

This ACCV 2024 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2024 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

# PrimeDepth: Efficient Monocular Depth Estimation with a Stable Diffusion Preimage

Denis Zavadski\*, Damjan Kalšan\*, and Carsten Rother

Computer Vision and Learning Lab, IWR, Heidelberg University, Germany {name.surname}@iwr.uni-heidelberg.de

Abstract. This work addresses the task of zero-shot monocular depth estimation. A recent advance in this field has been the idea of utilising Text-to-Image foundation models, such as Stable Diffusion [51]. Foundation models provide a rich and generic image representation, and therefore, little training data is required to reformulate them as a depth estimation model that predicts highly-detailed depth maps and has good generalisation capabilities. However, the realisation of this idea has so far led to approaches which are, unfortunately, highly inefficient at test-time due to the underlying iterative denoising process. In this work, we propose a different realisation of this idea and present PrimeDepth, a method that is highly efficient at test time while keeping, or even enhancing, the positive aspects of diffusion-based approaches. Our key idea is to extract from Stable Diffusion a rich, but frozen, image representation by running a single denoising step. This representation, we term preimage, is then fed into a refiner network with an architectural inductive bias, before entering the downstream task. We validate experimentally that PrimeDepth is two orders of magnitude faster than the leading diffusion-based method, Marigold [28], while being more robust for challenging scenarios and quantitatively marginally superior. Thereby, we reduce the gap to the currently leading data-driven approach, Depth Anything [72], which is still quantitatively superior, but predicts less detailed depth maps and requires 20 times more labelled data. Due to the complementary nature of our approach, even a simple averaging between PrimeDepth and Depth Anything predictions can improve upon both methods and sets a new state-of-the-art in zero-shot monocular depth estimation. In future, data-driven approaches may also benefit from integrating our preimage.

Keywords: Monocular Depth Estimation  $\cdot$  Diffusion Models

## 1 Introduction

Depth estimation from a single image is a long-standing problem in computer vision. Although it is formally ill-posed, metric and affine-invariant depth estimation are well-defined tasks, and humans can also solve it approximately from cues such as texture-gradient, shading and shadows, or relative size [45]. Monocular-image depth estimation can also serve as an essential component for

<sup>\*</sup>Equal Contribution



Our PrimeDepth (0.57 sec)

Fig. 1: Results of Depth Anything [72] (top row), Marigold [28] (middle row) and our PrimeDepth (bottom row), for a challenging scene from the ETH3D Dataset [59]. While our non-optimised method is fast at test time (0.57 sec<sup>1</sup>), as well as Depth Anything (0.13 sec), Marigold is rather slow (62.84 sec). Runtimes were measured on an A100 GPU. Visually, our result shows most details, i.e. more detailed than Depth Anything and less grainy than Marigold, which however is not reflected in the quantitative numbers for this image (Depth Anything  $\delta_1 = 99.95\%$ , Marigold  $\delta_1 =$ 99.88% and Ours  $\delta_1 = 99.18\%$ ). The reason is the sparse ground truth LIDAR data with holes for objects with fine details (see bottom, left). While the data-driven method, Depth Anything, requires a large corpus of training data (1.5M labelled and 62M unlabelled images), ours and Marigold only need 74K synthetic training images.

monocular-video depth estimation of dynamic scenes [6] or in case of a forwardmoving camera, as in a famous "Hitchcock shot".

In the area prior to deep learning, depth estimation was only performed on a coarse superpixel-resolution, e.g. [54], or by approximating the scene with vertical planes, as in the automatic photo pop-up work [23]. In recent years, the research field underwent a major transformation and current methods achieve a level of detail that is oftentimes not even captured in the respective ground

 $<sup>^1\</sup>mathrm{With}$  precomputed captions. On the fly, BLIP-2 [36] image captioning adds on average additional 0.13 sec in runtime.

truth data, due to physical limitations of depth measurement devices, such as LIDAR.

In this work, we address the task of zero-shot, monocular depth estimation. This means that, once trained, our approach should work well on a broad range of image domains, such as indoor and outdoor, as well as in a variety of conditions, such as imagery of night scenes. Arguably, the first work that attempted zero-shot depth estimation has been MiDaS [49]. The first MiDaS version uses a high-capacity encoder, which got replaced by a Vision-Transformer [11] later on (MiDaS v3.0 [48]). The key strategy to success has been to use a very large corpus of labelled training images, real and synthetic, from different domains. The latest work with this data-driven philosophy is Depth Anything [72], which achieves impressive results using 1.5M labelled and additionally 62M unlabelled training images. It is a transformer-based architecture which is very fast at test time. Fig. 1 shows a result of Depth Anything [72] for a  $768 \times 512$  image.

Given the rise of generative Text-to-Image Diffusion models, with Stable Diffusion (SD) [51] as one of the first foundation models, a new paradigm for monocular depth estimation has emerged. It has been observed in various works [39,65] that SD "understands the intrinsic representation of the real 3D world" to some extent in order to generate images. For instance, in cross- and self-attention maps single objects are visible. Equipped with this insight, various works [28, 32, 34, 46, 66, 80] have used the power of SD, or other diffusion models, to perform depth estimation from a given image. There are, in general, two lines-of-work for doing so. The most prominent line-of-work [28, 34] takes a diffusion model and "repurposes" it for the task of depth estimation, such as denoising a random noise-image to a depth-image output, guided by the given image. In contrast to pure data-driven approaches, it estimates a more detailed depth and, according to [34], also generalises better to very different domains such as art paintings. Furthermore, since the knowledge of "realism" originates from the diffusion model, it can be trained well with as little as 74K synthetic training images, e.g. [28], which is only 4.9% of labelled training images compared to the data-driven Depth Anything [72] method. Fig. 1 shows a result from Marigold [28], which is, arguably, the leading method of this line-of-work.

