
Repurposing Marigold for Zero-Shot Metric Depth Estimation via Defocus Blur Cues

Chinmay Talegaonkar*, Nikhil Gandudi Suresh, Zachary Novack, Yash Belhe,
Priyanka Nagasamudra, Nicholas Antipa
University of California San Diego

Abstract

Recent monocular metric depth estimation (MMDE) methods have made notable progress towards zero-shot generalization. However, they still exhibit a significant performance drop on out-of-distribution datasets. We address this limitation by injecting defocus blur cues at inference time into Marigold, a *pre-trained* diffusion model for zero-shot, scale-invariant monocular depth estimation (MDE). Our method effectively turns Marigold into a metric depth predictor in a training-free manner. To incorporate defocus cues, we capture two images with a small and a large aperture from the same viewpoint. To recover metric depth, we then optimize the metric depth scaling parameters and the noise latents of Marigold at inference time using gradients from a loss function based on the defocus-blur image formation model. We compare our method against existing state-of-the-art zero-shot MMDE methods on a self-collected real dataset, showing quantitative and qualitative improvements. Our implementation is available at <https://github.com/chinmay0301ucsd/DiffusionCam>.

1 Introduction

Estimating metric depth from a single camera viewpoint is a central problem in computer vision with numerous downstream applications, including 3D reconstruction [24], autonomous driving [50], and endoscopy [31]. This task, known as *monocular metric depth estimation* (MMDE), is fundamentally ill-posed due to inherent depth-scale ambiguity [46]. Multi-view methods [65] avoid this ambiguity but are often expensive and impractical in settings like endoscopy or microscopy. Training data-driven MMDE methods is challenging, as it requires accounting for a diverse set of camera parameters and metric depth scales. As a result, existing MMDE models struggle in zero-shot settings, i.e., they generalize poorly to unseen datasets. Recent advances in zero-shot MMDE [75] have demonstrated improved generalization, but there is still a considerable performance drop on unseen datasets.

In contrast to MMDE, monocular *relative* depth estimation (MDE) methods recover a relative depth map, factoring out the physical depth scale. This enables using large-scale datasets with diverse depth ranges [41] for training data-driven MDE methods. As a result, MDE methods achieve better zero-shot generalization at significantly lower training cost than MMDE methods, as shown by recent results [70, 71, 26]. However, despite favorable performance on benchmarks, the absence of metric scale in MDE outputs precludes their applicability in downstream tasks requiring absolute depth.

Existing data-driven MMDE methods commonly suffer from two failure modes: undesirable coupling between image texture and depth predictions (fig. 3), and inaccurate estimation of the scene’s physical scale (see fig. 4). The first issue of *texture coupling* also affects MDE methods (see fig. 3) unless explicitly mitigated through complex training procedures that account for texture variation [22, 55].

*Corresponding author: ctalegaonkar@ucsd.edu

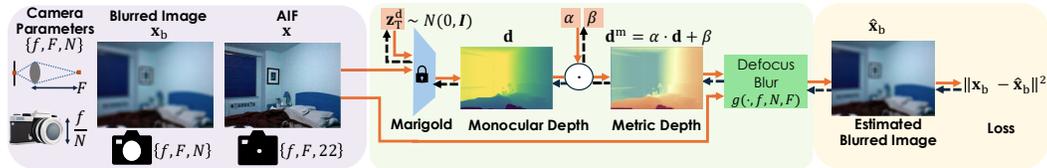


Figure 1: **Method overview.** We capture two images (same viewpoint) from a camera with focal length f and focused at a distance F : an all-in-focus (AIF) image x (F-stop: $N = 22$) and a blurred image x_b (F-stop: $N < 22$). Using the AIF x and an initial *learnable* noise vector $z_T^{(d)}$, Marigold predicts the relative depth d . We then affine transform d with *learnable* parameters (α, β) , obtaining the metric depth d^m . Given the AIF x , depth d^m , and camera parameters (f, F, N) , we synthesize a blurred image \hat{x}_b using the defocus blur forward model. To update the learnable parameters, we compute their gradients w.r.t the L2 loss between \hat{x}_b and x_b .

We demonstrate that incorporating camera physics, particularly defocus blur, with a data-driven MDE model at inference time can effectively address both failure modes *without any re-training*. We use Marigold [26], a diffusion-based MDE model, as it enables backpropagation-based inference time optimization using defocus cues, while staying on its learned manifold of plausible depth maps. To provide these defocus blur cues, we capture two images from a single viewpoint: a small-aperture all-in-focus (AIF) image, which is given as input to Marigold, and a large-aperture image that provides physical cues for metric depth. Our approach requires only a variable aperture camera, such as a DSLR, and knowledge of the lens focal length, focus distance, and F-stop, which are readily available from image metadata; this avoids the need for extrinsic calibration required by multiview setups.

Defocus blur is known to provide coarse metric depth cues in monocular settings [33, 59, 69], but has rarely been explored alongside recent data-driven advances in MMDE. We show in our work that revisiting this classical cue and integrating it with modern relative depth methods like Marigold improves metric depth estimation performance, without any retraining. While most depth from defocus methods use multiple measurements [21, 30, 63], we require only two. Prior work has also explored other hardware novelties to encode depth cues, such as dual pixel cameras [19, 68] and coded apertures [28]. However, these depth cues are coarse, and training MMDE models for these systems is limited by the lack of datasets captured with a sufficient camera parameter diversity. This motivates the need for an approach that combines strong data-driven pre-trained MDE/MMDE models trained on existing RGB datasets with physics-based cues from non-pinhole cameras.

To this end, our **contributions** are as follows: We formulate MMDE as an inverse problem under the defocus blur image formation model. Using inference-time optimization, we recover the metric depth scale and correct for texture-depth coupling – without retraining Marigold. We devise a hardware setup consisting of a rigidly coupled DSLR (RGB camera) and an Intel RealSense depth camera to capture RGB images and the ground truth metric depth. We collect a dataset of 7 diverse real-world indoor scenes, captured at different defocus blur levels. We compare our method on this dataset with current MMDE methods and demonstrate quantitative improvements on real data.

2 Related Work

Zero-shot depth models Recent approaches for Monocular (relative) Depth Estimation (MDE) and Monocular Metric Depth Estimation (MMDE) can be classified as discriminative or generative. Discriminative approaches largely rely on vision transformer-based architectures [7, 5, 6, 74, 39] trained on large-scale datasets. Transformer-based methods have been more successful for MDE [70, 71] compared to MMDE accuracy-wise, as MMDE is a more ill-posed task than MDE. These approaches have very low inference times, but for MMDE, the performance degrades on out-of-distribution test scenes². Most of these methods are purely data-driven and ignore depth-based visual effects in images such as defocus cues. It is also difficult to incorporate physics-based depth refinement in these methods at test time without any re-training [73]. Since we use a generative MDE, our approach can incorporate physical cues at test time in a training-free manner.

²Refer to [75] for a detailed survey on MMDE methods.

Current state-of-the-art *generative* MDE/MMDE methods [47, 13] are predominantly diffusion-based. Several previous methods [13, 47, 32] incorporate diffusion-based depth denoising in their pipelines, achieving highly detailed depth maps – but are not zero-shot. [48] achieves zero-shot MMDE with a diffusion-based approach by incorporating diverse field of view (FOV) augmentations in training, but it is not open source and lags behind transformer-based methods in performance. Marigold [26] is trained by fine-tuning Stable Diffusion-v2 on synthetic depth data. It achieves high-quality zero-shot MDE and supports test time refinement, but is not applicable natively for MMDE. Our approach uses defocus cues to refine the relative depth predictions from Marigold (or similar methods [16]), enabling its application to MMDE. Prior work [60] also proposes a similar strategy for dense MMDE from a sparse metric depth map using Marigold and test time optimization. In contrast, our method does not require a sparse depth map as input, and solely relies on RGB images and a priori known scene bounds. Using defocus blur cues, our method resolves inaccuracies in monocular depth while also estimating the global scaling parameters for metric depth.

