to relight, but ours are latent, with no explicit physical meaning.

Color Constancy. Image color is ambiguous: a green pixel could be the result of a white light on a green surface, or a green light on a white surface. Humans are unaffected by this ambiguity (eg [21, 4]; recent review in [50]). There is extensive computer vision literature; a recent review appears in [32]. We do not estimate illumination color but estimate a single color correction (Section 4.2).

Lighting Estimation and Representation. Accurate lighting representation is crucial for tasks like object insertion and relighting. Traditional methods used parametric models such as environment maps and spherical harmonics to represent illumination [13, 42]. Debevec’s seminal work [13] on recovering environment maps from images of mirrored spheres set the foundation for many subsequent works. Methods by Karsch et al. [26, 27], Gardner et al. [17, 18], Garon et al. [19] and Weber et al. [49] advanced the field by using learned models to recover parametric, semi-parametric or panoramic representations of illumination. Recent approaches include representing illumination fields as dense 2D grids of spherical harmonic sources [34, 36] or learning 3D volumes of spherical Gaussians [48]. These methods can model complex light-dependent effects but require extensive CGI datasets for training [43, 35]. Our approach diverges by not relying on labeled illumination representations or CGI data, instead producing abstract representations of illumination through deep features without specific physical interpretations.

Image-based Relighting. Other works focus on portrait relighting using deep learning [45, 55, 40, 44], which are typically specialized to faces and trained on paired or light-stage data. Self-supervised methods for outdoor image relighting leverage single-image decomposition with parametric outdoor illumination, benefiting from simpler lighting conditions dominated by sky and sunlight [54, 37]. [22] introduced a self-attention autoencoder model to re-render a source image to match the illumination of a guide image, focusing on separating scene representation and lighting estimation with a self-attention mechanism for targeted relighting. Similarly, [51] proposed a depth-guided image relighting, which combines source and guide images along with their depth maps to generate relit images. In contrast, our work shows that intrinsic properties relevant to relighting can emerge naturally from training to relight, facilitating complex scene relighting without the need for explicit lighting estimation. We compare with both [22] and [51] for relighting capabilities on real scenes.

Emergent Intrinsic Properties. Bhattad et al. [8] and Du et al. [14] demonstrate that intrinsic images can be extracted from generative models using a small intrinsic image dataset obtained from pretrained off-the-shelf intrinsic image models. Our work explores how intrinsic image properties emerge as a result of training a model for relighting, without the need for an intrinsic image dataset.

3 Learning Latent Intrinsic from Relighting.

Our relighting model can be seen as a form of autoencoder. One encoder computes a latent representation of scene intrinsics from an image of a target scene; another computes a latent representation of scene extrinsics from an image of a placeholder scene in the reference lighting. These are combined, then decoded into a final image of the target scene in the reference lighting. Losses impose the requirements that (a) the final image is right and (b) the latent intrinsics computed for a scene are not affected by illumination. The procedure for combining intrinsics and extrinsics is carefully designed to make it very difficult for intrinsic features of the placeholder scene to “leak” into the final image.

3.1 Model structure

Decoder setup: Write $I_s^l \in \mathbf{R}^{H \times W \times 3}$ for the input image, captured from scene s with lighting configuration l . Training uses pairs $I_s^{l_1}$ and $I_s^{l_2}$, representing the same scene s under different lighting

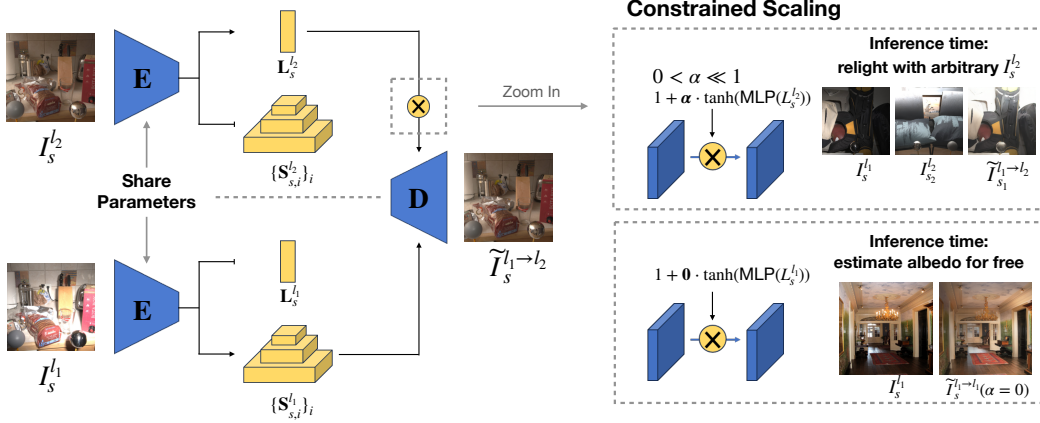


Figure 2: The network diagram of our relighting model. The model functions as an autoencoder, comprising an encoder E and a decoder D . **Left Half:** The encoder E maps input image I_s^l , captured under scene s and lighting l , to low-dimensional extrinsic features L_s^l and set of intrinsic features map $\{S_{s,i}^l\}_i$. The decoder D then generates new images based on these intrinsic and extrinsic representations. **Right Half:** We employ *constrained scaling* for the injection of L_s^l , utilizing $0 < \alpha \ll 1$ to regularize the information passed from L_s^l , thereby enforcing a low-dimensional parameterization of the extrinsic features. We train our system to relight target images given input paired with images captured under the same scene s . During inference, our model demonstrates the ability to generalize to arbitrary reference images for relighting and can estimate albedo for free.

conditions l_1 and l_2 . The model *does not see* detailed lighting information (for example, the index of the lighting) during training, because standardizing lighting settings across various scenes is often impractical.