However, this line-of-work has two major drawbacks: (i) It is rather slow, since it requires multiple denoising steps to map random noise to a depth map. Furthermore, due to the probabilistic nature, Marigold [28] has to average over an ensemble output; (ii) If a latent Diffusion Model is used, e.g. SD, the depth-based loss function has to be defined in latent domain. We show later that this leads to sub-optimal performance. These two drawbacks can be redeemed by simply using SD solely as a feature representation, i.e the features of a *single* denoising step of SD are extracted and fed into a depth prediction module. This marks the second line-of-work [32,33,46,80] which, to the best of our knowledge, has so far not been applied to zero-shot depth estimation. Indeed, when integrating SD-features in a straightforward manner to a downstream task then the performance is considerably below state-of-the-art, as shown later. The focus of this work is to look into the details of *what* should be extracted from SD and *how* it should be

integrated into a downstream network. This analysis pays-off, and we arrive at a method that is quantitatively marginally superior to the leading diffusion-based method, Marigold [28], while being on average over  $100 \times$  faster. When comparing with Marigold, our results are also more robust to extreme imaging conditions such as night scenes with very limited illumination. We conjecture that this is due to Marigold changing the weights of the Stable Diffusion prior while we keep the rich representation unchanged. We call our method PrimeDepth, and show a result in Fig. 1. To summarise, our contributions are as follows:

- Preimage, an image representation of feature maps and attention maps that are derived from the single, last denoising step of Stable Diffusion [51]. The preimage can effectively be used for a downstream task when combined with a refiner network that has an inductive bias by design.
- PrimeDepth, a single-step diffusion-based approach for zero-shot depth estimation. It has two conceptual advantages over multi-step diffusion approaches: i) very fast inference, e.g. on average 100× faster than Marigold [28];
   ii) utilizing a loss function in pixel domain and not latent domain.
- PrimeDepth is in our quantitative evaluation the runner-up, and only inferior to Depth Anything [72] that is, however, a data-driven approach and requires over 20× more labelled training data, as well as roughly 25× longer training times based on the information provided in [72]. PrimeDepth also predicts more detailed depth maps than Depth Anything and is more robust than Marigold [28] for challenging scenes such as at nighttime.

## 2 Related Work

#### 2.1 Denoising Diffusion Models

Since their introduction [21, 62, 64], diffusion models have vastly advanced the field of unconditional and conditional image generation [9, 22, 26]. By construction, these models have the theoretical capability to represent any arbitrary data distribution [62]. This makes them the spearhead in Text-to-Image (T2I) generation [31, 47, 51, 81] and tasks requiring a rich and faithful image representation, such as image editing [7, 8, 18, 25] or image-to-image translation [52, 67]. Many works have been dedicated to speed-up the generation process starting from optimised sampling [10, 42, 63, 79] up to the distillation of the internal representation of pre-trained diffusion models into single step methods [37, 53, 73]. In this work, we use the rich representation of the pre-trained Stable Diffusion and apply it to the task of zero-shot monocular depth estimation using only one single iteration.

#### 2.2 Zero-Shot Monocular Depth Estimation

In contrast to earlier works, which addressed in-domain monocular depth estimation [1,14–16,35,80], zero-shot depth estimators aim to generalise to unseen data distributions. Although there have been efforts on ordinal [5,70] and metric depth estimation [75], in this work, we focus on the affine-invariant setting [49]. Thus far, two main research directions have been explored. The first leverages strong

925

5

priors, such as T2I diffusion models (see Sec. 2.3), while the second, discussed in this section, is data-driven [13,48,49,72,76,78]. MiDaS [49] proposes training on diverse datasets, while DPT [48] and Omnidata [13] further expand the number of training data and use transformers. LeReS [76] proposes a framework that can additionally estimate the depth-shift and focal length, while HDN [78] improves on detailed depth prediction with a multi-scale depth normalisation method. Finally, Depth Anything [72] expands the datasets to large-scale unlabelled data with a challenging student-teacher optimisation target. Since Depth Anything is, to the best of our knowledge, quantitatively the leading method in this research direction, we perform a detailed comparison to it.

### 2.3 Diffusion-Based Monocular Depth Estimators

Diffusion models for depth prediction can be roughly assigned to three different paradigms.

**Diffusion paradigm** includes methods that take a not pre-trained diffusion framework and train it from scratch to denoise depth maps [12, 27, 55–57]. DepthGen [57] conditions diffusion on an RGB image and operates in pixel domain. The authors addressed training on incomplete noisy ground truth data using simple interpolation and unsupervised auxiliary pre-training. DDVM [55] extended pre-training to include synthetic data and reduced the computational overhead through coarse-to-fine refinement in a patch-wise manner. DMD [56] extends DDVM to a joint indoor-outdoor predictor by learning depth in log space. DDP [27] and DiffusionDepth [12] encode the conditional image with an off-the-shelf feature extractor to a smaller resolution prior to diffusion. Furthermore, such methods yield a posterior uncertainty by design based on unstable pixel regions [27, 55–57].

**Pre-trained diffusion paradigm** uses a pre-trained network, oftentimes a T2I model, and fine-tunes it on the downstream task [28,34]. This paradigm has recently garnered popularity with the release of Marigold [28]. Marigold performs diffusion and depth predictions in latent domain and is trained exclusively on limited synthetic data. Its main drawback is the computational overhead with a multi-step denoising approach, combined with an ensemble scheme. Due to conceptual similarity, we compare extensively to Marigold. The goal of DMP [34] is to show the generalisation of Stable Diffusion [51] for downstream tasks such as depth estimation of highly creative synthetic data. It is done by training a LoRA [24] and reformulating denoising as blending the image and depth map to achieve deterministic prediction. This paradigm demonstrates good zero-shot performance with detailed predictions, however, it is slow at inference time.

**Representation extractor paradigm**, in contrast to the previous, considers the pre-trained T2I network primarily as a feature representation akin to traditional backbones [32,33,46,66,69,80]. The representation is extracted using image inversion and fed to a trainable depth predictor. Our method, PrimeDepth, follows this paradigm. For depth estimation, the pioneering work VPD [80] extracts intermediate image features at four resolutions from Stable Diffusion [51] and fuses them into a single low-resolution feature block using a semantic Feature Pyramid Network (FPN) [29]. The block is then processed by a lightweight

upsampling decoder to form the final dense prediction. Notably, VPD demonstrated that diffusion models provide a competitive representation that surpasses traditional visual backbones like SwinV2 [40]. VPD considered only the task of in-domain monocular depth estimation. In order to still compare to VPD, we build a VPD-like architecture within our framework and show that it is inferior. Recently, TADP [32] investigated the text guidance employed by VPD, showing that text-image alignment promotes better performance on downstream tasks. In their in-depth analysis of various captioning approaches they found that automated image captioning such as BLIP-2 [36] yield best performance. A concurrent work, ECoDepth [46] claimed that text guidance lacks expressiveness for perception tasks and replaced the text embeddings with ViT [11] embeddings extracted from the input image. Finally, MetaPrompts [66] completely obviates the need for captions by making the text embeddings trainable. In our work, we use BLIP-2 for caption generation, since improving text embeddings is not our focus. The big advantage of this paradigm compared to the previous is inference speed, as it requires only one single network forward pass.

