arXiv:2308.02283v1 [cs.CV] 4 Aug 2023

Diffusion-Augmented Depth Prediction with Sparse Annotations

Jiaqi Li
lijiaqi_mail@hust.edu.cn
School of AIA, Huazhong University of Science and Technology

Yiran Wang*
wangyiran@hust.edu.cn
School of AIA, Huazhong University of Science and Technology

Zihao Huang
zihaohuang@hust.edu.cn
School of AIA, Huazhong University of Science and Technology

Jinghong Zheng
deepzheng@hust.edu.cn
School of AIA, Huazhong University of Science and Technology

Ke Xian
ke.xian@ntu.edu.sg
S-Lab, Nanyang Technological University

Zhiguo Cao
zgcao@hust.edu.cn
School of AIA, Huazhong University of Science and Technology

Jianming Zhang
jianmzha@adobe.com
Adobe Research

ABSTRACT

Depth estimation aims to predict dense depth maps. In autonomous driving scenes, sparsity of annotations makes the task challenging. Supervised models produce concave objects due to insufficient structural information. They overfit to valid pixels and fail to restore spatial structures. Self-supervised methods are proposed for the problem. Their robustness is limited by pose estimation, leading to erroneous results in natural scenes. In this paper, we propose a supervised framework termed **D**iffusion-**A**ugmented **D**epth **P**rediction **(DADP)**. We leverage the structural characteristics of diffusion model to enforce depth structures of depth models in a plug-and-play manner. An object-guided integrality loss is also proposed to further enhance regional structure integrality by fetching objective information. We evaluate DADP on three driving benchmarks and achieve significant improvements in depth structures and robustness. Our work provides a new perspective on depth estimation with sparse annotations in autonomous driving scenes.

CCS CONCEPTS

• **Computing methodologies → Scene understanding**.

KEYWORDS

depth prediction, diffusion model, autonomous driving

ACM Reference Format:
Jiaqi Li, Yiran Wang, Zihao Huang, Jinghong Zheng, Ke Xian, Zhiguo Cao, and Jianming Zhang. 2023. Diffusion-Augmented Depth Prediction with Sparse Annotations. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23), October 29-November 3, 2023, Ottawa, ON, Canada.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3581783.3611807

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MM '23, October 29-November 3, 2023, Ottawa, ON, Canada
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0108-5/23/10…$15.00
https://doi.org/10.1145/3581783.3611807

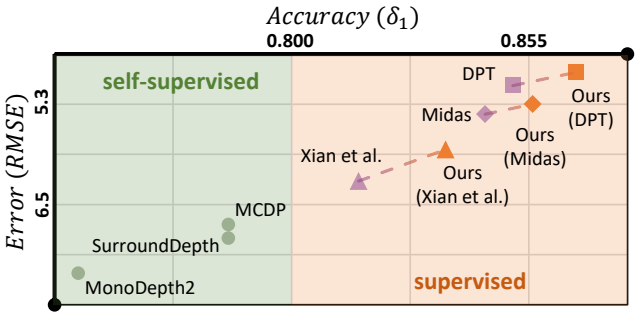

Figure 1: Comparisons with supervised [35, 36, 53] or self-supervised [14, 51, 55] methods on nuScenes [6]. The X-axis represents δ_1. The Y-axis represents RMSE. Our DADP achieves state-of-the-art performance and improvements with different depth predictors in a plug-and-play manner.

1 INTRODUCTION

Depth prediction aims to predict dense depth maps, whether in supervised or self-supervised paradigm. However, sparse depth annotations in driving scenes pose an obstacle for the task.

LiDAR is one primary acquisition equipment for driving datasets but only generates sparse annotations. NuScenes [6], DDAD [16], and KITTI [13] only have 0.24%, 1.85%, and 15.8% pixels with valid ground truth respectively. In contrast, depth models are supposed to predict dense results with both accurate details and integral spatial structures. Due to insufficient structural information, supervised [35, 36, 53] or self-supervised [14, 51, 55] methods produce failure predictions with concave objects, erroneous outcomes, or noticeable artifacts on autonomous driving scenarios [6, 16].

Supervised methods [35, 36, 48, 53, 57] employ various loss functions [10, 26, 28, 36, 47, 53] to measure the discrepancy between output depth and ground truth. However, models fail to acquire sufficient structural information from sparse annotations of driving scenes. They overfit to pixels with valid ground truth and cannot preserve regional structures (*i.e.*, complete shape of objects) and holistic structures. As shown in Fig. 5, supervised models [35, 36, 53] trained on nuScenes [6] produce concave areas on

objects. They even exhibit holistically wrong spatial structures with striped artifacts, overfitting to horizontally-distributed annotations from LiDAR. As proved in Sec. 4.2, these defects cannot be solved by common techniques for overfitting such as data augmentations or weight decay [32]. Previous supervised models and loss [35, 36, 53] show limitations on sparse depth annotations of driving scenes.

To enhance spatial structures, recent works [14, 17, 19, 51, 55] explore self-supervised manner on driving scenes [6, 16]. SurroundDepth [51] employs pseudo labels from Structure-from-Motion [40] to pretrain their model. They utilize pose estimation and photometric loss [12] between six cameras to restore depth structures. MCDP [55] conducts multi-camera prediction by projections between different views. However, self-supervised methods rely on pose estimation [14], which is inaccurate in natural scenes and limits the robustness of those methods. As shown in Fig. 6, they produce erroneous results on night or rainy scenes. Besides, deploying multiple cameras on driving cars [51] is inflexible and costly.