Write E for the encoder, D for the decoder. The encoder must produce the intrinsic and extrinsic representations from the input image. Write $S_{s,i}^l \in \mathbf{R}^{(H_i \times W_i) \times C_i}$ for spatial feature maps yielding the intrinsic representation, with i for the layer index, and $L_s^l \in \mathbf{R}^C$ for extrinsic features; we have:

$$E(I_s^l) := \{S_{s,i}^l\}_i, L_s^l \quad (1)$$

We apply L2 normalization along the feature channel to both sets of features. During training, we add random Gaussian noise to the input image to enhance semantic scene understanding capabilities:

$$E(I_s^l + \sigma\epsilon) := \{S_{s,i}^l\}_i, L_s^l \quad (2)$$

Decoder setup: The decoder D relights $I_s^{l_1}$ using extrinsic features extracted from $I_s^{l_2}$:

$$D(\{S_{s,i}^{l_1}\}, L_s^{l_2}) := \tilde{I}_s^{l_1 \rightarrow l_2} \quad (3)$$

We optimize the autoencoder using a pixel-wise loss on both relighted and reconstructed images:

$$\mathcal{L}_{\text{relight}} := \mathcal{L}_{\text{pixel}}(\tilde{I}_s^{l_1 \rightarrow l_2}, I_s^{l_2}) + \mathcal{L}_{\text{pixel}}(\tilde{I}_s^{l_2 \rightarrow l_2}, I_s^{l_2}) \quad (4)$$

where $\mathcal{L}_{\text{pixel}}$ represents the pixel-wise losses: L2 distance on pixels; structural similarity index (SSIM) [47]; and l2 distance on image spatial gradient (weights 10, 0.1 and 1 respectively).

3.2 Intrinsicness

Intrinsicness: Our model should report the same latent intrinsic for the same scene in different lightings, so we apply the following loss to the encoder:

$$\mathcal{L}_{\text{intrinsic}} := \sum_i \|S_{s,i}^{l_1} - S_{s,i}^{l_2}\|_2 + 1e-3 \cdot \mathcal{L}_{\text{reg}}(S_{s,i}^{l_1}) \quad (5)$$

where \mathcal{L}_{reg} is a regularization term on intrinsic features, defined as follows:

$$\mathcal{L}_{\text{reg}}(\mathbf{S}) := \|\mathbf{R}(\mathbf{S}) - \mathbf{R}(\hat{\mathbf{S}})\|_2 \quad (6)$$

$$\mathbf{R}(\mathbf{S}) := \log \det \left(\mathbf{I} + \frac{d}{n\lambda^2} \mathbf{S}^\top \mathbf{S} \right) \quad (7)$$

where $R(\mathbf{S})$ is the coding rate [53] for a matrix $\mathbf{S} \in \mathbb{R}^{n \times d}$ with each row l2 normalized, under a distortion constant λ . $\hat{\mathbf{S}}$ is a random matrix with the same shape of \mathbf{S} and each row of $\hat{\mathbf{S}}$ is sampled from uniform hyperspherical distribution at the start of learning. In Eqn.5, $R(\hat{\mathbf{S}})$ serves as the optimization target of $R(\mathbf{S})$ to encourage the \mathbf{S} to uniformly spread out in the hyperspherical space. This strategy is now widely used in self-supervised learning; without the regularization term, the model can minimize the feature distance by simply collapsing the distribution of $\mathbf{S}_{s,i}^l$ with small variance, which will not yield effective lighting invariance.

3.3 Combining intrinsics and extrinsics

The placeholder scene is necessary to communicate illumination to the model, but has important nuisance features. Intrinsic information from this scene could “leak” into the final image, spoiling results. We introduce *constrained scaling*, a structural bottleneck that restricts the amount of information transmitted from the learned extrinsic features.

Write $\mathbf{F} \in \mathbb{R}^{h \times w \times c}$ for the feature map fed to the decoder. Constrained scaling combines intrinsic and extrinsic features by

$$\tilde{\mathbf{F}} := \mathbf{F} \odot (1 + \alpha \cdot \tanh(\text{MLP}(\mathbf{L}_s^l))) \quad (8)$$

where MLP, a series of fully connected layers with non-linear activation, aligns \mathbf{L}_s^l to the latent channel dimension of \mathbf{F} and $\alpha \ll 1$ is a small non-negative scalar (we use 5e-3). This approach means that any single extrinsic feature vector has little effect on the feature – for an effect, the extrinsics must be pooled over multiple locations. Illumination fields tend to be spatially smooth, supporting the insight that enforced pooling is a good idea.