Diffusion model priors for inverse problems We frame MMDE as an inverse problem under the defocus blur image formation model. This framing closely relates to the recent work on solving linear inverse problems using *pre-trained* diffusion models [12]. [9, 11, 54, 10] incorporate the forward model constraints while sampling pixel-space diffusion models pre-trained on smaller datasets. As a result, these methods require many steps while sampling the diffusion model and have limited generalizability as priors. [53, 44] use latent variable diffusion models (LDMs) as priors, resulting in better in-the-wild generalizability. We use Marigold as the LDM, but instead of incorporating the defocus forward model during sampling, we optimize the latent noise vector based on the error between the observed and predicted image using the forward model. Our approach is inspired by recent methods [37, 36, 77, 61, 17] which uses noise optimization in conjunction with a differentiable auxiliary guidance loss to improve the sampling quality of the diffusion model based on text input. While these methods use trained models as differentiable proxies for guidance, we use a physics-based imaging forward model. [34] also uses a physics-based forward model with a diffusion prior, but requires re-training the diffusion prior from scratch.

Depth estimation from passive camera physics cues A substantial body of research on MMDE leverages camera physics, including methods like depth-from-defocus (DfD), [64, 20, 58, 2], phase/aperture masks [28, 78, 76, 3, 67], and dual pixel sensors [18, 68, 1]. Classical approaches produce very coarse depth maps. Most DfD methods need a well-aligned [66] multi-image focal stack [29, 57], to achieve good depth quality. Optical mask-based methods pose a harder inverse problem of jointly estimating both AIF and the depth map. We capture only 2 images at the same focus distance and different apertures, requiring no image alignment and AIF estimation. This simplifies depth refinement with minimal added capture time. Previous work has also used variable apertures for classical [15] and learning-based [56] depth estimation methods, but it is not zero-shot. While learning-based approaches [19, 8, 66, 18] help improve the depth map quality for these methods, they don’t generalize well to out-of-distribution scenes. Dual pixel-based methods are popular for phone cameras, [18, 38], but are tied to the specific camera architecture. While we show results with a standard DSLR sensor, our approach can be adapted to dual-pixel-style camera architectures as well.

3 Preliminaries

Diffusion-based monocular depth estimation Our approach is built on top of Marigold [26], which is a monocular depth estimator trained by fine-tuning the denoising U-Net of StableDiffusion-v2 (SDv2) [43] on synthetic depth data. Marigold allows sampling the conditional distribution of the monocular depth given an input image, $p(\mathbf{d}_0|\mathbf{x})$. In particular, Marigold attempts to generate clean monocular depth maps, \mathbf{d}_0 given a clean input image \mathbf{x} by first sampling $\mathbf{d}_T \approx \mathcal{N}(0, \mathbf{I})$ from an i.i.d Gaussian distribution, and then iteratively denoising it (where each intermediate step is denoted by \mathbf{d}_t) according to a fixed noise schedule with parameters α_t, σ_t . As SDv2 is a *latent*-diffusion model, the entire generative process happens on encoded latent depth maps $\mathbf{z}_0^{(d)}$ with latent images $\mathbf{z}^{(x)}$ as conditioning. Training Marigold³ (which we denote as $\hat{\mathbf{x}}_\phi(\cdot)$) involves using the standard denoising loss common in image diffusion works, but on depth maps rather than images, and passes the image

³While technically SDv2 is trained in ϵ -prediction mode rather than \mathbf{x} -prediction, we use \mathbf{x} -prediction given their mathematical equivalence and alignment with Marigold-LCM, which uses \mathbf{x} -prediction.

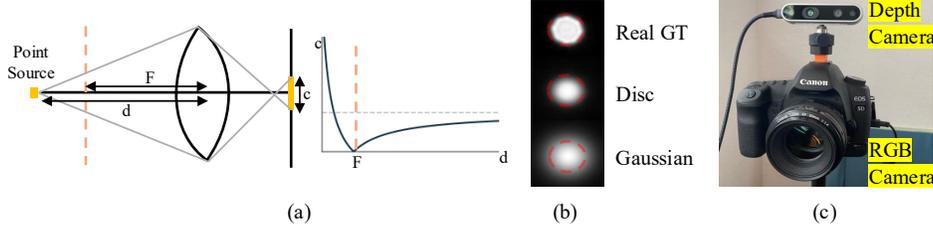


Figure 2: **Comparing simulated PSFs with the PSF captured from our camera setup.** (a) A point source placed d distance away from a thin lens focused at a focus distance F produces a blurred image (PSF) with a diameter c , also known as the circle of confusion. The variation of c with source distance d is shown in the plot. (b) The Disc approximation to the camera PSF lies roughly within the same bounds (dotted red circle) as the PSF captured from the RGB camera (Real GT) in (c). The Gaussian PSF significantly exceeds the bounds. Slight differences between the real and Disc PSF stem from the octagonal aperture and diffraction ignored in our model. (c) We rigidly mount an Intel RealSense on a DSLR to capture ground truth depth, and calibrate both cameras to align predicted depth from the RGB image with the ground truth depth for evaluation.

input into the model as well through channel-wise concatenation:

$$\mathbb{E}_{\mathbf{z}_0^{(d)}, \mathbf{z}^{(x)} \sim \mathcal{D}, t \sim p(t), \epsilon \sim \mathcal{N}(0, \mathbf{I})} [w(t) \|\mathbf{z}_0^{(d)} - \hat{\mathbf{x}}_\phi(\alpha_t \mathbf{z}_0^{(d)} + \sigma_t \epsilon, \mathbf{z}^{(x)}, t)\|_2^2], \quad (1)$$

where $w(t)$ is a time-dependent weighting function. At inference time, \mathbf{x} is first mapped to a lower dimensional latent vector $\mathbf{z}^{(x)} = \mathcal{E}(\mathbf{x})$ through VAE encoder \mathcal{E} in SDv2. The inference process starts with sampling a noisy latent depth vector $\mathbf{z}_T^{(d)} \sim \mathcal{N}(0, \mathbf{I})$, which is iteratively refined by applying the denoiser $\hat{\mathbf{x}}_\phi(\mathbf{z}_i^{(d)}, \mathbf{z}^{(x)}, t)$ to obtain $\mathbf{z}_0^{(d)}$ over T sampling steps. $\mathbf{z}_0^{(d)}$ is then decoded through the SD decoder \mathcal{D} to produce the output depth map $\mathbf{d} = \mathcal{D}(\mathbf{z}_0^{(d)})$. Marigold assumes both $\mathbf{z}^{(x)}, \mathbf{z}^{(d)} \in \mathbb{R}^M$. For faster inference, we use the latent consistency model version of Marigold, Marigold-LCM, and obtain $\mathbf{z}_0^{(d)}$ with a single inference step from $\mathbf{z}_T^{(d)}$. Monocular depth \mathbf{d} can be mapped to metric depth \mathbf{d}^m by an affine transform; more details in section 4.