To overcome these challenges, we propose a novel supervised framework with sparse annotations termed **D**iffusion-**A**ugmented **D**epth **P**rediction **(DADP)**. Our method does not rely on pose estimation and multiple cameras, achieving better robustness than self-supervised methods [14, 17, 19, 51, 55], especially for challenging night or rainy scenes. DADP consists of a noise predictor and a depth predictor. The depth predictor can be different supervised single-image depth models. The core task is to enhance depth structures. Recent diffusion models in other tasks [1, 25, 37, 52] showcases favorable structural properties, which can span coherent parts of objects as shown in Fig. 4. To acquire integral spatial structures, we introduce the noise predictor similar to diffusion models [8, 23] but in a plug-and-play manner. To be specific, we add Gaussian noise to input images. The noise predictor is trained to predict noise components. We fuse the structure-aware features from noise predictor and the detail-aware features from depth predictors, predicting depth maps with both accurate details and complete structures. The noise predictor can be adopted to off-the-shelf depth predictors in a plug-and-play manner. Besides, to further improve regional structure integrality of objects, we design our object-guided integrality loss that fetches objective structural information.

Experiments are conducted on prevailing driving benchmarks nuScenes [6], DDAD [16], and KITTI [13]. DADP effectively alleviates concave objects and artifacts produced by supervised depth predictors [35, 36, 53]. Compared with self-supervised methods [14, 51, 55], quantitative and qualitative results prove the robustness of our DADP on challenging driving scenes with glare, reflections, rain, or weak-textured areas. As shown in Fig. 1, DADP achieves state-of-the-art performance over previous supervised or self-supervised methods. We also adopt three different depth predictors [35, 36, 53] and demonstrate the effectiveness of our plug-and-play paradigm. The main contributions can be summarized as follows:

- We present a plug-and-play framework with sparse depth annotations termed **D**iffusion-**A**ugmented **D**epth **P**rediction **(DADP)**.
- We propose a noise predictor to enforce depth structures in different depth predictors, utilizing structural properties of diffusion models to remedy sparse annotations in autonomous driving.
- We design object-guided integrality loss to further enhance the completeness of objects with objective structural information.

2 RELATED WORK

Depth Datasets for Autonomous Driving. Most driving depth datasets are captured by LiDAR and suffer from sparsity of ground truth. NuScenes [6] is a large-scale dataset captured by a sensor suite (with six cameras, one LiDAR, five RADAR, GPS, and IMU). It contains 1,000 driving scenes in Boston and Singapore. NuScenes [6] is quite challenging for its low density of depth annotations (0.24%) and its complex in-the-wild scenes. DDAD [16] is an urban driving dataset captured by six synchronized cameras and a high-resolution Luminar-H2 LiDAR. It is designed for long-range depth estimation in diverse urban scenes. The annotation density in DDAD [16] is 1.85%. KITTI [13] is another widely-used autonomous driving dataset with annotation density of 15.8%. It contains 61 outdoor scenes with binocular views from a driving car.

Supervised Depth Estimation. In recent years, supervised depth models [4, 27, 36, 49, 53] have significantly improved the depth accuracy. DPT [35] utilizes the vision transformer [9] for depth prediction and semantic segmentation. The structure-guided ranking loss is proposed by Xian *et al.* [53] with a novel sampling strategy for learning from pseudo-depth data. Midas [36] adopts a multi-objective learning strategy and trains depth models on mixing datasets. PackNet-SAN [15] improve spatial structural integrity by jointly learning depth estimation and depth completion.

Under sparse annotations of driving scenes, those supervised models produce incomplete objects and artifacts. In this work, we present our novel supervised framework DADP. Thanks to the structure information from our noise predictor and the structure guidance from our object-guided integrality loss, the above-mentioned defects can be solved in a plug-and-play manner.

Self-supervised Depth Estimation. To enhance depth structures on driving scenes [6, 16], prior arts [17, 19, 20, 51, 55] seek for the self-supervised paradigm. They jointly optimize a pose module and a depth module by photometric loss [12]. Monodepth2 [14] introduces the multi-scale sampling strategy and minimum re-projection loss to deal with obscuring situations. The cross-view transformer is proposed by SurroundDepth [51] to improve self-supervised methods on surrounding views. MCDP [55] iteratively optimizes depth results with neighboring cameras for consistent structures.

However, self-supervised methods [17, 19, 20, 51, 55] rely on pose estimation, which is unreliable and significantly limits their robustness. In contrast, our DADP achieves better robustness without pose estimation and multiple cameras.

Diffusion Models. Sohl-Dickstein *et al.* [42] first design the diffusion probabilistic model, learning to invert the diffusion process. Ho *et al.* [23] propose the DDPM framework to generate high-quality image samples. By learning the process of denoising, the diffusion model can generate images with both vivid details and reasonable structures. A series of improvements on model structures [8, 34] and sampling strategies [2, 39, 43] lead to the great success of diffusion models in generation tasks. They can synthesize realistic effects with reasonable spatial relations and structures according to prior conditions [8, 24, 37, 38]. Those results indicate that diffusion models acquire a strong capability of structural representations.

In our work, we leverage the structural information embedded in the diffusion models to enhance regional and holistic spatial structures for depth estimation on sparse driving scenes [6, 13, 16].

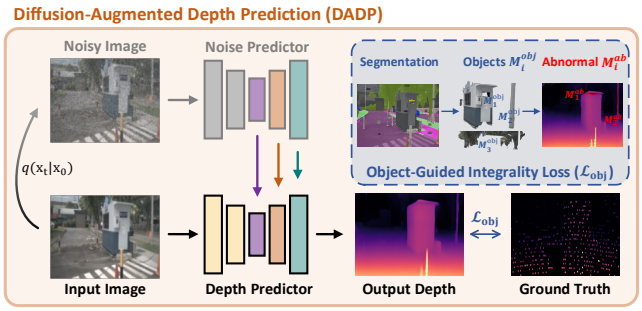

Figure 2: Overview of DADP framework. DADP contains a noise predictor and a depth predictor. The depth predictor can be different single-image depth models. The noise predictor is an Unet [8] to predict the noise component of diffusion step $q(x_t|x_0)$. We fuse structure-aware features from noise predictor and detail-aware features of depth predictor to enhance depth structures with sparse annotations on driving scenes. We also propose object-guided integrality loss to improve objective structural integrity. Our loss guides abnormal regions $\mathbf{M}_i^{ab}$ with incorrect depth variation to normal depth range in different objects. Better viewed when zoomed in.