Constrained scaling compresses latent vectors into a very small numerical range, making learning difficult. We use a regularizer to promote a uniform distribution of \mathbf{L}_s^l , which improves optimization. In particular, we have

$$\mathcal{L}_{\text{extrinsic}} := \mathcal{L}_{\text{reg}}(\mathbf{L}_s^l) \quad (9)$$

By choosing $\alpha \ll 1$ and training model with uniform regularization term Eqn.9, we effectively push the lighting code to uniformly spread over $[-\alpha, \alpha]$ where the absolute value of each channel indicates the strength of the light. As a side effect, by setting $\alpha = 0$ to disable the contribution of the lighting code, we get image albedo estimation from our model for free.

Our final training objective is weighted combination of all individual loss terms:

$$\mathcal{L} := \mathcal{L}_{\text{relight}} + 1e-1 \cdot \mathcal{L}_{\text{intrinsic}} + 1e-4 \cdot \mathcal{L}_{\text{extrinsic}} \quad (10)$$

4 Experiments

We will first provide a brief description of our experimental procedure (Sec 4.1), followed by a discussion on how we evaluate the various relighting capabilities of our approach, including its strong generalization across datasets with different distributions (Sec 4.2). Finally, we will present the emergent albedo that is recovered from the latent intrinsic without using any albedo-like images (Sec 4.3).

4.1 Experiment Details

Training Details We train our model using the MIT multi-illumination dataset [38], which includes images of 1,015 indoor scenes captured under 25 fixed lighting, totaling 25,375 images. We follow the official data split and train our model on the 985 training scenes. During training, we randomly sample pairs of images from the same scene under different lighting conditions and perform random spatial cropping, with the crop ratio randomly selected between 0.2 and 1.0, followed by resizing the cropped image to a resolution of 256x256. For further details, please refer to our appendix.

4.2 Evaluating image relighting

Relighting on the Multi-illumination dataset: We relight images of scenes in the test set using reference images from the test set, then compare to the correct known relighting from the test set

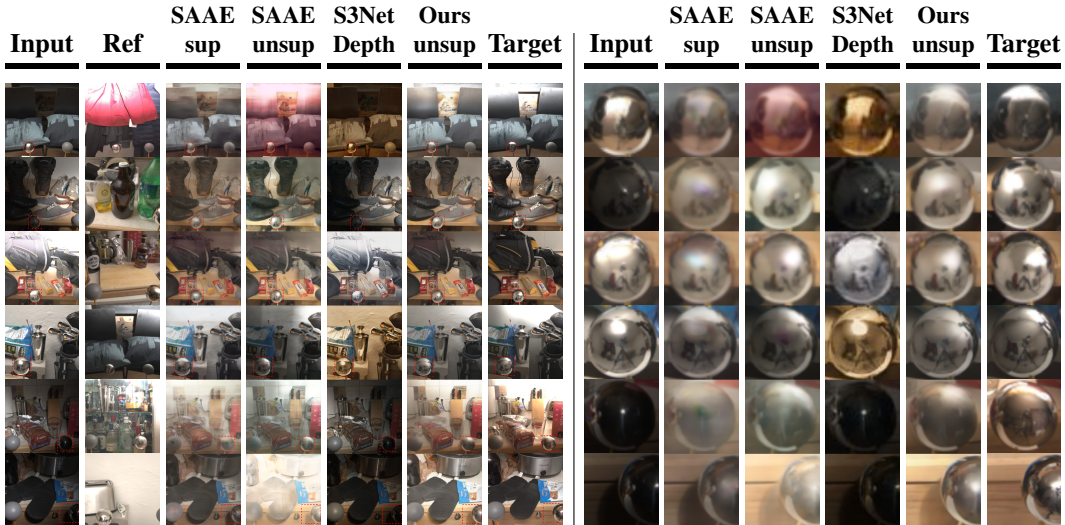


Figure 3: Our method outperforms all other approaches in estimating light and rendering the scene. The Unsupervised SA-AE [22] method fails by incorporating intrinsic elements from reference images. The S3Net [51] approach struggles with rendering when using unpaired reference images. **Right:** A zoomed-in view of the chrome ball was used as a probe to evaluate detail preservation in the environment map. Our method effectively retains the intricate room layout and accurately renders the appropriate lighting effects.

Methods	Labels	Raw Output		Color Correction	
		RMSE↓	SSIM↑	RMSE↓	SSIM↑
Input Img	-	0.384	0.438	0.312	0.492
SA-AE [22]	Light	0.288	0.484	0.232	0.559
SA-AE [22]	-	0.443	0.300	0.317	0.431
S3Net [51]	Depth	0.512	0.331	0.418	0.374
S3Net [51]	-	0.499	0.336	0.414	0.377
Ours($\sigma = 0$)	-	0.326	0.232	0.242	0.541
Ours(w/o \mathcal{L}_{reg})	-	0.315	0.462	0.232	0.550
Ours	-	0.297	0.473	0.222	0.571

Table 1: We assess the quality of image relighting using the multi-illumination dataset [38]. Our method, when evaluated on raw output, significantly outperforms all other unsupervised approaches and achieves competitive results compared to the supervised SA-SA [22], which requires ground truth light supervision. When we correct the colors by eliminating global color drift caused by light ambiguity, our method surpasses all other approaches. Additionally, warming up the model as a denoising auto-encoder proves beneficial compared to when it is not warmed up ($\sigma = 0$).