### 3 Method

#### 3.1 Stable Diffusion

As a latent Text-to-Image (T2I) diffusion model, Stable Diffusion (SD) generates images by successively transforming Gaussian noise  $z_T$  to a latent encoding  $z_0$  in 1000 overall steps. This generative process in latent domain is performed by a U-Net f consisting of different neural blocks, *i.e.* residual, self- and crossattention blocks, that are executed sequentially and at different scales. The excessive training of SD on the diverse LAION-5B [60] dataset allows for a rich representation of arbitrary imagery. The motivation for our SD representation stems from [20,43], which has shown that for the semantic correspondence task, the last denoising step  $f(z_1)$  contains rich information. This intuitively makes sense since at that time-step the representation contains the least amount of "potential hallucinations" of the generative model.

#### 3.2 Preimage Representation

We define our preimage as all the intermediate, multi-scale feature maps, crossand self-attention maps of every neural block in the last denoising step of SD version 2.1, see Fig. 2 (left). In contrast to prior work VPD [80], we also utilise self- and cross-attention maps which encode similar but complementary learned structural and semantic concepts. While the feature activation maps can be used right away, the attention maps require further processing to reduce the respective channel-size and possible redundancies. With  $\mathcal{H}$  attention heads, height h and width w, the self-attention maps have the dimension  $\mathcal{H} \times h \times w \times N_p$ , with  $N_p =$  $h \cdot w$  being the number of pixels, since all pixels attend to each other. Similarly, the cross-attention maps, attending to 77 text tokens, have the dimension  $\mathcal{H} \times h \times w \times 77$ . For an image with the resolution of  $512 \times 512$ , the corresponding largest self-attention map in latent domain has the dimension  $\mathcal{H} \times 64 \times 64 \times 4096$ . Using all maps is not only computationally infeasible, but also impossible, as the

927



**Fig. 2:** (Left) Stable Diffusion preimage consisting of intermediate feature maps, crossand self-attention maps for every neural block of the last denoising step. (Right) Examples of self-attention maps, with respect to the red square, and cross-attention maps. Below is the fusion model shown, for the aggregation of the preimage parts.

channel size  $N_p$  changes with the resolution of the input. We circumvent both problems by dividing the attention maps in  $8 \times 8$  regions and averaging over the respective attention-regions. Because each head can potentially attend to, and therefore encode, different kind of content, we do not average over heads. This means formally,

$$\mathbb{R}^{\mathcal{H} \times h \times w \times N_p} \xrightarrow{\text{regional}} \mathbb{R}^{\mathcal{H} \times h \times w \times 64} \longrightarrow \mathbb{R}^{h \times w \times (64 \cdot \mathcal{H})} \qquad \text{(self-attention)} \qquad (1)$$
$$\mathbb{R}^{\mathcal{H} \times h \times w \times 77} \longrightarrow \mathbb{R}^{h \times w \times (77 \cdot \mathcal{H})} \qquad \text{(cross-attention)}. \qquad (2)$$

Fig. 2 (right) shows examples for attention maps after averaging over regions. Before passing the extracted parts of the preimage to the refining network, the features need to be aggregated. The fusion module for the aggregation is illustrated in Fig. 2 (bottom, right) and comprises a sequence of concatenations and convolution projections of all feature and attention maps.

#### 3.3 PrimeDepth Architecture

The full architecture design of PrimeDepth is shown in Fig. 3. The depth prediction pipeline starts by encoding the image x with the SD Pixel-to-Latent (P2L) encoder to its latent domain representation  $z_0$  and adding one step of noise to receive the slightly noisy latent  $z_1$ . In the denoising step of the frozen U-Net, the preimage parts, shown with red arrows, are passed through fusion modules (Fig. 2, bottom right) and provided to the preimage refining decoder via concatenations at the respective blocks. The output of the refining decoder is a latent representation that can be used by downstream decoder-heads to get the final output in pixel domain. In contrast to diffusion-based approaches that use multiple-denoising steps, like Marigold [28], we define the loss function in the output domain. The downside of defining the loss function in latent domain like Marigold is twofold: i) The task decoder cannot be trained at the same time; ii) The loss function in latent domain is a surrogate of the one in pixel domain.

The downstream heads are trainable copies of the SD latent-to-pixel (L2P) decoder with an increased bottleneck size from 4 input channels to 512. We use two downstream heads for training, one for depth prediction and the other



Fig. 3: PrimeDepth Pipeline. The input image is first encoded to latent domain, augmented with one noise step and processed by the frozen U-Net of Stable Diffusion. The intermediate parts of the preimage (red arrows) are aggregated with the fusion module (see Fig. 2) and provided to the preimage refiner network at the respective intermediate stages. The output of the refiner is fed to two downstream heads for the respective downstream tasks.

one for semantic segmentation as regularisation. We utilise 150 semantic classes as in [82] and see later that it improves the performance. Note that previous works have also demonstrated the benefits of auxiliary tasks for regularisation, including semantics [4, 19, 30, 71] or normal prediction [41, 74].

To summarise, PrimeDepth fulfils three conceptual requirements. (i) The base model weights remain unaltered. We believe that this is an advantage, since SD was trained extensively on a large amount of data. By adapting the weights, as e.g. done in Marigold [28] and VPD [80], the representation capabilities may decrease and with this the generalisation power, as seen later in reduced robustness of e.g. Marigold. (ii) Utilisation of the full preimage representation. As we see later, using the full representation of SD, i.e. using all intermediate feature maps, self- and and cross-attention maps, is beneficial for the depth prediction task. (iii) Exploiting architectural inductive bias for the preimage representation. The refining decoder has the same architecture as the SD decoder, but with only 50% of the original channel sizes. Hence it introduces a natural inductive bias to the architecture for the gradual processing of the preimage. All components of the preimage are provided to the refiner at its respective resolution stages.