3 PROPOSED METHOD

3.1 Overview

Fig. 2 showcases the technical pipeline of our Diffusion-Augmented Depth Prediction (DADP) framework. On autonomous driving scenarios only with sparse annotations [6, 16], the key problem is to enforce regional and holistic spatial structures without the unreliable pose estimation in prior arts [14, 19, 20, 51, 55]. DADP contains two main components: a depth predictor and a noise predictor. The depth predictor can be different single-image depth models such as Xian *et al.* [53], Midas [36], and DPT [35]. The noise predictor is a Unet [8, 34] to predict noise components of the diffusion process. The noise predictor and diffusion models [23] are trained to predict noise components and generate noise-free images from the noisy ones. In this way, the noise predictor acquires favorable structural properties as demonstrated in Sec. 4.2. We adopt the noise predictor to restore depth structures. Specifically, we fuse the structure-aware features from the noise predictor and the detail-aware features from the depth predictor by feature fusion modules (FFM) [30, 31]. Ultimately, depth predictors can restore both accurate details and integral spatial structures.

We conduct a two-stage training procedure. First, we train the noise predictor on RGB images with the subtask of unconditional image generation as DDPM [8, 23, 34]. Once the noise predictor is trained, its parameters are fixed. It can be directly adopted to the training of different depth predictors in a plug-and-play manner. The next step is the supervised training for depth predictors. To further improve the regional structural integrality, we design our object-guided integrality loss which fetches the spatial structural information of different objects. To be specific, we utilize a state-of-the-art panoptic segmentation model [7] to segment different objects. We focus on abnormal regions within each object, *i.e.*, concave areas with incorrect depth variation. Our loss guides those abnormal regions to integral depth structures.

During inference, along with the noise predictor, depth predictors predict depth maps with regionally and holistically complete spatial structures, removing the incomplete objects, concave areas, and artifacts produced by previous supervised depth predictors [35, 36, 53] on sparse driving scenes [6, 16]. Without relying on pose estimation, our DADP achieves better robustness than self-supervised frameworks [14, 51, 55] on challenging scenarios.

3.2 Noise Predictor

Preliminaries. Diffusion models [3, 8, 23, 43] are latent variable models for generative tasks, which are trained to denoise Gaussian-blurred images and reverse the diffusion process. If we denote a noise-free input RGB image as $x_0 \in \mathbb{R}^{h\times w\times 3}$ and the maximum diffusion step as T, the diffusion process of diffusion step $t \in \{0, 1, \cdots, T\}$ can be formulated as follows:

$$\begin{aligned} q\left(x_t \mid x_0\right) &:= \mathcal{N}\left(x_t; \sqrt{\bar{\alpha}_t} x_0, \left(1-\bar{\alpha}_t\right) I\right), \\ x_t &= \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1-\bar{\alpha}_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0,1), \end{aligned} \tag{1}$$

where $\alpha_t := 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=0}^{t} \alpha_s$. β is the noise variance schedule as DDPM [23]. $x_t \in \mathbb{R}^{h\times w\times 3}$ represents the resulting high-noise image of diffusion step t. $\mathcal{N}$ denotes Gaussian noise.

As for the denoising process, a noise predictor $\epsilon_\theta(x_t, t)$ is trained to reverse the diffusion process and restore the noise-free x_0. Each denoising step is approximated by a Gaussian distribution:

$$p_\theta\left(x_{t-1} \mid x_t\right) := \mathcal{N}\left(x_{t-1}; \mu_\theta\left(x_t, t\right), \Sigma_\theta\left(x_t, t\right)\right) . \tag{2}$$

The noise predictor $\epsilon_\theta(x_t, t)$ predicts the noise component of step t and obtains the mean value $\mu_\theta(x_t, t)$ of above-mentioned Gaussian distribution by linear combination of x_t. The covariance predictor $\Sigma_\theta\left(x_t, t\right)$ can be fixed [23] or learned [8]. By iterating the denoising process, the noise-free image x_0 can be recovered.

Depth Structure Augmentation. We adopt an Unet as the noise predictor with multi-resolution attention [8] and BigGAN up-/down-sampling [5]. It is trained on RGB images with mean square error (MSE) between the predicted and actual noise components. In Sec. 4.2, we prove that the noise predictor acquires a strong ability of structural representations by denoising RGB images. After the training of noise predictor, we freeze its parameters and utilize its structure-aware features to augment depth structures in a plug-and-play manner. See supplementary for details of the noise predictor.

As shown in Fig. 2, we fuse the structure-aware features from noise predictor and the detail-aware features from depth predictor by feature fusion modules (FFM) [30, 31]. The CNN decoder of the depth predictor gradually improves spatial resolutions and predicts depth results. Specifically, if we denote the block indexes of the noise predictor as b, we leverage structure-aware features $\{t = 50, b = 12\}$, $\{t = 100, b = 8\}$, and $\{t = 150, b = 5, 6, 7\}$ from the noise predictor. Resolutions of the structure-aware features are adjusted to $1/2$, $1/4$, and $1/8$ of input resolution respectively.

The ultimate step is to train depth predictors [35, 36, 53]. Directly training depth predictors on highly sparse driving scenes [6, 16] will produce incomplete objects, concave areas, and artifacts by previous supervised training procedures and loss functions [10, 36, 53]. With

the structure information from noise predictor, our DADP can remove those defects and predict depth maps with regionally and holistically integral spatial structures. The noise predictor can be adapted to different depth predictors [35, 36, 53] in a plug-and-play manner. As for the supervision between predicted depth and ground truth, we adopt the commonly-applied affinity invariant loss [36]. More importantly, to further improve the regional integrality of objects, we propose to use the objective structural information as guidance, which will be illustrated in the next section.