Modeling defocus cues Under the thin lens camera assumption, defocus blur manifests as a uniform blur kernel, with a depth-dependent diameter. The blur kernel diameter for a point source placed at a distance d away from the camera is determined by the circle of confusion (CoC) [40] equation

$$c(d) = \frac{f^2}{N} \frac{|d - F|}{d(F - f)s}, \quad (2)$$

where the focal length f , focus distance F , and F-stop N are camera parameters that are known from the image EXIF data. s denotes the pixel size in physical units (m). The CoC defines a depth-dependent point spread function (PSF) $h(i, j | u, v, d)$ that predicts the response at pixel coordinate (i, j) from a point source at lateral coordinate (u, v) and depth d under defocus blur. Note that (u, v) can represent depth-normalized pixel coordinates under an ideal pinhole projection, allowing us to write $d = \mathbf{d}^m[u, v]$ (valid for general non-volumetric scenes that assume a single depth value per coordinate). We assume that the PSF is shift-invariant for a given depth d , i.e. $h(i, j | u, v, d) = h(i - u, j - v | 0, 0, \mathbf{d}^m[u, v]) = h(i - u, j - v | \mathbf{d}^m[u, v])$, which is simply the on-axis PSF, modeled as a disc (assuming circular aperture) with radius given by CoC equations, translated to be centered at (i, j) . To ensure smooth optimization, we include a linear fall-off at the boundary of the discontinuous disc kernel, similar to [62]. The PSF can then be expressed as

$$h(i, j | u, v, d) = \frac{\widehat{W}(i, j | u, v, d)}{\sum_{i, j} \widehat{W}(i, j | u, v, d)}, \text{ where} \quad (3)$$

$$\widehat{W}(i, j | u, v, d) = \begin{cases} 1, & m \leq \frac{c(d)-1}{2} \\ \frac{c(d)+1}{2} - m, & \frac{c(d)-1}{2} < m \leq \frac{c(d)+1}{2} \\ 0, & \frac{c(d)+1}{2} > m \end{cases}, \quad (4)$$

$$m = \sqrt{(i - u)^2 + (j - v)^2}. \quad (5)$$

We assume the PSF to be normalized and explicitly account for exposure and energy balancing during image capture and processing (see section 4). While the above PSF can also be approximated as an isotropic 2D Gaussian [20], we opt for the disc parameterization in eq. (5) similar to [62, 51], as it better approximates the PSF of the real camera compared to the Gaussian approximation in [20], as shown in fig. 2. We can see from eq. (2) that an image captured with a very small aperture (high N) would have negligible defocus blur due to very small CoC values. Such an image is referred to as the all-in-focus (AIF) image, \mathbf{x} . Given the AIF \mathbf{x} , the blurred image \mathbf{x}_b can be approximated (neglecting occlusion) as a spatially varying convolution between h and \mathbf{x} ,

$$\mathbf{x}_b(i, j) = \iint \mathbf{x}(u, v) \cdot h(i - u, j - v; \mathbf{d}^m[u, v]) du dv \quad (6)$$

For simplicity, we denote the above image formation forward model as

$$\mathbf{x}_b = g(\mathbf{x}, \mathbf{d}^m, f, F, N). \quad (7)$$

The AIF image \mathbf{x} , captured at a high F-stop, serves as the blur-free input for both Marigold and the camera blur model eq. (7), while the low F-stop image \mathbf{x}_b provides defocus cues for MMDE.

Inference-time optimization Without loss of generality, any generative model (GAN, Diffusion or flow-based) $\phi : \mathbb{R}^M \rightarrow \mathbb{R}^N$ can be construed as a mechanism to map a simple probability distribution, such as an i.i.d. Gaussian distribution, $\varepsilon \sim \mathcal{N}(0, \mathbf{I}) \in \mathbb{R}^M$ to a non-linear N -dimensional manifold, such as images or audio, through a differentiable generative process $x = \phi(\varepsilon)$. Inference time optimization [37] refers to manipulating the generation process by updating the *initial noise* ε based on gradients from a differentiable loss function $\mathcal{L}(x)$ on the generated sample x ,

$$\varepsilon \rightarrow \varepsilon - \nabla_{\varepsilon} \mathcal{L}(x). \quad (8)$$

To further ensure that ε still lies close to the Gaussian manifold after the gradient updates, ε can be rescaled to have a L2 norm of \sqrt{M} as in [45], which is a valid approximation for samples drawn from a high dimensional Gaussian distribution as per the Gaussian annulus theorem [4]. This holds for most generative models, as the initial noise vectors are typically high-dimensional ($M > 50$). In our case, ε corresponds to the noise latent $\mathbf{z}_T^{(d)}$ in Marigold, which we optimize using a loss function (eq. (11)) governed by the defocus blur forward model eq. (7).

4 Method

We capture two images per scene: \mathbf{x} , with F-stop ($N_{\text{aif}} = 22$) and exposure time t_{aif} , serves as the blur-free all-in-focus (AIF) image. A second image, \mathbf{x}_b , is captured at a lower F-stop ($N_b = 8$) with exposure time t_b , thereby providing strong depth-varying defocus cues. The forward model in eq. (7) assumes radiometrically linear images (no gamma correction or non-linear processing) and energy constancy between the AIF and blurred images. We use raw images to satisfy these assumptions. Since total captured energy scales with exposure time (t) and aperture area ($\propto (f/N)^2$) [25], we scale \mathbf{x}_b by the factor $\frac{t_{\text{aif}}}{t_b} \cdot \frac{N_b^2}{N_{\text{aif}}^2}$ to match the energy in \mathbf{x} . Note that we vary the exposure time to ensure well-exposed measurements across F-stop settings, while fixing the camera gain.

We frame metric depth estimation as an inverse problem, with the defocus blur image formation process in eq. (6) as the forward model, and Marigold as the monocular depth prior. To obtain scale-invariant monocular depth $\mathbf{d} \in [0, 1]$, we use Marigold-LCM, which takes in as input the AIF \mathbf{x} (encoded to $\mathbf{z}^{(x)}$), and a learnable depth latent vector $\mathbf{z}_T^{(d)} \sim \mathcal{N}(0, \mathbf{I})$. A single inference step of Marigold-LCM gives us the denoised depth latent $\mathbf{z}_0^{(d)} = \hat{\mathbf{x}}_{\phi}(\mathbf{z}_T^{(d)}, \mathbf{z}^{(x)}, 1)$, which is decoded to monocular depth $\mathbf{d} = \mathcal{D}(\mathbf{z}_0^{(d)})$. The predicted monocular depth $\mathbf{d} \in [0, 1]$ is then mapped to metric depth by affine transforming \mathbf{d} with a learnable metric scale (α) and offset (β) per scene, $\mathbf{d}^m = \alpha \cdot \mathbf{d} + \beta$. To ensure that α, β remain bounded and differentiable, we parameterize them as $\alpha = s_{\text{max}} \cdot \sigma(a)$ and $\beta = s_{\text{min}} \cdot \sigma(b)$, where $\sigma(\cdot)$ is the sigmoid function; a, b are unconstrained learnable parameters initialized to 0, and s_{max} and s_{min} are the upper and lower scene depth bounds, respectively, which we assume are known a priori (valid for indoor scenes). To summarize, the metric

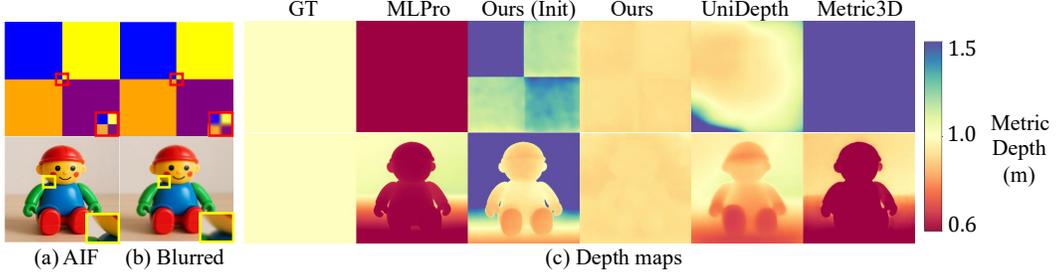


Figure 3: **Correcting texture-depth coupling.** We assess MMDE performance on textured fronto-parallel 2D planes with constant ground truth depths (GT). Using an all-in-focus (a) and blurred image (zoom into insets) (b), our method (RMSE: 0.01) recovers the correct depth maps (c) for the two textured planes. We resolve the texture coupling in the Marigold prediction (Ours Init) and recover the correct metric scale. Competing methods (RMSE: 0.2-0.5) fail to predict both the constant relative depth map (except MLPro and Metric3D in row 1) and the correct scale.