α	Raw Output		Color Correction	
	RMSE ↓	SSIM ↑	RMSE ↓	SSIM ↑
∞	0.471	0.287	0.352	0.407
1e-2	0.314	0.444	0.238	0.546
5e-3	0.297	0.473	0.222	0.571
1e-3	0.312	0.453	0.256	0.524
5e-4	0.309	0.460	0.253	0.533

Table 2: We analyze the impact of α on relighting quality using the multi-illumination dataset [38]. Setting α to ∞ , which removes the scaling constraints, results in poor relighting quality, indicating that restricting information from extrinsic sources significantly improves generation quality. Within a limited parameter search, 5e-3 yields the best results.

using various metrics. For each input image, we randomly sample reference images from different scenes and lighting conditions. To reduce the effect of randomness in comparing different relighting strategies, we select 12 random reference images for each input image, and maintain the same image-reference pairs when evaluating different models. We report the results, measured in RMSE and SSIM, in Table 1. We report these metrics both for absolute predictions and for predictions where any global color shift is corrected by a single, least-squares scale of each predicted color layer (i.e. one scale for R; one for G; one for B). This color correction allows us to distinguish between spatial errors and global color shifts; these appear to have a significant effect, possibly because there are visible color shifts present in some of the dataset images.

In Table 1, we compare to SA-AE [22], a model that requires a ground truth light index for supervision, and S3-Net, which needs a ground truth depth map as a conditional input. For S3-Net, we use a state-of-the-art depth estimator to provide pseudo-GT on the relighting dataset as input. For a fair comparison to our model, which does not require any supervision outside of the ground truth



Figure 4: Latent extrinsics can be interpolated successfully; **leftmost** and **rightmost** columns are images from the multi-illumination dataset, and intermediate images are obtained by linear interpolation on the latent extrinsics (light-dependent representations), then decoding. Note how the light seems to "move" across space.

relighting, we also report results for modified versions of the baselines trained without additional supervision. For SA-AE, we train their light estimation model and relighting model end-to-end by removing the loss from light supervision. For S3-Net [51], we simply remove the depth from the model’s input.

Without color correction, only light-supervised SA-AE slightly outperforms our model, while all other baselines are significantly worse. The unsupervised version of SA-AE performs much worse because their light estimator struggles to distinguish the extrinsic from the intrinsic components. Specifically, SA-AE also parameterizes the extrinsic as a lower-dimensional representation but without the constrained scaling that our model uses. As a result, the estimated extrinsic from their unsupervised model also carries intrinsic information, and one can see “leaks”. S3-Net performs worse in both versions since they concatenate input and reference images before feeding them into the models, which significantly affects the model’s generalization ability, especially during test time when we use images from different scenes as references.

On color-corrected images, our approach outperforms all methods, including the light-supervised version of SA-AE, indicating that, up to the constant color drift, our extrinsic estimation network is at least as good as, or even better than, a light estimation network trained with supervision. Removing the denoising setup from our model ($\sigma = 0$) results in worse performance in both cases due to inferior semantic scene understanding. We additionally provide ablation studies on the choices of α in Table 2 and find $\alpha = 5e - 3$ produces the best results.

Each image in the multi-illumination dataset shows a chrome ball, which gives a good estimate of an environment map for that image. Correctly rendering the effects of lighting changes on these chrome balls appears to be extremely difficult; the changes are substantial, and concentrated in a small region of the image (so correct representation of these changes has little effect on typical image losses). Figure 3 shows a crop of our results around this chrome ball. Our method represents these changes well; we are aware of no other results reported for this effect. Compared to other approaches, our model accurately preserves the room layout, even in cases of extreme light changes.

Unlike classical rendering models that use a specific parameterized form to represent extrinsics, our framework learns an implicit extrinsic representation. However, we can still parameterize the learned extrinsic representation to create new light sources. In Figure 4, we demonstrate this capability by rendering images using interpolated extrinsic representations.

Relighting synthetically relighted images from StyLitGAN: StyLitGAN [9] is a recent method that can produce multiple illuminations of a single generated room scene by manipulating StyleGAN latents appropriately. In the multi-illumination dataset, reference light and target images tend to share a strong spatial correlation in light patterns. In contrast, StyLitGAN generates extremely challenging images where very significant changes in lighting occur. Furthermore, StyLitGAN images have visible luminaires. To relight the input, the model must infer high-level concepts rather than simply copying the spatially corresponding light patterns from the reference. We train our model using StyLitGAN images to evaluate generalization qualitatively (quantitative evaluation would be of dubious value, because StyLitGAN images are generated rather than real). Figure 5 shows results. Notice how our method successfully relights from references, achieves brighter illuminations by turning on luminaires (here bedside lights), achieves darker scenes by turning off luminaires, and is somewhat less inclined to invent luminaires than StyLitGAN is. The model knows that light must come from somewhere, and how the effects of light are distributed.

Zero-Shot Relighting: In Figure.6, we show our model’s strong generalization by applying the model solely trained on multi-illumination dataset—without additional training or fine-tuning—to



Figure 5: Qualitative results for relighting interior scenes using our relighter trained on images obtained from StyLitGAN (which produces multiple illuminations of a generated scene). StyLitGAN has a strong tendency to increase or decrease illumination by adjusting luminaires, typically bedside lights but also light coming through French windows, etc. On the **left**, where the reference lighting tends to be brighter and more concentrated, notice how for the two top images, our relighter has identified and "turned up" the bedside lights; for the third, it has resisted StyLitGAN's tendency to invent helpful luminaires (there isn't a bedside light where StyLitGAN imputed one, as close inspection shows). On the **right**, where the reference lighting is much more uniform, our relighter has achieved this by "turning down" bedside lights. This is an emergent phenomenon; the method is not supplied with any explicit luminaire model or labeled data.