It is worth to note that using an inductive architectural bias is not new, and for instance has shown to be beneficial in the context of controlled image generation for SD [77]. We believe that transformer architectures could also benefit from the integration of the generative preimage with architectural bias and sketch a possible integration to a DPT-based [48] refiner in the supplement.

### 3.4 Training Protocol

Although the preimage refiner network has the same architecture as the SDdecoder, it has to be initialised with random weights since the channel sizes

9

differ. To compensate for the random initialisation, we first pre-train the model with unlabelled data. Note that this is a common strategy, also used for instance by Depth Anything [72], and gives a marginal quantitative boost in performance as shown later. For the pre-training, we use a subset of 600K images from the LAION-Aesthetics dataset, which contains images of high visual quality from arbitrary domains. We compute a pseudo ground truth for depth using Depth Anything [72] and for segmentation maps with InternImage-H [68]. We use BLIP-2 [36] to captionise all images. The model is first pre-trained on the pseudo ground truth data and then trained on 74K images of synthetically labelled data. Since our model uses a preimage that is computed in one single forward pass, we are not restricted by latent losses. We use the scale and shift invariant loss of MiDaS [49] for the depth estimation in pixel domain. We first shift and scale the ground truth depth  $d^* = \frac{d-t(d)}{s(d)}$  with t(d) = median(d) and  $s(d) = \frac{1}{hw} \sum_{i=1}^{hw} |d_i - t(d)|$  and then compute the mean squared error between the affine-invariant ground truth  $d^*$  and the depth prediction  $\hat{d}$ , which is aligned to  $d^*$  by least squared errors  $\alpha(d^*, \hat{d})$ , to get the shift and scale invariant loss

$$\mathcal{L}_{\rm ssi} = {\rm MSE}\left(d^*, \alpha(d^*, \hat{d})\right).$$
(3)

For the regularisation through semantic segmentation, we use a combination of the Dice Loss [44] and the Focal Loss [38] with respect to the one-hot encoded semantic segmentation maps for C = 150 classes

$$\mathcal{L}_{\text{dice}} = 1 - \frac{1}{C} \sum_{c=0}^{C-1} \frac{2\sum_{i=1}^{hw} y_i^{*c} \cdot \hat{y}_i^c}{\sum_{i=1}^{hw} (y_i^{*c} + \hat{y}_i^c)}, \qquad \mathcal{L}_{\text{focal}} = -\sum_{i=1}^{hw} (1 - y_i^* \cdot \hat{y}_i)^2 \log(y_i^* \cdot \hat{y}_i) \quad (4)$$

with  $y^*, \hat{y}$  being the ground truth and the predicted segmentation maps, respectively. We dynamically weigh the depth loss to the segmentation loss and arrive at the final training objective

$$\mathcal{L} = \mathcal{L}_{dice} + \mathcal{L}_{focal} + \lambda \mathcal{L}_{ssi}, \quad \text{with} \quad \lambda = sg\left(\frac{\mathcal{L}_{dice} + \mathcal{L}_{focal}}{\mathcal{L}_{ssi}}\right), \tag{5}$$

where  $sg(\cdot)$  represents the stop-gradient operator.

### 4 Experiments

#### 4.1 Datasets and Evaluation Metrics

**Training datasets.** We utilised the same datasets as Marigold [28] to allow for a fair comparison. These are the two synthetic and densely-labelled datasets Hypersim [50] and Virtual KITTI 2 [2], which represent indoor and outdoor scenarios. The LAION-Aesthetics dataset [60] is used for pre-training our method, see details in Sec. 3.4, and gives a marginal boost in performance, see below.

**Test datasets.** We evaluate in total on five unseen datasets: KITTI [17] using Eigen split [14] (654 images), NYUv2 [61] (656 images), ETH3D [59] (456



Fig. 4: Qualitative results of two competing methods (Depth Anything [72] and Marigold [28]) for 4 datasets, while results for ETH3D are shown in Fig. 1. The main visual artefacts of the respective methods are indicated by arrows. The prominent observations across many images are as follows. Depth Anything has less sharp depth maps (KITTI, nuScenes-C, Fig. 1) and can see inside mirrors (NYUv2) and through transparent surfaces (supplement). Marigold predicts sharp depth maps but sometimes with grainy artefacts (KITTI, Fig. 1). It also struggles to predict sky (KITTI, nuScenes-C) and objects at mid-distance (rabbitai). Our method gives sharper depth maps than Depth Anything (KITTI, NuScenes-C), but can also struggle with sky (supplement).

images), rabbitai [58] (60 images)<sup>2</sup>, and a manually chosen subset of nuScenes [3], we term nuScenes-C (3489 images). Since KITTI, ETH3D and NuScenes-C were captured with LIDAR sensors, the ground truth measurements are sparse with occasional holes, *e.g.* in the sky and objects with fine details. The NYUv2 dataset was captured with an active RGBD-kinect sensor. It has dense ground truth apart from occasionally missing regions at depth-borders. While the rabbitai dataset has fewest number of images, it represents, at the same time, a dataset with dense and highly detailed ground truth depth. It was captured with a passive light-field camera setup. We created the nuScenes-C dataset by selecting 87 scenes (in total 3489 images) from nuScenes that represent challenging scenarios, such as images in bad weather conditions, nighttime, and adversarial visual cues including reflections and artefacts on the lens. We further categorised the nuScenes-C dataset into 6 categories, see Fig. 5.

**Metrics.** As in training, a predicted depth map  $\hat{d}$  is first affinely aligned with ground truth  $d^*$ . We use two common metrics: i) absolute relative error  $AbsRel = \frac{1}{hw} \sum_{hw} \frac{|d^* - \hat{d}|}{d^*}$  and ii) accuracy under threshold  $\delta_1 = \frac{1}{hw} \sum_{hw} [max(\frac{\hat{d}}{d^*}, \frac{d^*}{\hat{d}}) < 1.25]$ , measuring the fraction of pixels that satisfy the condition.

<sup>&</sup>lt;sup>2</sup>See Robust Vision Challenge 2020 (http://www.robustvision.net/rvc2020.php).