depth (\mathbf{d}^m) can be expressed using the optimizable parameters $a, b, \mathbf{z}_T^{(d)}$ as:

$$\begin{aligned} \mathbf{d}^m &= s_{\max} \cdot \sigma(a) \cdot \mathcal{D} \left(\hat{\mathbf{x}}_\phi \left(\mathbf{z}_T^{(d)}, \mathbf{z}^{(x)}, 1 \right) \right) + s_{\min} \cdot \sigma(b) \\ &:= y \left(a, b, \mathbf{z}_T^{(d)} \right). \end{aligned} \quad (9)$$

The optimized metric depth $\hat{\mathbf{d}}^m$ can then be recovered by solving:

$$\hat{\mathbf{d}}^m = \arg \min_{\mathbf{d}^m = y(a, b, \mathbf{z}_T^{(d)})} \|\mathbf{x}_b - g(\mathbf{x}, \mathbf{d}^m, f, F, N)\|_2^2 \quad (10)$$

$$\text{subject to } \|\mathbf{z}_T^{(d)}\|_2 = \sqrt{M}, \quad (11)$$

where $g(\cdot)$ denotes the defocus blur forward model (eq. (7)), \mathbf{x}_b and \mathbf{x} are the captured blurred and AIF images, respectively. By optimizing the learnable parameters $a, b, \mathbf{z}_T^{(d)}$, we incorporate defocus blur cues for both correct metric scale recovery (a, b) and refining the initial depth estimate by Marigold ($\mathbf{z}_T^{(d)}$). See fig. 1 for the overview of our method.

Motivating Synthetic Toy Examples We demonstrate that our approach resolves texture-depth coupling, accurately recovering metric depth in a synthetic scene with a textured plane at constant depth. While such a plane may seem simple, distinguishing it from a flat 2D image/poster or an actual 3D scene is challenging when viewed from a single viewpoint. Data-driven methods are biased towards predicting depths that reflect surface variations even in the absence of true depth changes, i.e., their outputs are strongly *texture coupled*. However, supplementing the AIF image with a simulated blurred image provides defocus cues that help our method disambiguate a flat poster from a 3D scene, as demonstrated in fig. 3 using toy examples of textured planes with constant depth. While learning-based methods and initial Marigold outputs suffer from texture coupling and scale errors, our method corrects both, producing accurate, constant-depth maps that outperform all baselines.

Accelerating inference We use the distilled latent consistency model version of Marigold (Marigold-LCM) to reduce the number of sampling steps significantly. We observe that a single sampling step suffices for our case, which significantly speeds up inference-time optimization [36, 14] relative to the normal 20-50 inference steps [37, 61] that Marigold [60] uses. We show an ablation study with more sampling steps in supplement.E. We also implement custom CUDA kernels for the Disc-PSF forward model. This provides a 2.5x speed up over the PyTorch implementation provided by [62] while being more memory efficient, allowing our method to scale to higher-resolution images. Using a single sampling step of Marigold-LCM allows us to compute gradients w.r.t $\mathbf{z}_T^{(d)}$ without gradient checkpointing as previously done in [37].

Optimization details We run the optimization for 200 iterations, which takes roughly 3.5-4 minutes on an NVIDIA A-40 GPU with peak memory usage of 15 GB. We use the Adam optimizer with a

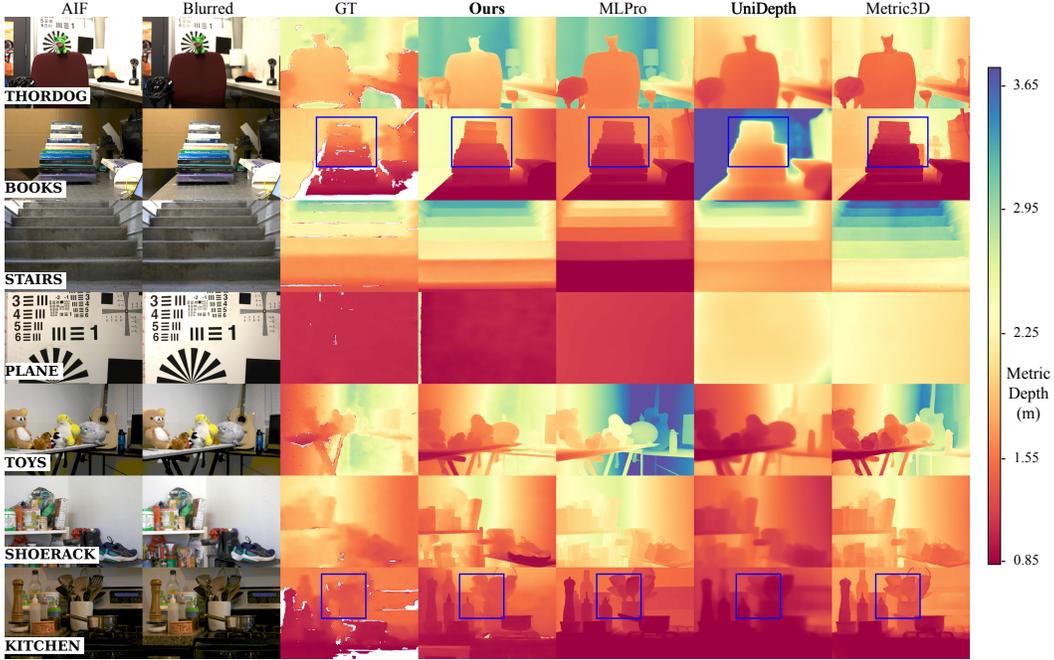


Figure 4: **Comparisons on our collected dataset.** Our method consistently estimates accurate metric depth across all the scenes. We also observe better relative depth recovery due to leveraging defocus cues (zoom in 4x on blurred) in some regions (blue boxes, STAIRS). While the competing methods perform comparably to ours in some cases (MLPro:PLANE, KITCHEN, UniDepth:STAIRS, Metric3D:THORDOG, BOOKS), they struggle with the rest of the scenes due to incorrect relative depth (TOYS, BOOKS) and metric scale (PLANE) recovery. recovers sharp details but fails at metric scale and relative depth accuracy for many of the scenes. Since the RealSense has a wider FOV than the DSLR, we show a roughly aligned crop of the GT depth for comparison.

learning rate of 1.5×10^{-3} for $\mathbf{z}_T^{(d)}$, 5×10^{-3} for (a, b) , and default values for optimizer parameters. Note that we use the *same* scene bounds $s_{\min} = 1.49$, $s_{\max} = 3.5$ for all the real scenes in our dataset. These values represent a conservative upper bound on the potential maximum scale and offset in the real dataset. Please see supplement.A for more details.