3.3 Object-Guided Integrality Loss

With the structural information from the noise predictor described above, our DADP is capable to restore spatial structures from sparse depth annotations of autonomous driving scenarios. To further improve the integrity of different objects, we propose the object-guided integrality loss as regional structural guidance.

The design of our object-guided integrality loss is based on a simple prior that the depth values inside a certain object should be continuous and smooth. Our loss focuses on the areas with abnormal depth variation inside each object and guides them back to normal depth values. Thus, we need to extract each object region from the input image. In our implementation, we utilize the state-of-the-art Mask2Former [7] to perform panoptic segmentation and extract the object mask of each instance. Here, we denote the object masks as $\mathbf{M}_i^{obj}, i \in \{1, 2, \cdots, K\}$, where K represents the number of objects in the input image.

Specifically, given the noise-free input image $\mathbf{x_0}$, we denote the predicted and ground truth depth as $\hat{\mathbf{d}}$ and $\hat{\mathbf{d}}^*$, which are aligned to zero translation and unit scale as MiDaS [36] to deal with varied scale and shift of different depth predictors [35, 36, 53]. For each object $\mathbf{M}_i^{obj}$, we segment the corresponding region of predicted depth, denoted by $\hat{\mathbf{d}}_i$. The operation can be easily achieved by dot product between object mask $\mathbf{M}_i^{obj}$ and predicted depth $\hat{\mathbf{d}}$. Similarly, We obtain the list $\hat{\mathbf{d}}_i^*$ of valid ground truth values within the object.

The next step is to find the areas with abnormal depth variation inside each object. We consider the normal depth range in object $\mathbf{M}_i^{obj}$ with upper bound $\mathbf{U}_i$ and lower bound $\mathbf{L}_i$ as follows:

$$\begin{aligned} \mathbf{U}_i &= (1+\alpha) * max(max(\hat{\mathbf{d}}_i^*), median(\hat{\mathbf{d}}_i)) \ , \\ \mathbf{L}_i &= (1-\alpha) * min(min(\hat{\mathbf{d}}_i^*), median(\hat{\mathbf{d}}_i)) \ , \end{aligned} \tag{3}$$

where α is the tolerance factor. To be mentioned, for some small objects with no valid ground truth, *i.e.*, $\hat{\mathbf{d}}_i^*$ is empty, we use the median value of $\hat{\mathbf{d}}_i$ with tolerance factor α for $\mathbf{U}_i$ and $\mathbf{L}_i$. In contrast, when pixels with valid ground truth exist in the object, the maximum and minimum value of $\hat{\mathbf{d}}_i^*$ will be adopted.

Our loss encourages depth values inside the object to be continuous and smooth. Pixels with depth values out of the normal depth range will be divided into abnormal areas. However, the segmentation model [7] cannot handle some complex occlusions, *e.g.*, sky regions with tree branches. Pixels that actually do not belong to the certain object should be removed from $\mathbf{M}_i^{obj}$. For this problem, we conduct k-means clustering [21] on the input image $\mathbf{x_0}$ within the object $\mathbf{M}_i^{obj}$. Pixels with the lowest 20% cosine similarity to their corresponding clustering centers are removed. If we denote those pixels to be removed as the mask $\mathbf{M}_i^{occ}$, the final abnormal region $\mathbf{M}_i^{ab}$ as shown in Fig. 2 is presented as follows:

$$\mathbf{M}_i^{ab} = (\hat{\mathbf{d}}_i > \mathbf{U}_i) \cup (\hat{\mathbf{d}}_i < \mathbf{L}_i) - \mathbf{M}_i^{occ} \ . \tag{4}$$

Finally, the abnormal depth values should be guided to the normal depth range. If we denote the pixels within $\mathbf{M}_i^{ab}$ as p, our object-guided integrality loss can be formulated as:

$$\mathcal{L}_{obj}\left(\hat{\mathbf{d}}, \mathbf{U}, \mathbf{L}\right) = \frac{1}{K}\sum_{i=1}^{K}\sum_{p=1}^{N_i} min\left(\left|\hat{\mathbf{d}}_i^{(p)} - \mathbf{U}_i\right|, \left|\hat{\mathbf{d}}_i^{(p)} - \mathbf{L}_i\right|\right) \ , \tag{5}$$

where N_i is the number of pixels in $\mathbf{M}_i^{ab}$. See supplementary for more visual results of the masks in our loss.

3.4 Implementation Details

Noise Predictor. For our noise predictor, the maximum diffusion step T is set to 1000. we resize images to 256×256 and adopt a batch size of 8 for training. We use the Adam optimizer with learning rate $1e-4$.

The mean square error (MSE) between the predicted and actual noise components is used for supervision. The parameters of our noise predictor are fixed after the first stage of training.

Depth Predictor. To prove our plug-and-play manner, we adopt Xian *et al.* [53], DPT-hybrid [35] and Midas-v2 [36] as depth predictors in our experiments. Xian *et al.* [53] is trained from scratch, while DPT [35] and Midas [36] are fine-tuned on nuScenes [6] and DDAD [16] datasets from their pretrained checkpoints.

For the training of depth predictors, following prior arts [15, 51, 55], we utilize the training resolution of 640×352 on nuScenes dataset [6] and 640×384 on DDAD dataset [16]. We train depth predictors for 5 epochs on the nuScenes [6] and 20 epochs on DDAD [16] with a batch size of 16. The initial learning rate is set to $5e-5$ and decreases by $1e-5$ for every five epochs. For KITTI [13] dataset, we follow the same standard training procedure as BTS [29]. Except for our object-guided integrality loss, we adopt the commonly-applied affinity invariant loss [35, 36] $\mathcal{L}_{af}$ between predicted depth and ground truth. See supplementary for more details. The overall loss $\mathcal{L}$ can be expressed as:

$$\mathcal{L} = \mathcal{L}_{af} + \lambda \mathcal{L}_{obj} \ , \tag{6}$$

where the coefficient λ is 0.1. We set the tolerance factor $\alpha = 0.1$.