**Table 1:** Comparison of 8 methods for *zero-shot* monocular depth estimation, where numbers are taken from Marigold [28], apart from the numbers of Depth Anything [72], taken from their article, and our numbers using our protocol. Since the methods may not be evaluated consistently, the average rank may be the best performance indicator.

Model	$\delta_1 \uparrow$	ITTI AbsRel↓	$\delta_1 \uparrow$	VUv2 AbsRel $\downarrow$	$\delta_1 \uparrow$	TH3D AbsRel↓	Average Rank	Num. Train Samples	Architecture Type
MiDaS [49]	63.0	23.6	88.5	11.1	75.2	18.4	8.0	2M	ResNeXt101
LeReS [76]	78.4	14.9	91.6	9.0	77.7	17.1	6.5	354K	ResNeXt101
Omnidata [13]	83.5	14.9	94.5	7.4	77.8	16.6	5.7	12.2M	Transformer
DPT [48]	90.1	10.0	90.3	9.8	94.6	7.8	4.7	1.4M	Transformer
HDN [78]	86.7	11.5	94.8	6.9	83.3	12.1	4.5	300K	Transformer
Marigold [28]	91.6	9.9	96.4	5.5	96.0	6.5	2.3	74K	Multi-step diffusion
Depth Anything [72]	94.7	7.6	98.1	4.3	88.2	12.7	2.2	63.5M	Transformer
Our	<u>93.7</u>	7.9	<u>96.6</u>	5.8	96.7	<u>6.8</u>	2.0	74K	Single-step diffusion

#### 4.2 Comparison to other Methods

While reproducing the results of previous methods, it became apparent that the evaluation for the same datasets can unfortunately differ substantially between articles. As an example, Depth Anything [72] reports  $\delta_1 = 88.2\%$  for the ETH3D dataset, while using our protocol, which tries to match the numbers of Marigold [28], the score is much better, i.e.  $\delta_1 = 98.1\%$  for ETH3D. Tab. 1 shows a comparison of 8 methods for three unseen datasets, where all numbers are from Marigold, apart from the numbers of Depth Anything [72] which are from their article and our numbers using our protocol. Since individual numbers might not be fully representative, we provide the average rank as an indicator for overall performance<sup>3</sup>. We observe that the two recent methods, Marigold and Depth Anything, have clearly the best rank, i.e. 2.3 and 2.2, alongside our approach with rank 2.0. Given this, we do a detailed analysis of these three methods for 5 unseen datasets using our evaluation protocol, see Tab. 2. We see that Depth Anything is superior to the other methods, while our method is clearly the runner-up. Our method without pre-training is on a par with Marigold. As mentioned in the introduction, Marigold and our approach only use a fraction of labelled training data compared to Depth Anything (74K versus 1.5M). Furthermore, since our approach uses only a single diffusion step, we are on average 100 times faster than Marigold at test time.<sup>4</sup> Despite the pre-training on Depth Anything predictions, the results of PrimeDepth and Depth Anything are complementary. To show this, we provide two simple approaches for combining the results of Depth Anything and PrimeDepth. The Pixel-wise-Average method simply averages the two results pixel-wise. It improves over both Depth Anything and PrimeDepth, meaning that the two signals are not fully correlated. We also run an Image-Oracle, which selects either the Depth Anything or PrimeDepth result. The oracle improves the results noticeably, and in 21%to 55% of cases the PrimeDepth result is selected. Qualitative results are shown

<sup>&</sup>lt;sup>3</sup>We take the average with respect to both measures,  $\delta_1$  and AbsRel.

<sup>&</sup>lt;sup>4</sup>Computed from 1000 runs of different image resolutions on an A100 GPU.

**Table 2:** Comparison to Depth Anything (DA) [72] and Marigold [28] for *zero-shot* monocular depth estimation using the same (our) evaluation protocol.



Fig. 5: Box plot for the  $\delta_1$  accuracy of challenging scenes from the nuScenes-C dataset, split into 6 categories, where long, vertical stripes provide the median values. Our PrimeDepth is consistently, marginally inferior to Depth Anything [72], but consistently and sometimes considerably superior to Marigold [28]. The variability, measured by IQR score i.e. size of a box, is considerably higher for Marigold than Depth Anything and our method. For nighttime scenes, the performance drops for all methods, however Marigold is clearly more affected (lowest median and highest variability).

and discussed in Fig. 1 and Fig. 4. In brief, the depth maps of Depth Anything exhibit blurriness, although this is not reflected in the quantitative results due to missing values in the ground truth. Also, Depth Anything has the same artefacts as in the ground truth, e.g. it sees through reflective objects. Marigold and our method are visually more detailed, however, again, giving missing ground truth it is difficult to clearly measure its correctness. Marigold is occasionally too detailed which results in grainy artefacts, and sometimes struggles to predict depth at mid-distance. Our method and Marigold sometimes do not predict sky at infinity.

#### 4.3 Comparison for challenging scenarios

To better understand the behavior of methods for challenging scenarios, we visualise their performance with respect to 6 different categories in form of a boxplot. Fig. 5 shows the  $\delta_1$  accuracy computed on a per-frame basis. The median value is marked with a long, vertical stripe. The variability in the results is shown by the size of the solid box, which is the interquartile range (IQR) containing 50% of all inlier predictions. We observe that the general trend matches that of the benchmark above. Our method is marginally inferior to Depth Anything [72],



Fig. 6: Challenging scenes for category "Day Rain" (top row), "Night Rain" (middle row) and "Night Dark" (bottom row). Depth Anything [72] wrongly predicts depth for reflections on a wet surface (top row). Marigold [28] struggles to predict the depth of objects at mid-depth (bottom row) and sky (all rows). A failure case for both methods, Marigold and ours, is a water drop on the camera lens which is likely predicted as a round object in the scene (middle row).

but consistently superior to Marigold [28]. The IQR scores of our method and Depth Anything are similar and relatively small, in contrast to Marigold. For nighttime scenes, the drop in median score and higher IQR scores is noticeable for all methods, however, considerably and consistently more prominent for Marigold. We conjecture that the low robustness of Marigold stems from the fact that Stable Diffusion itself is modified from being an image generation method to being a depth prediction method. In contrast, we retain the rich information of the frozen Stable Diffusion model. Qualitative samples of challenging scenes at night and in rain are shown in Fig. 6. It is noticeable that Marigold struggles to predict objects at mid-distance as well as the sky region. Depth Anything wrongly predicts depth for reflections on a wet surface. A failure case of both, Marigold and our approach, is a water drop on the camera lens, which is likely predicted as a round object in the scene.