5 Experiments and Results

Dataset details Our method requires 2 images captured with different apertures (but same view-point) to integrate defocus cues. Standard monocular depth datasets [49, 35] typically capture in-focus images at a single aperture per scene, making them unsuitable for evaluating our method. While one could simulate defocused images with eq. (7) and RGBD data, this does not capture the model mismatch (occlusion, diffraction) in the physical image formation process. Therefore, to fairly evaluate our method, we construct a hardware capture setup, shown in fig. 2, comprising an Intel RealSense depth camera rigidly mounted to a DSLR camera (Canon EOS 5D Mark II). We use this system to collect a custom real-world dataset of 7 unique scenes, evaluating our method against learning-based MMDE baselines. We choose scenes with diverse subjects and depth profiles, placed within the operating depth range of RealSense (0.3–3.8m) to ensure accurate ground truth depth. Note that the CoC changes negligibly with depth beyond these distances, making defocus cues unreliable. For each scene, we capture images at 6 different apertures, $f/4$, $f/8$, $f/11$, $f/13$, $f/16$, and $f/22$, with the latter serving as the AIF image \mathbf{x} . The blurred image, \mathbf{x}_b , is selected from the lower F-stop images (see fig. 5 for comparison of F-stop setting on depth map quality). As described in section 4, we try to maintain a similar ratio of the exposure time and the camera aperture area for all the measurements taken for a scene. The lens focal length (f) and F-stop (N) are provided by the camera EXIF data, and the focus distance F is read manually from the lens’s analog focus scale⁴.

⁴Alternatively, a lens with focus motor encoding would allow this to be known from EXIF

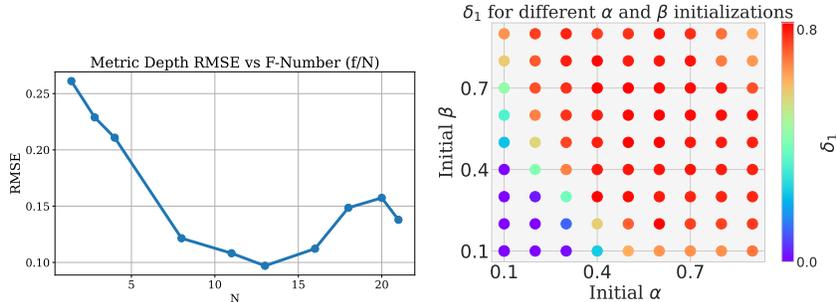


Figure 5: **Analyzing the effect of different aperture sizes and initializations.** *Left:* We use our forward model to simulate blurred images of a scene from the NYU-v2 dataset. We observe minimum depth error at $N = 13$, with errors increasing at more extreme aperture values. *Right:* We plot δ_1 at the end of optimization for various α, β initializations normalized between 0-1 ($\alpha \rightarrow \frac{\alpha}{s_{\max}}, \beta \rightarrow \frac{\beta}{s_{\min}}$), in the figure for visualization. While the performance degrades for small values of α, β (bottom left), it is relatively stable for a broad range of initializations.

Please see supplement.B for camera parameters and other dataset details, and supplement.H for more discussion on ground truth depth map quality.

Evaluation Metrics We evaluate the predicted metric depth, $\hat{\mathbf{d}}^m$, from our method against the RealSense ground truth depth \mathbf{d} (with overloaded notation), using the metrics in [5, 39]. Specifically, we compute absolute relative error (REL) = $\frac{1}{M} \sum_{i=1}^M \frac{|\mathbf{d}_i - \hat{\mathbf{d}}_i^m|}{\mathbf{d}_i}$, root mean squared error (RMSE) = $\sqrt{\frac{1}{M} \sum_{i=1}^M |\mathbf{d}_i - \hat{\mathbf{d}}_i^m|^2}$, average log error (log10) = $\frac{1}{M} \sum_{i=1}^M |\log_{10} \mathbf{d}_i - \log_{10} \hat{\mathbf{d}}_i^m|$, and accuracy thresholds δ_n = fraction of pixels where $\max\left(\frac{\mathbf{d}_i}{\hat{\mathbf{d}}_i^m}, \frac{\hat{\mathbf{d}}_i^m}{\mathbf{d}_i}\right) < (1.25)^n$ for $n = 1, 2, 3$. M denotes the number of pixels, and $\mathbf{d}_i, \hat{\mathbf{d}}_i^m$ denote the RealSense and predicted depths at pixel i , respectively. We also report point-cloud-based metrics, i.e., Chamfer Distance (CD) and the aggregated F1 score (FA) as defined in [39]. We align the ground truth and predicted depth map, similar to [19], and compute the metrics for pixels with non-zero values across the aligned depth maps. See supplement.C for more details on the depth map alignment/calibration procedure.

Quantitative and qualitative results We evaluate our method against some of the recently popular MMDE methods UniDepth [39], Metric3D [74], and MLPro [7] on our collected dataset (table 2) and the NYUv2 [35] test set. Both Metric3D and UniDepth are provided with the required camera parameters as input. We outperform existing methods qualitatively (fig. 4) and quantitatively (table 2) on all the evaluation metrics, averaged over all 7 scenes in our collected dataset. Please see supplement.D for per-scene quantitative metrics on our dataset and supplement.H for depth error visualizations. On the NYUv2 test set, our method is on par with the MMDE baselines. See supplement.J for more details. For our dataset, the MMDE methods (MLPro, Metric3D) recover sharper details and are on-par with our method on some scenes (see fig. 4), but they lack consistency in their overall performance across all scenes. Our method achieves better consistency across varying scene conditions. Leveraging defocus cues enables our approach to recover the correct depth scale while also resolving relative depth errors in some cases (fig. 4 insets). We also evaluate our method (row 4 in table 2) using a Gaussian PSF [20] in the forward model. While previous work [20] uses the Gaussian PSF for training models for unsupervised depth recovery, we find that the model mismatch between the Gaussian PSF and the PSF of the real camera (fig. 2) leads to severe performance degradation in our case, compared to using the Disc PSF, which matches the real camera PSF better. This highlights the value of a physically consistent forward model, even with strong learned priors.

Comparing with fine-tuning based methods Our method relies on optimizing noise-latents of a relative depth foundation model (RDFM) such as Marigold or Geowizard at test-time. Our approach (inference-time optimization) can be valuable in zero-shot settings, where fine-tuning-based methods would struggle due to the domain gap. We validate this by comparing our method with MMDE

Method	RMSE ↓	REL ↓	log10 ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑	CD ↓	FA ↑
[71] FT Hypersim	0.523	0.314	0.151	0.178	0.825	0.971	0.208	0.616
[71] FT NYUv2	0.407	0.252	0.096	0.654	0.832	0.988	0.183	0.733
Marigold FT NYUv2	0.475	0.231	0.098	0.634	0.878	0.963	0.175	0.808
Ours, No FT	0.273	0.125	0.052	0.879	0.975	0.991	0.103	0.870

Table 1: **Comparison with fine-tuning-based (FT) methods on our dataset.** Our method (row 4) outperforms (FT)-based methods across all metrics. NYUv2 (real, depth $\leq 10\text{m}$) has a lesser domain gap to our dataset (real, depth $\leq 3.5\text{m}$) than HyperSim (synthetic, depth $\leq 20\text{m}$); as evident by the improvement in DepthAnythingv2 [71] when FT on NYUv2 (row 2) compared to HyperSim (row 1).

models obtained by fine-tuning RDFMs (DepthAnythingv2 [71], Marigold) on Metric depth datasets with a restricted depth range, as required by our method (section 4). We evaluate the following fine-tuning-based baselines on our collected dataset in table 1 – DepthAnythingv2 [71] fine-tuned on the Hypersim[42] dataset (row 1), NYU-v2 train set (row 2), and Marigold Fine-Tuned on NYUv2 (row 3). Our method (row 4, table 1) outperforms the fine-tuning-based baselines. This validates the utility of our inference-time method in zero-shot settings, where fine-tuning-based methods can struggle due to a domain gap. Please see supplement.K for details on fine-tuning these baselines and additional results on the NYUv2 test set. In addition to fine-tuning-based methods, we also outperform [56], which estimates metric depth from multi-aperture inputs via self-supervised learning with a differentiable forward model. See Supplement.L for more details on the comparison with [56].