4 EXPERIMENTS

In this section, we evaluate our Diffusion-Augmented Depth Prediction (DADP) framework on prevailing autonomous driving datasets nuScenes [6], DDAD [16], and KITTI [13]. We first briefly describe the evaluation protocol and datasets in Sec. 4.1. Some experimental results are shown to further expound on our motivations in Sec. 4.2. The quantitative and qualitative comparisons with state-of-the-art approaches are shown in Sec. 4.3. We also conduct ablation studies and prove the effectiveness of our design in Sec. 4.4.

4.1 Evaluation Protocol and Datasets

Evaluation Protocol. For supervised depth models, we compare our DADP with Xian *et al.* [53], Midas [36] , and DPT [35] as different depth predictors, which demonstrates our effective plug-and-play manner. Those models [35, 36, 53] are trained on nuScenes [6]

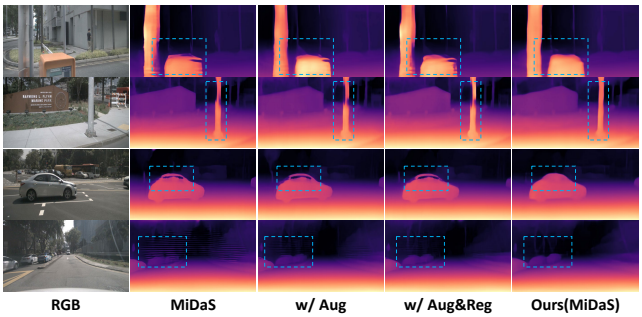

Figure 3: Visual results with data augmentations and weight decay [32]. We train Midas [36] on nuScenes dataset [6] with several data augmentations (Aug) and L_2 regularization (Reg) for weight decay. We highlight regions with prominent difference in dashed rectangular. The incomplete objects, concave areas, and artifacts cannot be solved by those common techniques for overfitting. Better viewed when zoomed in.

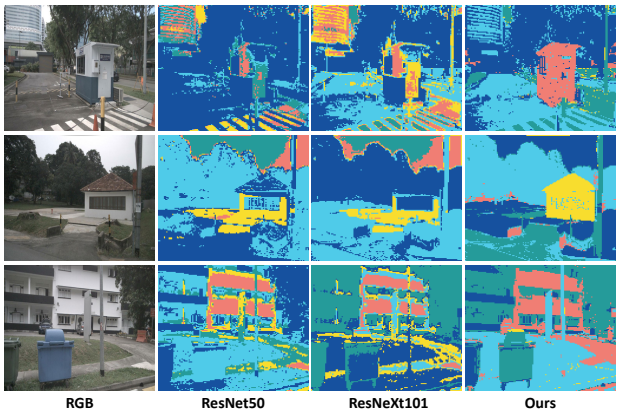

Figure 4: Visual comparisons of features. We visualize feature maps from CNN encoders [22, 54] and the structure-aware features in our noise predictor by k-means [21] clustering (k=5). Our noise predictor shows better representations of spatial structures, which can span coherent parts of objects.

and DDAD [16] with their original training pipeline. On DDAD [16] dataset, we also compare our approach with previous state-of-the-art supervised framework PackNet-SAN [15]. Following PackNet-SAN [15], we evaluate supervised methods [15, 35, 36, 53] on four views of DDAD [16], due to self-occlusion in the other two views.

Self-supervised methods [14, 17, 19, 51, 55] with pose estimation of multiple views are prevailing on autonomous driving datasets [6, 16]. Consequently, we also compare our DADP with them, even though the quantitative comparisons between supervised and self-supervised methods might be unfair. We mainly compare and demonstrate the robustness of our DADP especially in challenging driving scenes with glare and reflections at night, rainy scenes, or weak-textured areas. Our approach does not utilize the unreliable camera poses. Following previous state-of-the-art self-supervised MCDP [55] on DDAD [16] dataset, we compare with self-supervised approaches [15, 35, 36, 53] on six views of DDAD [16].

Datasets. We mainly compare those supervised or self-supervised methods on nuScenes [6] and DDAD [16] datasets. To demonstrate the robustness of our framework, we further evaluate different approaches with daytime and nighttime scenes on nuScenes [6] dataset. For sufficient comparisons and evaluations, we also compare with those methods on KITTI [13] dataset. See supplementary for the experimental results on KITTI [13] dataset.

Evaluation Metrics. We evaluate the performance of different approaches with the commonly-applied depth metrics including *Abs Rel*, *Sq Rel*, *RMSE*, and δ_i $(i = 1, 2, 3)$.

4.2 Motivation Review

Dominant autonomous driving depth datasets NuScenes [6] and DDAD [16] only have 0.24% and 1.85% pixels with depth ground truth. Directly training fully supervised depth estimation models [35, 36, 53] on those highly sparse driving scenes produces concave regions, incomplete objects, or even some artifacts in predicted depth maps. Due to the lack of structural information in the sparse annotations, models could overfit to pixels with valid ground truth and fail to establish regional and holistic spatial structures.

Experimental results prove that those defects cannot be solved by common techniques for overfitting such as data augmentations or weight decay [32]. We train Midas [36] with weight decay [32] and several data augmentations on nuScenes [6] dataset. Specifically, we adopt random cropping, horizontal flipping, and color jittering as data augmentations (Aug). We also apply L_2 regularization (Reg) for weight decay. As shown in Fig. 3, the above-mentioned defects still remain with data augmentations and weight decay [32]. The incomplete objects and concave areas cannot be improved without sufficient structural information and guidance.