#### 4.4 Ablation Study

To evaluate the effectiveness of the individual components of our method, we ablate in Tab. 3 each design choice after pre-training on pseudo ground truth images only. The PrimeDepth architecture is called V.Full and performs best across the board. By training the model in latent domain (V.Latent), as done by methods like Marigold [28], we see a decrease in performance across all datasets. Omitting the regularisation with respect to the segmentation loss (V.NoSegmen), results in a similar decrease in the overall accuracy. Using V.NoSegmen and additionally removing either the feature maps (V.NoFeature) or the attention maps (V.NoAttent), gives a further decline in performance. Overall, feature maps appear to be more informative than attention maps, but attention maps still contain useful additional information. We also compare our network design with designs suggested in previous works [32,80]. For this we use our full preimage and

**Table 3:** Ablation Study of different design choices, noted by Version.\* (V.\*). All versions are only pre-trained on a 600K image subset of LAION, and not trained on ground truth data as our variants in Tab. 1 and Tab. 2.

					ł	KITTI	r	rabbitai		cenes-C	Avg. Rank
Model	Segm	Atten	Feat	PixLoss	$\delta_1 \uparrow$	AbsRel	$\downarrow \delta_1 \uparrow$	AbsRel	$\downarrow \delta_1 \uparrow$	AbsRel	k
V.Full	x	x	x	x	92.8	8.7	73.5	21.9	80.6	14.2	1.0
V.Latent	x	x	х	-	92.5	8.8	70.1	24.4	79.8	15.0	3.0
V.NoSegmen	-	x	х	x	92.8	8.9	72.9	24.6	79.8	15.0	3.0
V.NoAttent	-	-	х	x	<u>92.7</u>	9.3	72.7	25.4	79.1	16.3	4.2
V.NoFeature	-	x	-	x	88.0	11.4	67.9	23.9	76.5	16.9	4.8
V.BlockPreIm	-	х	х	x	86.6	12.4	65.3	26.5	80.1	14.8	4.7

re-scale it as a single "block" with uniform resolution and provide it as regular input to the downstream task. The key difference to our design is that the natural inductive bias of the preimage refiner is not exploited. To be in spirit close to [80], we inserted the preimage of V.NoSegmen as a block into our preimage refiner. We refer to this design as V.BlockPreIm (see full network design in supplement). The performance of V.BlockPreIm, compared to V.NoSegmen, is clearly inferior for the KITTI dataset and the densely labelled rabbitai dataset, while for the sparse nuScenes-C data, the performance is basically on par.

## 5 Conclusion and Future Work

In this work, we presented PrimeDepth, a method for harnessing and utilising the complete, rich preimage of Stable Diffusion for zero-shot monocular depth estimation. Our model achieves competitive results on a variety of datasets while being considerably faster than competitive diffusion-based approaches, like Marigold [28], owing to a single diffusion step. By keeping the diffusion prior fixed, our method achieves robust results while being trained on a fraction of synthetically labelled data compared to competing data-driven methods, like Depth Anything [72]. The competitive performance of PrimeDepth advocates for the richness of the preimage of large-scale generative models as a starting point for downstream tasks, potentially also beyond depth estimation. One avenue of future research is to explore the complementary nature of data-driven and diffusion-based approaches, since the latter may generalise better, see [34].We sketch in supplement how a preimage can be integrated into a DPT-based [48] approach, as used in Depth Anything.

## Acknowledgements

We thank Yannick Pauler, Nicolas Bender and Friedrich Feiden for their help. The project has been supported by the Konrad Zuse School of Excellence in Learning and Intelligent Systems (ELIZA) funded by the German Academic Exchange Service (DAAD). The authors gratefully acknowledge the support by the Ministry of Science, Research and the Arts Baden-Württemberg (MWK) through bwHPC, SDS@hd and the German Research Foundation (DFG) through the grants INST 35/1597-1 FUGG and INST 35/1503-1 FUGG.

## References

- Bhat, S.F., Alhashim, I., Wonka, P.: Adabins: Depth estimation using adaptive bins. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 4009–4018 (2021) 4
- 2. Cabon, Y., Murray, N., Humenberger, M.: Virtual kitti 2. arXiv preprint arXiv:2001.10773 (2020) 9
- Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 11621–11631 (2020) 10
- Chen, P.Y., Liu, A.H., Liu, Y.C., Wang, Y.C.F.: Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 2624–2632 (2019) 8
- Chen, W., Fu, Z., Yang, D., Deng, J.: Single-image depth perception in the wild. Adv. Neural Inform. Process. Syst. 29 (2016) 4
- Cheng, J., Yin, W., Wang, K., Chen, X., Wang, S., Yang, X.: Adaptive fusion of single-view and multi-view depth for autonomous driving. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 10138–10147 (2024) 2
- Choi, J., Choi, Y., Kim, Y., Kim, J., Yoon, S.: Custom-edit: Text-guided image editing with customized diffusion models. arXiv preprint arXiv:2305.15779 (2023)
   4
- Couairon, G., Verbeek, J., Schwenk, H., Cord, M.: Diffedit: Diffusion-based semantic image editing with mask guidance. arXiv preprint arXiv:2210.11427 (2022) 4
- Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Adv. Neural Inform. Process. Syst. 34, 8780–8794 (2021) 4
- Dockhorn, T., Vahdat, A., Kreis, K.: Genie: Higher-order denoising diffusion solvers. Adv. Neural Inform. Process. Syst. 35, 30150–30166 (2022) 4
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020) 3, 6
- Duan, Y., Guo, X., Zhu, Z.: Diffusiondepth: Diffusion denoising approach for monocular depth estimation. arXiv preprint arXiv:2303.05021 (2023) 5
- Eftekhar, A., Sax, A., Malik, J., Zamir, A.: Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In: Int. Conf. Comput. Vis. pp. 10786–10796 (2021) 5, 11
- 14. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. Adv. Neural Inform. Process. Syst. **27** (2014) **4**, **9**
- Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 2002–2011 (2018) 4
- Garg, R., Bg, V.K., Carneiro, G., Reid, I.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: Eur. Conf. Comput. Vis. pp. 740–756. Springer (2016) 4
- Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 3354– 3361. IEEE (2012) 9
- Goel, V., Peruzzo, E., Jiang, Y., Xu, D., Sebe, N., Darrell, T., Wang, Z., Shi, H.: Pair-diffusion: Object-level image editing with structure-and-appearance paired diffusion models. arXiv preprint arXiv:2303.17546 (2023) 4