Extending our method beyond Marigold Our method is plug-and-play and, in principle, works with any diffusion-based model like Marigold that exposes a differentiable mapping from its latent space to depth estimates. This enables test-time latent optimization, which allows for incorporating defocus cues without retraining. We demonstrate this by using Geowizard as the diffusion backbone [16] instead of Marigold. GeoWizard is a stable-diffusion-based monocular depth + normal predictor, trained on a more complex data distribution. Our method achieves similar performance with either of the backbones (Rows 5,6 in table 2), but shows minor patch-level artifacts with GeoWizard, likely due to stronger texture-depth coupling. See Supplement.I for visualizations and hyperparameter details.

Method	RMSE ↓	REL ↓	log10 ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑	CD ↓	FA ↑
MLPro	0.468	0.246	0.105	0.597	0.821	0.990	0.205	0.696
UniDepth	0.574	0.358	0.152	0.263	0.757	0.902	0.260	0.612
Metric3D	0.349	0.195	0.087	0.611	0.958	0.983	0.135	0.814
Ours - Gaussian	0.528	0.279	0.142	0.422	0.695	0.928	0.241	0.652
Ours - Disc	0.273	0.125	0.052	0.879	0.975	0.991	0.103	0.870
Ours (Disc) with GeoWizard	0.291	0.137	0.061	0.824	0.966	0.990	0.105	0.874

Table 2: **Comparison with learning based MMDE methods on our dataset.** Our method with the Disc PSF outperforms all the MMDE baselines averaged over all scenes in our dataset. The disc PSF, being more consistent with the real camera PSF, outperforms the Gaussian PSF. UniDepth, Metric3D, and our method are provided with camera intrinsics parameters during inference.

5.1 Ablation Studies

Only single blurred image as input: In our method, Marigold takes the all-in-focus (AIF) image (F/22) as input to predict relative depth, which is combined with learnable scale–offset parameters to synthesize a blurred image. The synthesized image is compared with a captured blurred image (F/8) that provides metric depth cues [52]. To assess the contribution of the AIF, we remove the AIF and instead provide a single moderately blurred image (F/16) as input to both Marigold and the loss function (eq. (11)). This ablation also evaluates whether the diffusion prior alone is strong enough to operate with a single modestly blurred image. Removing the AIF input leads to a large RMSE increase (1.36 vs. 0.346) and visible artifacts (fig. 6), indicating that while the diffusion prior captures coarse depth cues, the AIF input is essential for accurate metric depth estimation.

Sensitivity to α, β initialization: Gradient-based methods are known to be susceptible to local minima for non-convex optimizations. We thus evaluate the sensitivity of our method to the initialization for α, β . We run the optimization in (eq. (11)) with a fixed $\mathbf{z}_T^{(d)}$ while grid-searching over initial

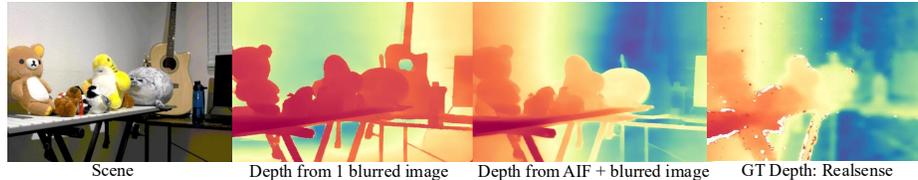


Figure 6: **Degradation in depth quality on using only a single blurred image.** For the TOYS scene, we observe that using a single blurred image as input results in severely inaccurate relative depth (middle), with all toys, guitar, and the monitor at similar relative depths. Our proposed method (right) recovers the depth ordering between the objects more accurately.

values of α and β . In fig. 5 (right), we observe that performance (measured by δ_1) drops for small initialization values but remains stable across a broad range around $\alpha = 0.5, \beta = 0.5$, supporting the robustness of our chosen initialization. We also ablate on sensitivity to initial $\mathbf{z}_T^{(d)}$ in supplement.E.

Improvements in relative depth from defocus cues We quantitatively evaluate the effect of defocus cues in our method on improving the relative depth quality. Optimizing only α, β while holding $\mathbf{z}_T^{(d)}$ constant leads to a performance drop (table 3). This highlights the role of defocus cues in refining the relative depth initially predicted by Marigold. Please see supplement.F for visualizations.

Method	RMSE ↓	REL ↓	log10 ↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Ours α, β opt	0.297	0.156	0.069	0.743	0.957	0.99
Ours	0.273	0.125	0.052	0.879	0.975	0.991

Table 3: **Relative depth quality** Optimizing the noise latent along with the affine parameters (ours) performs better than optimizing only the affine parameters (α, β opt).

Different aperture sizes We analyze how the aperture (F-stop) used for capturing the blurred image affects our performance. To do this, we use a scene from the NYU-v2 [35], an indoor RGBD dataset with high-quality ground truth depth annotations. This synthetic setup allows evaluating large F-stops that cannot be captured with our camera, while isolating aperture size from forward model mismatches in a real setup. Using the ground truth depth and our forward model (eq. (7)), we simulate the blurred images \mathbf{x}_b at varying F-stops (N values) and compute the error metrics between the ground truth and our predicted depth. We observe in fig. 5 that the performance (measured in RMSE) degrades for extreme aperture sizes. This is expected, as extreme blur (high or low) makes the inverse problem ill-posed, and an optimal blur level is key for accurate depth recovery. While $N = 13$ appears to be optimal in simulation (ignoring model mismatch), it underperforms $N = 8$ on average for real scenes, likely due to low contrast and insufficient blur cues in some of the scenes (STAIRS, PLANE). Please see supplement.G for results across all apertures captured in the real dataset.

6 Limitations and Future Work

While our method outperforms data-driven MMDE baselines, it remains significantly slower at inference time. Our method is best suited to scenes with a small depth range for which defocus blur offers high depth sensitivity. We observe that if the initial Marigold prediction is severely incorrect in some regions (visualized in supplement.F), the optimization may not always be able to fully correct them (ours for SHOERACK, THORDOG in fig. 4). A possible extension of our method is to jointly estimate the AIF and depth map from a single blurred input, as previously explored in non-zero-shot approaches [19, 1]. Our framework can broaden the utility of pre-trained depth priors to scientific applications involving depth-dependent imaging processes such as hyperspectral imaging [27], endoscopy [31], and microscopy [72]. While we avoid discretizing the depth map [23], our forward model loses accuracy at occlusion boundaries. We envision further improvements through better PSF engineering (coded aperture masks) and more accurate forward modeling of defocus blur. Another promising direction is to adapt feed-forward methods such as [71] for handling multi-aperture inputs through test-time adaptation and fine-tuning [73]. Large-scale multi-aperture datasets could make such methods practical for these settings.

Acknowledgements: We thank Alankar Kotwal and Bhargav Ghanekar for helpful discussions, Namrata Mantri for helping collect the real dataset, and Agastya Kalra for proofreading the draft.