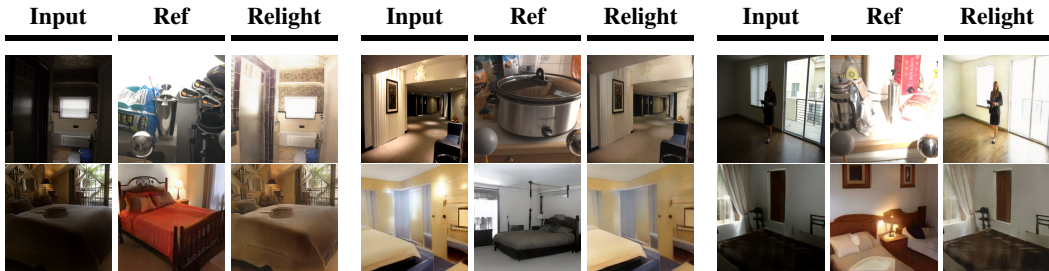


Figure 6: **Zero-Shot Relighting.** Our relighting model, trained only on the multi-illumination dataset, generalizes well to out-of-distribution images, as shown on the IIW dataset (first row) and StyleGAN images (second row). It accurately infers scene geometry and lighting. Note that it identifies and turns on the bedside lamps in StyleGAN images despite having no training in bedroom images. This demonstrates the model's strong generalization ability and the model clearly "knows" something about light sources.

relight IIW and StyleGAN-generated images. Despite the significant distribution shift in lighting patterns and room setup, our model accurately identifies luminaires and relights images.

4.3 Zero-shot albedo evaluation

Constrained scaling allows us to infer albedo without any decoding (and without any albedo data!) by setting $\alpha = 0$ during inference. We benchmark these albedo estimates using the WHDR metric on the IIW [5] dataset (Section 2). We use WHDR because it is widely used and allows comparisons, but existing literature records significant problems in interpreting the measure [16, 6, 29]. Among other irritating features, the metric seems to prefer odd colors, and can be hacked by heavily quantized albedo maps. As is standard, we obtain lightness by averaging R, G, and B albedo and compute relative lightness of two pixel locations i_1, i_2 by comparing to a confidence threshold δ :

$$\tilde{J}_{i,\delta}(\bar{R}) = \begin{cases} 1 & \text{if } \bar{R}_{i_1}/\bar{R}_{i_2} > 1 + \delta \\ 2 & \text{if } \bar{R}_{i_2}/\bar{R}_{i_1} > 1 + \delta \\ E & \text{otherwise} \end{cases} \quad (11)$$

Methods	labels	Flat	Tune δ	WHDR
Intrinsic Diffusion [29]	CG	No	No	22.61
Intrinsic Diffusion[29]	CG	Yes	Yes	17.10
Inverser Render[52]	No	No	No	21.40
BBA[16]	No	No	Yes	17.04
Ours	No	No	No	28.97
Ours	No	No	Yes	19.09
Ours	No	Yes	Yes	15.81

Table 3: We benchmark our albedo estimation on test set of IIW dataset [5] and compare with others, though the reliability has been questioned by recent papers [16]. Flat denotes postprocessing images with flattening [10]. Despite our model never being trained on albedo maps or CG data, our best configuration significantly outperforms all other methods suggesting our model learns high-quality intrinsic representations

α	WHDR			
	$\delta = 0.1$		optimal δ	
	w/ F	w/o F	w/ F	w/o F
1e-2	17.64	28.97	15.81	19.09
5e-3	18.93	31.81	16.02	19.53
1e-3	18.00	29.77	15.84	19.13
5e-4	18.04	29.62	15.85	19.12

Table 4: We conduct ablation experiments to assess the impact of α on the quality of albedo. "w/F" and "w/o F" denote post-processing images with and without flattening [10], respectively. The setting of $\delta = 0.1$ and w/o F is the most affected by α . Despite this, all values of α achieve high performance in our optimal configurations.