Previous supervised learning approaches [15, 35, 36, 53] cannot handle the sparsity problem for autonomous driving. They do not conduct experiments on nuScenes [6] dataset while only PackNet-SAN [15] shows results on DDAD [16] dataset. To establish complete spatial structures on those sparse driving scenes [6, 16], most previous works [14, 17, 19, 51, 55] seek for the self-supervised learning manner with pose estimation in multiple views. However, camera poses are inaccurate and unreliable on many challenging scenes and limit the robustness of those self-supervised approaches [14, 17, 19, 51, 55]. As shown in Fig. 6, previous Monodepth2 [14] and SurroundDepth [51] produce erroneous predictions on night or rainy scenes.

We are devoted to solving this dilemma. We prefer not to rely on the pose estimation and self-supervised manner considering their limited robustness on natural scenes. The key problem is to enforce integral spatial structures under the condition of highly sparse depth annotations. An intuitive idea is to fuse spatial features extracted from RGB images by convolutional encoders [22, 54]. Specifically, we replace the noise predictor with widely-used CNN encoders [22, 54] and fuse the extracted features into the depth predictor. However, concave objects and artifacts are not settled. More quantitative and visual depth comparisons can be found in our supplementary. We visualize the feature maps of the CNN encoders [22, 54] by k-means clustering [21]. As shown in Fig. 4,

Method	NuScenes						DDAD					
	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1 ↑	δ_2 ↑	δ_3 ↑	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Xian *et al.* [53]	0.147	1.375	6.266	0.799	0.913	0.957	0.150	3.281	12.784	0.810	0.914	0.954
PackNet [18]	–	–	–	–	–	–	0.125	2.158	11.245	0.836	0.929	0.962
MiDaS [36]	0.122	1.106	5.485	0.844	0.933	0.964	0.137	2.904	11.870	0.836	0.929	0.964
DPT [35]	0.121	1.074	5.315	0.851	0.934	0.965	0.134	2.630	11.913	0.837	0.930	0.964
PackNet-SAN [15]	–	–	–	–	–	–	0.119	**1.931**	10.852	0.850	0.936	**0.977**
Ours(Xian *et al.*)	0.131	1.141	5.952	0.824	0.924	0.963	0.138	2.297	11.150	0.838	0.934	0.967
Ours(MiDaS)	0.117	1.084	5.370	0.856	0.938	0.967	0.122	2.227	10.406	0.866	0.945	0.970
Ours(DPT)	**0.112**	**1.010**	**5.181**	**0.867**	**0.941**	**0.968**	**0.118**	2.140	**10.130**	**0.870**	**0.946**	0.972

Table 1: Comparisons with supervised depth estimation approaches on nuScenes [6] and DDAD [16]. The first five rows are results of previous state-of-the-art supervised models. We show our results with different depth predictors in the last three rows, which demonstrate the effectiveness of our plug-and-play manner. Following PackNet-SAN [15], we evaluate supervised approaches on four views of DDAD [16]. Best performance is in boldface. Second best is underlined.

Method	NuScenes						DDAD					
	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1 ↑	δ_2 ↑	δ_3 ↑	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1 ↑	δ_2 ↑	δ_3 ↑
Monodepth2 [14] (-M)	0.287	3.349	7.184	0.641	0.845	0.925	0.217	3.641	12.962	0.699	0.877	0.939
PackNet-SfM [17] (-M)	0.309	2.891	7.994	0.547	0.796	0.899	0.234	3.802	13.253	0.672	0.860	0.931
FSM [19]	0.334	2.845	7.786	0.508	0.761	0.894	0.229	4.589	13.520	0.677	0.867	0.936
TransDSSL [20]	–	–	–	–	–	–	0.151	3.591	14.350	–	–	–
SurroundDepth [51]	0.245	3.067	6.835	0.719	0.878	0.935	0.200	3.392	12.270	0.740	0.894	0.947
MCDP [55]	0.237	3.030	6.822	0.719	–	–	0.193	3.111	12.264	0.811	–	–
Ours(Xian *et al.*)	0.131	1.141	5.952	0.824	0.924	0.963	0.146	2.352	10.778	0.826	0.927	0.962
Ours(MiDaS)	0.117	1.084	5.370	0.856	0.938	0.967	0.132	**2.312**	10.087	0.856	0.938	0.966
Ours(DPT)	**0.112**	**1.010**	**5.181**	**0.867**	**0.941**	**0.968**	**0.130**	2.338	**9.994**	**0.860**	**0.939**	**0.966**

Table 2: Comparisons with self-supervised depth estimation approaches on nuScenes [6] and DDAD [16] datasets. The first six rows are results of previous state-of-the-art self-supervised models. Our results with different depth predictors are in the last three rows. Best results are highlighted in bold. Second best is underlined. $(-M)$ indicates occlusion masking [51] in DDAD [16].

detailed RGB information rather than spatial structures are presented by CNNs [22, 54]. For better spatial structures, motivated by recent diffusion models [3, 8, 33, 43, 44], we propose our DADP with the noise predictor. The noise predictor is trained to denoise and generate noise-free images from noisy ones. This subtask can enforce better representations of spatial structures in the structure-aware features as shown in Fig. 4. With the structural information from the noise predictor, our DADP predicts depth results with integral regional and holistic structures, significantly alleviating the incomplete objects and artifacts.

Besides, to further guide integral regional structures of objects, we design the object-guided integrality loss. Our loss focuses on the areas with abnormal depth variation inside a certain object and guides the depth values to normal depth range. Overall, our DADP and object-guided integrality loss significantly improve depth structures on sparse autonomous driving scenarios [6, 16] and obtain better robustness without using pose estimation of multiple views.

4.3 Comparisons with state-of-the-art results

Comparisons with supervised methods. We compare our DADP with previous supervised depth estimation approaches [15, 18, 35, 36, 53]. We present quantitative comparisons in Table 1 and qualitative results in Fig. 5. For PackNet [18] and PackNet-SAN [15], we report their official depth metrics on DDAD dataset [16]. For Xian *et al.* [53], Midas [36], and DPT [35], we train their models on nuScenes and DDAD datasets [6, 16] with their original training pipeline. In Fig. 5, we can observe that previous supervised frameworks [35, 36, 53] produce obvious concave areas, incomplete objects, or even artifacts. With the structural information from the noise predictor, our DADP effectively releases those defects and achieves state-of-the-art performance as shown in Table 1.