- 16 D.Zavadski et al.
- Guizilini, V., Hou, R., Li, J., Ambrus, R., Gaidon, A.: Semantically-guided representation learning for self-supervised monocular depth. arXiv preprint arXiv:2002.12319 (2020) 8
- Hedlin, E., Sharma, G., Mahajan, S., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: Unsupervised semantic correspondence using stable diffusion. Adv. Neural Inform. Process. Syst. 36 (2024) 6
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Adv. Neural Inform. Process. Syst. 33, 6840–6851 (2020) 4
- Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. J. Mach. Learn. Res. 23(1), 2249–2281 (2022) 4
- Hoiem, D., Efros, A.A., Hebert, M.: Automatic photo pop-up. In: ACM SIG-GRAPH, pp. 577–584 (2005) 2
- Hu, E.J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: LoRA: Low-rank adaptation of large language models. In: Int. Conf. Learn. Represent. (2022) 5
- Hu, H., Chan, K.C., Su, Y.C., Chen, W., Li, Y., Sohn, K., Zhao, Y., Ben, X., Gong, B., Cohen, W., et al.: Instruct-imagen: Image generation with multi-modal instruction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 4754–4763 (2024) 4
- Hu, V.T., Zhang, D.W., Asano, Y.M., Burghouts, G.J., Snoek, C.G.: Self-guided diffusion models. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 18413–18422 (2023) 4
- Ji, Y., Chen, Z., Xie, E., Hong, L., Liu, X., Liu, Z., Lu, T., Li, Z., Luo, P.: Ddp: Diffusion model for dense visual prediction. In: Int. Conf. Comput. Vis. pp. 21741– 21752 (October 2023) 5
- Ke, B., Obukhov, A., Huang, S., Metzger, N., Daudt, R.C., Schindler, K.: Repurposing diffusion-based image generators for monocular depth estimation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 9492–9502 (June 2024) 1, 2, 3, 4, 5, 7, 8, 9, 10, 11, 12, 13, 14
- Kirillov, A., Girshick, R., He, K., Dollar, P.: Panoptic feature pyramid networks. In: IEEE Conf. Comput. Vis. Pattern Recog. (June 2019) 5
- Klingner, M., Termöhlen, J.A., Mikolajczyk, J., Fingscheidt, T.: Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In: Eur. Conf. Comput. Vis. pp. 582–600. Springer (2020) 8
- Koh, J.Y., Park, S.H., Song, J.: Improving text generation on images with synthetic captions. arXiv preprint arXiv:2406.00505 (2024) 4
- Kondapaneni, N., Marks, M., Knott, M., Guimaraes, R., Perona, P.: Text-image alignment for diffusion-based perception. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 13883–13893 (June 2024) 3, 5, 6, 13
- Lavreniuk, M., Bhat, S.F., Müller, M., Wonka, P.: Evp: Enhanced visual perception using inverse multi-attentive feature refinement and regularized image-text alignment. arXiv preprint arXiv:2312.08548 (2023) 3, 5
- Lee, H.Y., Tseng, H.Y., Lee, H.Y., Yang, M.H.: Exploiting diffusion prior for generalizable dense prediction. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 7861– 7871 (June 2024) 3, 5, 14
- Lee, J.H., Han, M.K., Ko, D.W., Suh, I.H.: From big to small: Multi-scale local planar guidance for monocular depth estimation. arXiv preprint arXiv:1907.10326 (2019) 4
- 36. Li, J., Li, D., Savarese, S., Hoi, S.: BLIP-2: Bootstrapping language-image pretraining with frozen image encoders and large language models. In: Krause, A.,

Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., Scarlett, J. (eds.) Int. Conf. Mach. Learn. Proceedings of Machine Learning Research, vol. 202, pp. 19730– 19742. PMLR (23–29 Jul 2023) 2, 6, 9

- Lin, S., Wang, A., Yang, X.: Sdxl-lightning: Progressive adversarial diffusion distillation. arXiv preprint arXiv:2402.13929 (2024) 4
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Int. Conf. Comput. Vis. pp. 2980–2988 (2017) 9
- Liu, B., Wang, C., Cao, T., Jia, K., Huang, J.: Towards understanding cross and self-attention in stable diffusion for text-guided image editing. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 7817–7826 (2024) 3
- Liu, Z., Hu, H., Lin, Y., Yao, Z., Xie, Z., Wei, Y., Ning, J., Cao, Y., Zhang, Z., Dong, L., Wei, F., Guo, B.: Swin transformer v2: Scaling up capacity and resolution. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 12009–12019 (June 2022) 6
- Long, X., Lin, C., Liu, L., Li, W., Theobalt, C., Yang, R., Wang, W.: Adaptive surface normal constraint for depth estimation. In: Int. Conf. Comput. Vis. pp. 12849–12858 (2021) 8
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. Adv. Neural Inform. Process. Syst. 35, 5775–5787 (2022) 4
- Luo, G., Dunlap, L., Park, D.H., Holynski, A., Darrell, T.: Diffusion hyperfeatures: Searching through time and space for semantic correspondence. Adv. Neural Inform. Process. Syst. 36 (2024) 6
- Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: Int. Conf. on 3D Imaging, Modeling, Processing, Visualization and Transmission. pp. 565–571. IEEE (2016) 9
- 45. Palmer, S.E.: Vision science: Photons to phenomenology. MIT press (1999) 1
- Patni, S., Agarwal, A., Arora, C.: Ecodepth: Effective conditioning of diffusion models for monocular depth estimation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 28285–28295 (June 2024) 3, 5, 6
- 47. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023) 4
- Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: Int. Conf. Comput. Vis. pp. 12179–12188 (2021) 3, 5, 8, 11, 14
- Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. IEEE Trans. Pattern Anal. Mach. Intell. 44(3), 1623–1637 (2020) 3, 4, 5, 9, 11
- Roberts, M., Ramapuram, J., Ranjan, A., Kumar, A., Bautista, M.A., Paczan, N., Webb, R., Susskind, J.M.: Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In: Int. Conf. Comput. Vis. (2021) 9
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 10684–10695 (June 2022) 1, 3, 4, 5
- Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH. pp. 1–10 (2022) 4
- Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation. arXiv preprint arXiv:2311.17042 (2023)
- Saxena, A., Sun, M., Ng, A.Y.: Make3d: Learning 3d scene structure from a single still image. IEEE Trans. Pattern Anal. Mach. Intell. 31(5), 824–840 (2008) 2