References

- [1] Abdullah Abuolaim and Michael S Brown. Defocus deblurring using dual-pixel data. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 111–126. Springer, 2020.
- [2] Emma Alexander, Qi Guo, Sanjeev Koppal, Steven Gortler, and Todd Zickler. Focal flow: Measuring distance and velocity with defocus and differential motion. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, pages 667–682. Springer, 2016.
- [3] Nick Antipa, Grace Kuo, Reinhard Heckel, Ben Mildenhall, Emrah Bostan, Ren Ng, and Laura Waller. Diffusercam: lensless single-exposure 3d imaging. *Optica*, 5(1):1–9, 2017.
- [4] Imre Bárány, Van Vu, et al. Central limit theorems for gaussian polytopes. *The Annals of Probability*, 35(4):1593–1621, 2007.
- [5] Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023.
- [6] Reiner Birkl, Diana Wofk, and Matthias Müller. Midas v3. 1—a model zoo for robust monocular relative depth estimation. *arXiv preprint arXiv:2307.14460*, 2023.
- [7] Aleksei Bochkovskii, Amaël Delaunoy, Hugo Germain, Marcel Santos, Yichao Zhou, Stephan R Richter, and Vladlen Koltun. Depth pro: Sharp monocular metric depth in less than a second. *arXiv preprint arXiv:2410.02073*, 2024.
- [8] Marcela Carvalho, Bertrand Le Saux, Pauline Trouvé-Peloux, Andrés Almansa, and Frédéric Champagnat. Deep depth from defocus: how can defocus blur improve 3d estimation using dense neural networks? In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018.
- [9] Hyungjin Chung, Jeongsol Kim, Michael T. McCann, Marc L. Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *International Conference on Learning Representations (ICLR)*, 2023.
- [10] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [11] Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12403–12412, 2022.
- [12] Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milanfar, Alexandros G Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse problems. *arXiv preprint arXiv:2410.00083*, 2024.
- [13] Yiquan Duan, Xianda Guo, and Zheng Zhu. Diffusiondepth: Diffusion denoising approach for monocular depth estimation. In *European Conference on Computer Vision*, pages 432–449. Springer, 2024.
- [14] Luca Eyring, Shyamgopal Karthik, Karsten Roth, Alexey Dosovitskiy, and Zeynep Akata. Reno: Enhancing one-step text-to-image models through reward-based noise optimization. *Advances in Neural Information Processing Systems*, 37:125487–125519, 2024.
- [15] Hany Farid and Eero P Simoncelli. Range estimation by optical differentiation. *Journal of the Optical Society of America A*, 15(7):1777–1786, 1998.
- [16] Xiao Fu, Wei Yin, Mu Hu, Kaixuan Wang, Yuexin Ma, Ping Tan, Shaojie Shen, Dahua Lin, and Xiaoxiao Long. Geowizard: Unleashing the diffusion priors for 3d geometry estimation from a single image. In *European Conference on Computer Vision*, pages 241–258. Springer, 2024.
- [17] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H. Bermano, Tal Hassner, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022.
- [18] Rahul Garg, Neal Wadhwa, Sameer Ansari, and Jonathan T Barron. Learning single camera depth estimation using dual-pixels. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7628–7637, 2019.

- [19] Bhargav Ghanekar, Salman Siddique Khan, Pranav Sharma, Shreyas Singh, Vivek Boominathan, Kaushik Mitra, and Ashok Veeraraghavan. Passive snapshot coded aperture dual-pixel rgb-d imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25348–25357, 2024.
- [20] Shir Gur and Lior Wolf. Single image depth estimation trained via depth from defocus cues. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7683–7692, 2019.
- [21] Caner Hazirbas, Sebastian Georg Soyer, Maximilian Christian Staab, Laura Leal-Taixé, and Daniel Cremers. Deep depth from focus. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*, pages 525–541. Springer, 2019.
- [22] Jing He, Haodong Li, Wei Yin, Yixun Liang, Leheng Li, Kaiqiang Zhou, Hongbo Zhang, Bingbing Liu, and Ying-Cong Chen. Lotus: Diffusion-based visual foundation model for high-quality dense prediction. *arXiv preprint arXiv:2409.18124*, 2024.
- [23] Hayato Ikoma, Cindy M Nguyen, Christopher A Metzler, Yifan Peng, and Gordon Wetzstein. Depth from defocus with learned optics for imaging and occlusion-aware depth estimation. In *2021 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12. IEEE, 2021.
- [24] Kaiwen Jiang, Yang Fu, Mukund Varma T, Yash Belhe, Xiaolong Wang, Hao Su, and Ravi Ramamoorthi. A construct-optimize approach to sparse view synthesis without camera pose. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024.
- [25] Loyd A Jones, Emery Huse, and Vincent C Hall. On the relation between time and intensity in photographic exposure. *Journal of the Optical Society of America*, 12(4):321–348, 1926.
- [26] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024.
- [27] Esther Kho, Lisanne L de Boer, Anouk L Post, Koen K Van de Vijver, Katarzyna Józwiak, Henricus JCM Sterenberg, and Theo JM Ruers. Imaging depth variations in hyperspectral imaging: development of a method to detect tumor up to the required tumor-free margin width. *Journal of biophotonics*, 12(11):e201900086, 2019.
- [28] Anat Levin, Rob Fergus, Frédo Durand, and William T Freeman. Image and depth from a conventional camera with a coded aperture. *ACM transactions on graphics (TOG)*, 26(3):70–es, 2007.
- [29] Haiting Lin, Can Chen, Sing Bing Kang, and Jingyi Yu. Depth recovery from light field using focal stack symmetry. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3451–3459, 2015.
- [30] Chao Liu, Srinivasa G Narasimhan, and Artur W Dubrawski. Matting and depth recovery of thin structures using a focal stack. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6970–6978, 2017.
- [31] Yuying Liu and Siyang Zuo. Self-supervised monocular depth estimation for gastrointestinal endoscopy. *Computer Methods and Programs in Biomedicine*, 238:107619, 2023.
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [33] Maxim Maximov, Kevin Galim, and Laura Leal-Taixé. Focus on defocus: bridging the synthetic to real domain gap for depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1071–1080, 2020.
- [34] Opher Bar Nathan, Deborah Levy, Tali Treibitz, and Dan Rosenbaum. Osmosis: Rgb-d diffusion prior for underwater image restoration. In *European Conference on Computer Vision*, pages 302–319. Springer, 2024.
- [35] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012.
- [36] Zachary Novack, Julian McAuley, Taylor Berg-Kirkpatrick, and Nicholas J. Bryan. DITTO-2: Distilled diffusion inference-time t-optimization for music generation. In *International Society of Music Information Retrieval (ISMIR)*, 2024.