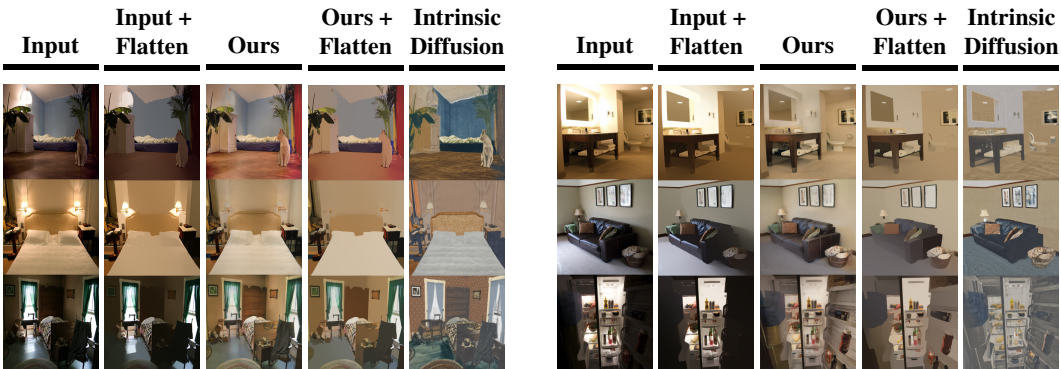


Figure 7: Qualitative Comparison of **Emergent Albedo from Latent Intrinsic** on the IIW Dataset. Although our model has never been trained on any albedo-like maps, it effectively removes the effects of external light and dark shadows from the input. In contrast, Intrinsic Diffusion [29], a supervised method trained on large computer graphics data, often produces color-drifted estimations, likely due to the domain shift between CG data and real images. Observe the subdued lighting around the mirrors (top row, right) in our recovered albedo. Also, pay attention to all the details inside the refrigerator, which are visible in our recovered albedos (bottom row; right) compared to intrinsic diffusion. For comparison, we also display naive flattening (in the second column), which by itself cannot effectively reduce the strong lighting effects.

The resulting classification (one lighter than two; two lighter than one; equivalent) is then compared to human annotations J using the confidence score w_i for each annotation pair. We report WHDR on the IIW test split in Table 3 to facilitate comparison with other approaches. Since our model is not trained with any albedo maps or computer-generated images, we need to adjust the threshold for the optimal performance. Following prior work, we optimize δ on the training split, which significantly improves our performance from 28.97 to 19.09. Additionally, we enhance our performance by post-processing our albedo map using flattening [10], an optimization technique to further reduce color variations. With this improvement, our results reach 15.81, substantially outperforming the intrinsic diffusion model [29], a diffusion-based albedo regression model trained on computer graphics data. In Figure 7, we show some qualitative comparisons to intrinsic diffusion. We observe that our method effectively removes external lighting effects and does not suffer from color drift due to domain gap unlike intrinsic diffusion, which is trained on CG data.

Sensitivity to light changes: Albedo are scene properties that are independent of lighting changes. In Figure. 8, we qualitatively assess this characteristic by varying lighting conditions, comparing our approach with the state-of-the-art supervised method, Intrinsic Diffusion [29]. Our method demonstrates consistent and accurate estimations that remain stable even under extreme lighting variations. In contrast, Intrinsic Diffusion [29] shows significant deviation from the natural color distribution and are sensitive to lighting changes.



Figure 8: Qualitative comparison of albedo stability under varying lighting conditions. Images shown are from the multi-illumination dataset test split. The top row features images under different lighting environments. The middle row presents estimated albedos obtained from Intrinsic Diffusion [29], while the bottom row shows the recovered albedos from the latent intrinsic representation. Intrinsic Diffusion has large color drift and is sensitive to changes in lighting. In contrast, the **albedos recovered from latent intrinsics remain stable under lighting changes, even in extreme conditions.**

5 Discussion, Limitations and Future Work

Our method presents an important advancement in image relighting by demonstrating that intrinsic properties such as albedo can emerge naturally from training on relighting tasks without explicit supervision. This finding simplifies the relighting process, eliminating the need for detailed geometric and surface models and enhancing the model’s ability to generalize across diverse and unseen scenes. By encoding scene and illumination properties as latent variables, we achieve accurate and flexible relighting. Our findings will have implications for various fields such as virtual reality and cinematic post-production. This approach reduces the learning process’s complexity and offers a new perspective on designing deep learning models to capture and utilize intrinsic scene properties. These findings can guide future research toward a more efficient and scalable relighting approach, encouraging the development of models that can handle various lighting conditions and scene complexities.

The current taxonomy of surface intrinsics—typically, depth, normal, albedo, and perhaps specular albedo and roughness—is quite limiting (compare human language for surface properties [4]). Our method, which computes latent intrinsic and extrinsic representations from images and combines these to transfer lighting conditions across scenes, captures physical concepts like luminaire and albedo without explicit physical parametrization. This ability to represent significant image effects without choosing a surface model offers substantial flexibility.

However, our method has several limitations. It relies on pairs of relighted data captured in the same scene, which can be resource-intensive to obtain. Additionally, it does not cope well with saturated pixel values common in LDR images. The intrinsic information being latent is another limitation since many applications require explicit intrinsic information like depth and normals.

Nonetheless, there is good evidence that explicit intrinsic information can be extracted from our latent intrinsics. Our method clearly “knows” albedo, and this information can be elicited without examples. Similarly, it “knows” something about luminaires, such as their locations and effects. It is intriguing to speculate that it “knows” other information relevant to relighting, such as depth or surface microstructure. Future work will pursue this line of inquiry and also focus on developing a purely unsupervised framework to infer intrinsic and extrinsic properties from collections of in-the-wild images. This will include refining probing techniques for better extraction of explicit intrinsics and identifying additional intrinsic properties crucial for relighting that do not align with the current taxonomy. We believe this will improve the applicability and robustness of our approach, making it suitable for a wider range of real-world scenarios.

Acknowledgment

AB thanks Stephan R. Richter for the discussions that led to the consideration of intrinsic images as latent variables. This material in part is based upon work supported by the National Science Foundation under Grant No. 2106825, and by a gift from Boeing.