Besides, we also fit the three different supervised depth models [35, 36, 53] as depth predictors into our DADP framework. Our approach can effectively deal with different depth predictors in a plug-and-play manner. With Xian *et al.* [53], Midas [36], and DPT [35] as depth predictors on DDAD [16] dataset, our DADP shows 2.8%, 3.0%, and 3.3% improvements of δ_1 respectively. The quantitative metrics cannot fully reflect our improvements due to the sparse ground truth. Both the depth metrics and visualizations demonstrate the effectiveness of our plug-and-play manner.

Comparisons with self-supervised frameworks. Due to the challenging sparse annotations for autonomous driving, most previous depth estimation methods on nuScenes [6] and DDAD [16]

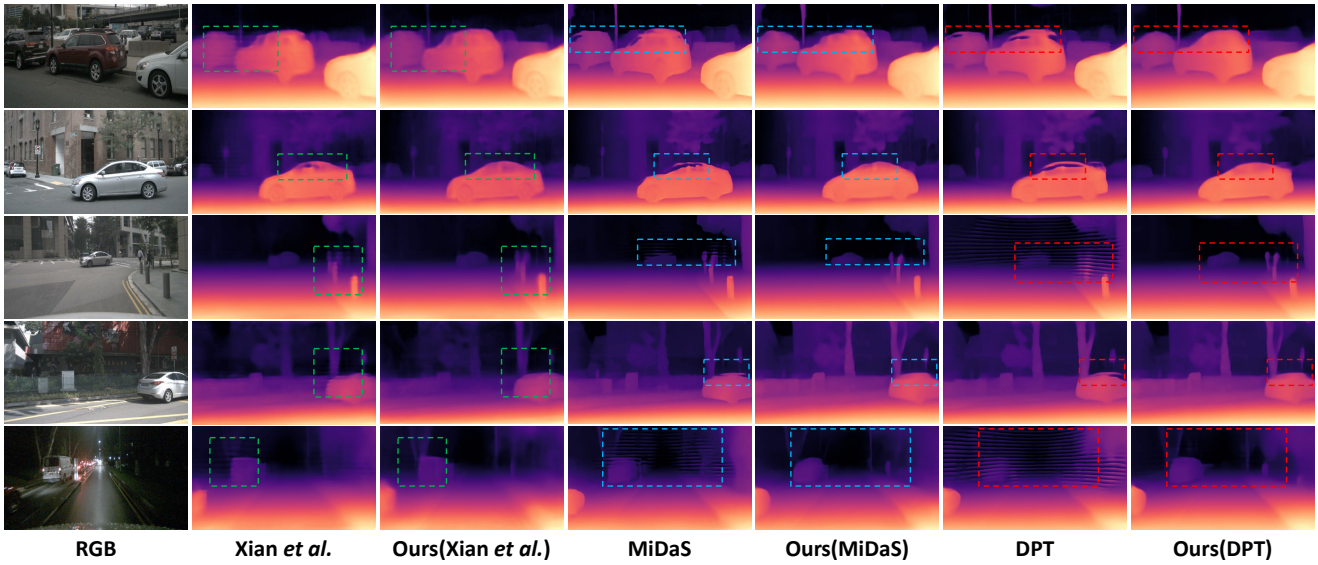

Figure 5: Visual comparisons with supervised depth prediction methods on nuScenes dataset [6]. We adopt supervised models [35, 36, 53] as different depth predictors in our DADP framework. Our approach effectively releases concave areas, incomplete objects, and artifacts produced by depth predictors. Regions with prominent difference are highlighted in dashed rectangular.

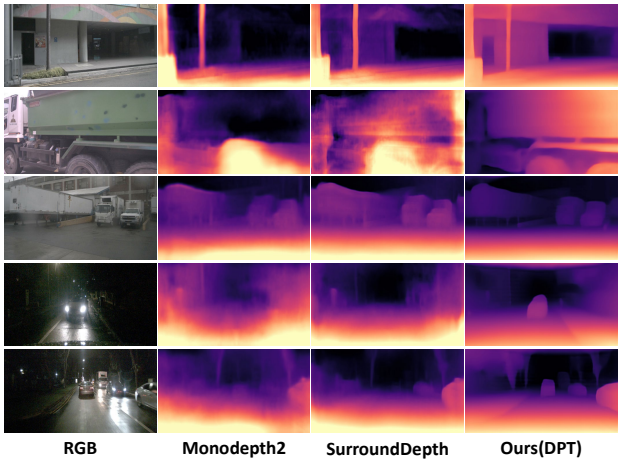

Figure 6: Visual comparisons with self-supervised methods [14, 51]. Without relying on pose estimation, our DADP shows significantly better robustness in challenging scenes with glare and reflections at night, rainy scenes, or weak-textured areas. Better viewed when zoomed in.

Table 3: Comparisons on daytime and nighttime scenes of nuScenes dataset [6]. We evaluate both supervised and self-supervised approaches [14, 35, 36, 51]. Our approach shows better robustness on both daytime and nighttime scenarios.