- 18 D.Zavadski et al.
- 55. Saxena, S., Herrmann, C., Hur, J., Kar, A., Norouzi, M., Sun, D., Fleet, D.J.: The surprising effectiveness of diffusion models for optical flow and monocular depth estimation. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Adv. Neural Inform. Process. Syst. vol. 36, pp. 39443–39469. Curran Associates, Inc. (2023) 5
- Saxena, S., Hur, J., Herrmann, C., Sun, D., Fleet, D.J.: Zero-shot metric depth with a field-of-view conditioned diffusion model. arXiv preprint arXiv:2312.13252 (2023) 5
- 57. Saxena, S., Kar, A., Norouzi, M., Fleet, D.J.: Monocular depth estimation using diffusion models. arXiv preprint arXiv:2302.14816 (2023) 5
- Schilling, H., Gutsche, M., Brock, A., Spath, D., Rother, C., Krispin, K.: Mind the gap-a benchmark for dense depth prediction beyond lidar. In: IEEE Conf. Comput. Vis. Pattern Recog. Worksh. pp. 338–339 (2020) 10
- Schops, T., Schonberger, J.L., Galliani, S., Sattler, T., Schindler, K., Pollefeys, M., Geiger, A.: A multi-view stereo benchmark with high-resolution images and multi-camera videos. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 3260–3269 (2017) 2, 9
- Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M.: Laion-5b: An open largescale dataset for training next generation image-text models. Adv. Neural Inform. Process. Syst. 35, 25278–25294 (2022) 6, 9
- Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgbd images. In: Eur. Conf. Comput. Vis. pp. 746–760. Springer (2012) 9
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: Int. Conf. Mach. Learn. pp. 2256–2265. PMLR (2015) 4
- 63. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020) 4
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Scorebased generative modeling through stochastic differential equations. In: Int. Conf. Learn. Represent. (2021) 4
- Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. Adv. Neural Inform. Process. Syst. 36, 1363–1389 (2023) 3
- Wan, Q., Huang, Z., Kang, B., Feng, J., Zhang, L.: Harnessing diffusion models for visual perception with meta prompts. arXiv preprint arXiv:2312.14733 (2023) 3, 5, 6
- 67. Wang, T., Zhang, T., Zhang, B., Ouyang, H., Chen, D., Chen, Q., Wen, F.: Pretraining is all you need for image-to-image translation. arXiv preprint arXiv:2205.12952 (2022) 4
- Wang, W., Dai, J., Chen, Z., Huang, Z., Li, Z., Zhu, X., Hu, X., Lu, T., Lu, L., Li, H., et al.: Internimage: Exploring large-scale vision foundation models with deformable convolutions. arXiv preprint arXiv:2211.05778 (2022) 9
- Wu, W., Zhao, Y., Chen, H., Gu, Y., Zhao, R., He, Y., Zhou, H., Shou, M.Z., Shen, C.: Datasetdm: Synthesizing data with perception annotations using diffusion models. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Adv. Neural Inform. Process. Syst. vol. 36, pp. 54683–54695. Curran Associates, Inc. (2023) 5
- Xian, K., Shen, C., Cao, Z., Lu, H., Xiao, Y., Li, R., Luo, Z.: Monocular relative depth perception with web stereo data supervision. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 311–320 (2018) 4

- Xu, X., Zhao, H., Vineet, V., Lim, S.N., Torralba, A.: Mtformer: Multi-task learning via transformer and cross-task reasoning. In: Eur. Conf. Comput. Vis. pp. 304–321. Springer (2022) 8
- Yang, L., Kang, B., Huang, Z., Xu, X., Feng, J., Zhao, H.: Depth anything: Unleashing the power of large-scale unlabeled data. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 10371–10381 (2024) 1, 2, 3, 4, 5, 9, 10, 11, 12, 13, 14
- Yin, T., Gharbi, M., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T., Park, T.: One-step diffusion with distribution matching distillation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 6613–6623 (2024) 4
- Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction. In: Int. Conf. Comput. Vis. pp. 5684–5693 (2019) 8
- 75. Yin, W., Zhang, C., Chen, H., Cai, Z., Yu, G., Wang, K., Chen, X., Shen, C.: Metric3d: Towards zero-shot metric 3d prediction from a single image. In: Int. Conf. Comput. Vis. pp. 9043–9053 (2023) 4
- Yin, W., Zhang, J., Wang, O., Niklaus, S., Mai, L., Chen, S., Shen, C.: Learning to recover 3d scene shape from a single image. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 204–213 (2021) 5, 11
- Zavadski, D., Feiden, J.F., Rother, C.: Controlnet-xs: Designing an efficient and effective architecture for controlling text-to-image diffusion models. arXiv preprint arXiv:2312.06573 (2023) 8
- Zhang, C., Yin, W., Wang, B., Yu, G., Fu, B., Shen, C.: Hierarchical normalization for robust monocular depth estimation. Adv. Neural Inform. Process. Syst. 35, 14128–14139 (2022) 5, 11
- 79. Zhang, Q., Chen, Y.: Fast sampling of diffusion models with exponential integrator. arXiv preprint arXiv:2204.13902 (2022) 4
- Zhao, W., Rao, Y., Liu, Z., Liu, B., Zhou, J., Lu, J.: Unleashing text-to-image diffusion models for visual perception. In: Int. Conf. Comput. Vis. pp. 5729–5739 (October 2023) 3, 4, 5, 6, 8, 13, 14
- Zheng, W., Teng, J., Yang, Z., Wang, W., Chen, J., Gu, X., Dong, Y., Ding, M., Tang, J.: Cogview3: Finer and faster text-to-image generation via relay diffusion. arXiv preprint arXiv:2403.05121 (2024) 4
- Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 633–641 (2017) 8