References

- [1] J. T. Barron and J. Malik. Shape, illumination, and reflectance from shading. *PAMI*, 2014.
- [2] J. T. Barron and B. Poole. The fast bilateral solver. In *ECCV*, 2016.
- [3] H. Barrow and J. Tenenbaum. Recovering intrinsic scene characteristics from images. In *ICVS*, 1978.
- [4] J. Beck. *Surface color perception*. Cornell University Press, 1972.
- [5] S. Bell, K. Bala, and N. Snavely. Intrinsic Images in the Wild. In *SIGGRAPH*, 2014.
- [6] A. Bhattad and D. A. Forsyth. Cut-and-paste object insertion by enabling deep image prior for reshading. In *3DV*, 2022.
- [7] A. Bhattad, V. Shah, D. Hoiem, and D. A. Forsyth. Make it so: Steering stylegan for any image inversion and editing. *arXiv:2304.14403*, 2023.
- [8] A. Bhattad, D. McKee, D. Hoiem, and D. Forsyth. Stylegan knows normal, depth, albedo, and more. In *NeurIPS*, 2024.
- [9] A. Bhattad, J. Soole, and D. Forsyth. Stylitgan: Image-based relighting via latent control. In *CVPR*, 2024.
- [10] S. Bi, X. Han, and Y. Yu. An l1 image transform for edge-preserving smoothing and scene-level intrinsic decomposition. In *SIGGRAPH*, 2015.
- [11] S. Bi, N. K. Kalantari, and R. Ramamoorthi. Deep hybrid real and synthetic training for intrinsic decomposition. In *EGSR*, 2018.
- [12] C. Careaga, S. M. H. Miangoleh, and Y. Aksoy. Intrinsic harmonization for illumination-aware image compositing. In *SIGGRAPH Asia*, 2023.
- [13] P. Debevec. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *PACMCGIT*, 1998.
- [14] X. Du, N. Kolkin, G. Shakhnarovich, and A. Bhattad. Generative models: What do they know? do they know things? let’s find out! *arXiv:2311.17137*, 2023.
- [15] Q. Fan, J. Yang, G. Hua, B. Chen, and D. Wipf. Revisiting deep intrinsic image decompositions. In *CVPR*, 2018.
- [16] D. Forsyth and J. J. Rock. Intrinsic image decomposition using paradigms. *PAMI*, 2021.
- [17] M.-A. Gardner, K. Sunkavalli, E. Yumer, X. Shen, E. Gambaretto, C. Gagné, and J.-F. Lalonde. Learning to predict indoor illumination from a single image. In *SIGGRAPH Asia*, 2017.
- [18] M.-A. Gardner, Y. Hold-Geoffroy, K. Sunkavalli, C. Gagné, and J.-F. Lalonde. Deep parametric indoor lighting estimation. In *CVPR*, 2019.
- [19] M. Garon, K. Sunkavalli, S. Hadap, N. Carr, and J.-F. Lalonde. Fast spatially-varying indoor lighting estimation. In *CVPR*, 2019.
- [20] A. Gilchrist. *Seeing Black and White*. Oxford University Press, 2006.
- [21] E. Hering. *Outlines of a theory of the light sense*. 1964. Translated from the German of 1874 by L.M Hurvich and D. Jameson.
- [22] Z. Hu, X. Huang, Y. Li, and Q. Wang. Sa-ae for any-to-any relighting. In *ECCV*, 2020.
- [23] M. Janner, J. Wu, T. D. Kulkarni, I. Yildirim, and J. Tenenbaum. Self-supervised intrinsic image decomposition. In *NeurIPS*, 2017.
- [24] O. F. Kar, T. Yeo, A. Atanov, and A. Zamir. 3d common corruptions and data augmentation. In *CVPR*, 2022.