- [37] Zachary Novack, Julian McAuley, Taylor Berg-Kirkpatrick, and Nicholas J. Bryan. DITTO: Diffusion inference-time t-optimization for music generation. In *International Conference on Machine Learning (ICML)*, 2024.
- [38] Liyuan Pan, Shah Chowdhury, Richard Hartley, Miaomiao Liu, Hongguang Zhang, and Hongdong Li. Dual pixel exploration: Simultaneous depth estimation and image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2021.
- [39] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10106–10116, 2024.
- [40] Michael Potmesil and Indranil Chakravarty. A lens and aperture camera model for synthetic image generation. In *Proceedings of the 8th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '81*, page 297–305, New York, NY, USA, 1981. Association for Computing Machinery.
- [41] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3):1623–1637, 2020.
- [42] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10912–10922, 2021.
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [44] Litu Rout, Negin Raof, Giannis Daras, Constantine Caramanis, Alex Dimakis, and Sanjay Shakkottai. Solving linear inverse problems provably via posterior sampling with latent diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- [45] Dvir Samuel, Rami Ben-Ari, Nir Darshan, Haggai Maron, and Gal Chechik. Norm-guided latent space exploration for text-to-image generation. *Advances in Neural Information Processing Systems*, 36, 2024.
- [46] Ashutosh Saxena, Sung Chung, and Andrew Ng. Learning depth from single monocular images. *Advances in neural information processing systems*, 18, 2005.
- [47] Saurabh Saxena, Charles Herrmann, Junhwa Hur, Abhishek Kar, Mohammad Norouzi, Deqing Sun, and David J Fleet. The surprising effectiveness of diffusion models for optical flow and monocular depth estimation. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 39443–39469. Curran Associates, Inc., 2023.
- [48] Saurabh Saxena, Junhwa Hur, Charles Herrmann, Deqing Sun, and David J Fleet. Zero-shot metric depth with a field-of-view conditioned diffusion model. *arXiv preprint arXiv:2312.13252*, 2023.
- [49] Daniel Scharstein and Richard Szeliski. High-accuracy stereo depth maps using structured light. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 195–202. IEEE, 2003.
- [50] Markus Schön, Michael Buchholz, and Klaus Dietmayer. Mgnnet: Monocular geometric scene understanding for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15804–15815, 2021.
- [51] Yichen Sheng, Zixun Yu, Lu Ling, Zhiwen Cao, Xuaner Zhang, Xin Lu, Ke Xian, Haiting Lin, and Bedrich Benes. Dr. bokeh: Differentiable occlusion-aware bokeh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4515–4525, 2024.
- [52] Haozhe Si, Bin Zhao, Dong Wang, Yunpeng Gao, Mulin Chen, Zhigang Wang, and Xuelong Li. Fully self-supervised depth estimation from defocus clue. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9140–9149, 2023.
- [53] Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Lijue Shen. Solving inverse problems with latent diffusion models via hard data consistency. *arXiv preprint arXiv:2307.08123*, 2023.
- [54] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*, 2023.

- [55] Ziyang Song, Zerong Wang, Bo Li, Hao Zhang, Ruijie Zhu, Li Liu, Peng-Tao Jiang, and Tianzhu Zhang. Depthmaster: Taming diffusion models for monocular depth estimation. *arXiv preprint arXiv:2501.02576*, 2025.
- [56] Pratul P Srinivasan, Rahul Garg, Neal Wadhwa, Ren Ng, and Jonathan T Barron. Aperture supervision for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6393–6401, 2018.
- [57] Michael Strecke, Anna Alperovich, and Bastian Goldluecke. Accurate depth and normal maps from occlusion-aware focal stack symmetry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2814–2822, 2017.
- [58] Murali Subbarao and Gopal Surya. Depth from defocus: A spatial domain approach. *International Journal of computer vision*, 13(3):271–294, 1994.
- [59] Huixuan Tang, Scott Cohen, Brian Price, Stephen Schiller, and Kiriakos N Kutulakos. Depth from defocus in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2740–2748, 2017.
- [60] Massimiliano Viola, Kevin Qu, Nando Metzger, Bingxin Ke, Alexander Becker, Konrad Schindler, and Anton Obukhov. Marigold-dc: Zero-shot monocular depth completion with guided diffusion. *arXiv preprint arXiv:2412.13389*, 2024.
- [61] Bram Wallace, Akash Gokul, Stefano Ermon, and Nikhil Vijay Naik. End-to-end diffusion latent optimization improves classifier guidance. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7246–7256, 2023.
- [62] Chao Wang, Krzysztof Wolski, Xingang Pan, Thomas Leimkühler, Bin Chen, Christian Theobalt, Karol Myszkowski, Hans-Peter Seidel, and Ana Serrano. An implicit neural representation for the image stack: Depth, all in focus, and high dynamic range. Technical report, 2023.
- [63] Ning-Hsu Wang, Ren Wang, Yu-Lun Liu, Yu-Hao Huang, Yu-Lin Chang, Chia-Ping Chen, and Kevin Jou. Bridging unsupervised and supervised depth from focus via all-in-focus supervision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12621–12631, 2021.
- [64] Masahiro Watanabe and Shree K Nayar. Rational filters for passive depth from defocus. *International Journal of Computer Vision*, 27:203–225, 1998.
- [65] Bowen Wen, Matthew Trepte, Joseph Aribido, Jan Kautz, Orazio Gallo, and Stan Birchfield. Foundation-stereo: Zero-shot stereo matching. *arXiv preprint arXiv:2501.09898*, 2025.
- [66] Changyeon Won and Hae-Gon Jeon. Learning depth from focus in the wild. In *European Conference on Computer Vision*, pages 1–18. Springer, 2022.
- [67] Yicheng Wu, Vivek Boominathan, Huaijin Chen, Aswin Sankaranarayanan, and Ashok Veeraraghavan. Phasecam3d—learning phase masks for passive single view depth estimation. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12. IEEE, 2019.
- [68] Shumian Xin, Neal Wadhwa, Tianfan Xue, Jonathan T Barron, Pratul P Srinivasan, Jiawen Chen, Ioannis Gkioulekas, and Rahul Garg. Defocus map estimation and deblurring from a single dual-pixel image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2228–2238, 2021.
- [69] Yalin Xiong and Steven A Shafer. Depth from focusing and defocusing. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 68–73. IEEE, 1993.
- [70] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10371–10381, 2024.
- [71] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024.
- [72] Kyrollos Yanny, Nick Antipa, William Liberti, Sam Dehaeck, Kristina Monakhova, Fanglin Linda Liu, Konlin Shen, Ren Ng, and Laura Waller. Miniscope3d: optimized single-shot miniature 3d fluorescence microscopy. *Light: Science & Applications*, 9(1):171, 2020.
- [73] Teresa Yeo, Oğuzhan Fatih Kar, Zahra Sodagar, and Amir Zamir. Rapid network adaptation: Learning to adapt neural networks using test-time feedback. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4674–4687, 2023.

- [74] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9043–9053, 2023.
- [75] Jiuling Zhang. Survey on monocular metric depth estimation, 2025.
- [76] Yucheng Zheng and M Salman Asif. Joint image and depth estimation with mask-based lensless cameras. *IEEE Transactions on Computational Imaging*, 6:1167–1178, 2020.
- [77] Yuxuan Zheng, Yifan Li, Yizhuo Zhang, Yiran Zhang, Lin Zhang, and Lei Zhang. TiNO-Edit: Timestep and noise optimization for robust diffusion-based image editing. *arXiv preprint arXiv:2304.06720*, 2023.
- [78] Changyin Zhou, Stephen Lin, and Shree Nayar. Coded aperture pairs for depth from defocus. In *2009 IEEE 12th international conference on computer vision*, pages 325–332. IEEE, 2009.

NeurIPS Paper Checklist

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 demonstrate evidence of the claims in the abstract in the section 5.

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: In section 6, we discuss the limitations of our method.

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]

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 will mention all the hyperparameters in the supplement.

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 the supplemental material?

Answer: [Yes]

Justification: We will release the code after acceptance. But the dataset will be provided as an anonymous link in the supplement.

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 specify implementation details in the paper; additional details will be included in the supplement.

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: [Yes]

Justification: We perform thorough ablations on different initialization, and characterize the effects of different aperture sizes on our method through a synthetic study. In the supplement, we characterize the effect of the stochasticity of the initial latent noise vector of the diffusion model by repeating an experiment 10 times with different random initializations for the latent noise vector, and report the variance in the observed final metrics.

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 mention the GPU used, memory usage, and optimization time for our algorithm.

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: We confirm to the ethics guidelines.

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: Our work has no particularly negative societal impacts.

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 work 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: [Yes]

Justification: We have cited the datasets and models we used for comparison. These datasets have been widely used in published papers on this topic.

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: [Yes]

Justification: We will release the self-collected dataset in the supplementary material with guidelines on how to use it, and all the camera parameters used while capturing it.

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: No crowdsourcing involved.

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: No human subjects involved

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.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: This paper is not related to LLM research or applications.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.