- [25] T. Karras, M. Aittala, T. Aila, and S. Laine. Elucidating the design space of diffusion-based generative models. In *NeurIPS*, 2022.
- [26] K. Karsch, V. Hedau, D. Forsyth, and D. Hoiem. Rendering synthetic objects into legacy photographs. In *SIGGRAPH Asia*, 2011.
- [27] K. Karsch, K. Sunkavalli, S. Hadap, N. Carr, H. Jin, R. Fonte, M. Sittig, and D. Forsyth. Automatic scene inference for 3d object compositing. *TOG*, 2014.
- [28] P. Kocsis, J. Philip, K. Sunkavalli, M. Nießner, and Y. Hold-Geoffroy. Lightit: Illumination modeling and control for diffusion models. *arXiv:2403.10615*, 2024.
- [29] P. Kocsis, V. Sitzmann, and M. Nießner. Intrinsic image diffusion for single-view material estimation. In *CVPR*, 2024.
- [30] E. Land. Color vision and the natural image: Part i. *PNAS*, 1959.
- [31] E. Land. Color vision and the natural image: Part ii. *PNAS*, 1959.
- [32] B. Li, H. Qin, W. Xiong, Y. Li, S. Feng, W. Hu, and S. Maybank. Ranking-based color constancy with limited training samples. *PAMI*, 2023.
- [33] Z. Li and N. Snavely. Cgintrinsics: Better intrinsic image decomposition through physically-based rendering. In *ECCV*, 2018.
- [34] Z. Li, M. Shafiei, R. Ramamoorthi, K. Sunkavalli, and M. Chandraker. Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image. In *CVPR*, 2020.
- [35] Z. Li, T.-W. Yu, S. Sang, S. Wang, M. Song, Y. Liu, Y.-Y. Yeh, R. Zhu, N. Gundavarapu, J. Shi, et al. Openrooms: An open framework for photorealistic indoor scene datasets. In *CVPR*, 2021.
- [36] Z. Li, J. Shi, S. Bi, R. Zhu, K. Sunkavalli, M. Hašan, Z. Xu, R. Ramamoorthi, and M. Chandraker. Physically-based editing of indoor scene lighting from a single image. In *ECCV*, 2022.
- [37] A. Liu, S. Ginosar, T. Zhou, A. A. Efros, and N. Snavely. Learning to factorize and relight a city. In *ECCV*, 2020.
- [38] L. Murmann, M. Gharbi, M. Aittala, and F. Durand. A dataset of multi-illumination images in the wild. In *CVPR*, 2019.
- [39] T. Nestmeyer and P. V. Gehler. Reflectance adaptive filtering improves intrinsic image estimation. In *CVPR*, 2017.
- [40] T. Nestmeyer, J.-F. Lalonde, I. Matthews, E. Games, A. Lehrmann, and A. Borealis. Learning physics-guided face relighting under directional light. In *CVPR*, 2020.
- [41] J. Philip, M. Gharbi, T. Zhou, A. A. Efros, and G. Drettakis. Multi-view relighting using a geometry-aware network. *TOG*, 2019.
- [42] R. Ramamoorthi and P. Hanrahan. On the relationship between radiance and irradiance: determining the illumination from images of a convex lambertian object. *JOSAA*, 2001.
- [43] M. Roberts, J. Ramapuram, A. Ranjan, A. Kumar, M. A. Bautista, N. Paczan, R. Webb, and J. M. Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *ICCV*, 2021.
- [44] S. Sengupta, B. Curless, I. Kemelmacher-Shlizerman, and S. M. Seitz. A light stage on every desk. In *ICCV*, 2021.
- [45] T. Sun, J. T. Barron, Y.-T. Tsai, Z. Xu, X. Yu, G. Fyffe, C. Rhemann, J. Busch, P. Debevec, and R. Ramamoorthi. Single image portrait relighting. *TOG*, 2019.
- [46] H. von Helmholtz. *Helmholtz' treatise on physiological optics*. 1924-1925. Translated from the 3rd German Edition of 1867, edited by J.P Southall.

- [47] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *TIP*, 2004.
- [48] Z. Wang, J. Phillion, S. Fidler, and J. Kautz. Learning indoor inverse rendering with 3d spatially-varying lighting. In *ICCV*, 2021.
- [49] H. Weber, M. Garon, and J.-F. Lalonde. Editable indoor lighting estimation. In *ECCV*, 2022.
- [50] C. Witzel and K. R. Gegenfurtner. Color perception: Objects, constancy, and categories. *Annu Rev Vis Sci*, 2018.
- [51] H.-H. Yang, W.-T. Chen, and S.-Y. Kuo. S3net: A single stream structure for depth guided image relighting. In *CVPR*, 2021.
- [52] Y. Yu and W. A. Smith. Inverserendernet: Learning single image inverse rendering. In *CVPR*, 2019.
- [53] Y. Yu, K. H. R. Chan, C. You, C. Song, and Y. Ma. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. In *NeurIPS*, 2020.
- [54] Y. Yu, A. Meka, M. Elgharib, H.-P. Seidel, C. Theobalt, and W. A. Smith. Self-supervised outdoor scene relighting. In *ECCV*, 2020.
- [55] H. Zhou, S. Hadap, K. Sunkavalli, and D. W. Jacobs. Deep single-image portrait relighting. In *ICCV*, 2019.

A Experiment Details

Training Details We train our model with a batch size of 256 for 1,000 epochs using the AdamW optimizer, with a constant learning rate of $2e-4$ and a weight decay ratio of $1e-2$. To improve the semantic representation, we corrupt images with Gaussian noise during the first 400 epochs and follow Karras et al. [25] to sample the standard deviation σ with $\ln(\sigma) \sim \mathcal{N}(-1.2, 1.2^2)$. In the later 600 epochs, we turn off the Gaussian noise to focus on enhancing the image quality. We train our model with 4A40 and a complete training requires 40 hours.

Model Details Our autoencoder employs a U-Net architecture, incorporating residual convolutional blocks as the fundamental components. Each block is composed of two convolutional layers, group normalization, and a nonlinear activation function. The structure specifies [1, 2, 2, 4, 4, 4] blocks at each resolution level, starting from a resolution of 256, with the resolution halving after each level. The corresponding configurations for latent channels at these levels are [32, 64, 128, 128, 256, 512].

The intrinsic features, denoted as $S_{s,i}^l$, are gathered from the output of the final block at each resolution level, starting from a resolution of 128x128 down to the bottleneck. For generating extrinsic features L_s^l , multiple MLP layers are applied to the bottleneck features of the encoder, followed by averaging across all spatial features. We limit the channel number of the extrinsic features to 16 to prevent them from conveying high-frequency components.



Figure 9: We visualize more examples for the image relighting task in multi-illumination dataset[38]. **Right:** Zoomed-in view of the chrome ball used as a probe to evaluate detail preservation in the environment map.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We states our contributions in Section. 1.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss our limitations in Section. 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide experimental details in appendix section. A.

Guidelines:

- The answer NA means that the paper does not include experiments.

- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the link to our project page.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide experimental details in appendix section A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We report the averaged results after several runs.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide those details in Section. A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: Our research conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: here is no societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: Our paper does not use existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Our paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.