Method	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1↑	δ_2↑	δ_3↑
Daytime Scenes						
Monodepth2 [14]	0.263	2.772	7.061	0.660	0.854	0.928
SurroundDepth [51]	0.236	3.002	6.754	0.738	0.889	0.940
MiDaS [36]	0.118	1.176	5.212	0.859	0.937	0.967
DPT [35]	0.115	1.049	5.269	0.863	0.938	0.967
Ours(MiDaS)	0.111	1.049	5.292	0.871	0.943	0.969
Ours(DPT)	**0.106**	**0.984**	**5.120**	**0.879**	**0.946**	**0.970**
Nighttime Scenes						
Monodepth2 [14]	0.585	18.715	12.119	0.484	0.732	0.851
SurroundDepth [51]	0.330	3.638	7.563	0.542	0.784	0.890
MiDaS [36]	0.198	1.699	6.197	0.697	0.871	0.940
DPT [35]	0.177	1.297	**5.727**	0.743	0.891	0.948
Ours(MiDaS)	0.182	1.397	6.068	0.722	0.884	0.948
Ours(DPT)	**0.166**	**1.243**	5.731	**0.757**	**0.896**	**0.952**

datasets are in the self-supervised paradigm with pose estimation. We compare our DADP with previous self-supervised frameworks [14, 17, 19, 20, 51, 55] in Table 2 and Fig. 6. We know that it might be unfair to compare supervised approaches with self-supervised methods. However, it is meaningful to explore supervised learning frameworks on sparse driving scenes. The robustness of self-supervised methods is limited due to the unreliable camera poses. Accurate pose estimation itself is a challenging task especially on in-the-wild images with glare and reflections at night, rainy scenes, or weak-textured areas. Whether compared with supervised or self-supervised methods, our DADP shows better qualitative and quantitative performance. Compared with self-supervised approaches, our DADP does not utilize camera poses and achieves better robustness on challenging driving scenes as shown in Fig. 6.

Robustness on Challenging Driving Scenes. To further demonstrate our robustness on challenging driving scenes, we evaluate previous supervised and self-supervised approaches [14, 35, 36, 51] with daytime and nighttime scenes respectively on nuScenes

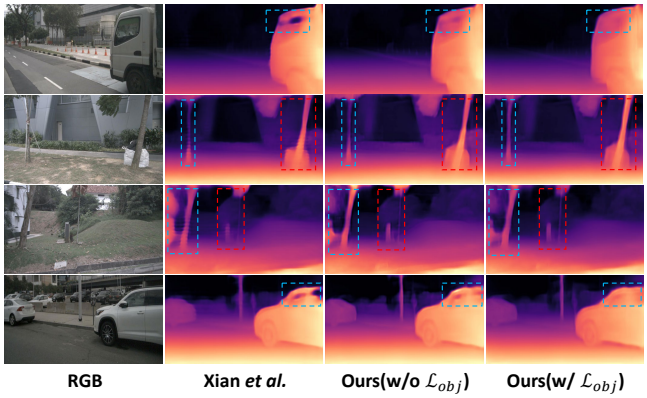

Figure 7: Qualitative ablation of DADP framework. The noise predictor enforces integral structures and removes stripped artifacts in the depth predictor [53]. The object-guided integrality loss further improves regional structural completeness of objects. Better viewed when zoomed in.

Table 4: Ablation of our noise predictor and object-guided integrality loss with Xian *et al.* [53] as the depth predictor.

Method	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1↑	δ_2↑	δ_3↑
Xian *et al.* [53]	0.183	1.375	6.266	0.799	0.913	0.957
Ours ($w/o\mathcal{L}_{obj}$)	0.140	1.292	6.001	0.820	0.921	0.960
Ours ($w/\mathcal{L}_{obj}$)	**0.132**	**1.141**	**5.951**	**0.824**	**0.925**	**0.963**

dataset [6]. The results are shown in Table 3. Compared with state-of-the-art self-supervised SurroundDepth [51] with pose estimation, our DADP showcases 21% δ_1 improvements in the challenging nighttime scenes. Compared with the supervised models, our method also outperforms DPT [35] by 4.2% in *Sq Rel* for nighttime. Whether compared with previous supervised or self-supervised approaches, our framework shows better robustness on challenging autonomous driving scenes with highly sparse depth annotations.

4.4 Ablation Studies

Effectiveness of DADP Framework. We ablate the noise predictor and our object-guided integrality loss to demonstrate the effectiveness of our design. In this experiment, we adopt Xian *et al.* [53] as the depth predictor on nuScenes dataset [6]. Quantitative results are shown in Table 4. Thanks to the structural information from the noise predictor, our DADP achieves 2.1% improvement on δ_1 compared with the depth predictor. By adding our object-guided integrality loss for supervision, the depth accuracy further improves. To be mentioned, the quantitative metrics with sparse ground truth cannot fully reflect our improvements in depth structures.

Visual comparisons are shown in Fig. 7. The noise predictor enforces spatial structures and removes the stripped artifacts in the depth predictor [53]. Meanwhile, our object-guided integrality loss can further improve the regional object completeness, for example, removing the concave areas on the window of the car. These structural improvements could not be presented by the depth metrics. Those concave areas might have few pixels with valid ground truth.

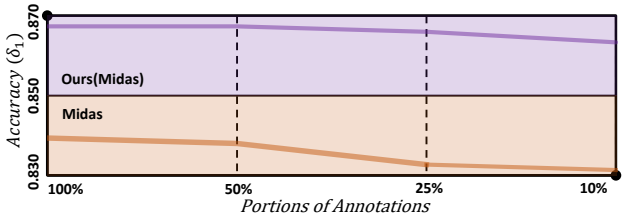

Figure 8: Ablation of annotation density on DDAD dataset. DDAD [16] dataset has 1.85% pixels with valid ground truth. We train our DADP with Midas [36] as the depth predictor using different portions of the valid pixels. Our DADP framework shows strong robustness on high annotation sparsity.

Table 5: Ablation of diffusion steps on nuScenes [6] dataset.

Diffusion Steps	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1↑	δ_2↑	δ_3↑
(0, 50, 100)	0.135	1.207	5.964	0.820	0.923	0.962
(100, 150, 200)	0.137	1.214	6.010	0.817	0.922	0.962
(50, 100, 150)	**0.132**	**1.141**	**5.951**	**0.824**	**0.925**	**0.963**

Density of Depth Annotations. Different autonomous driving datasets [6, 13, 16] possess different annotation density. To demonstrate the robustness of our framework on high sparsity, we randomly sample several portions of valid pixels for training on DDAD dataset [16]. As shown in Fig. 8, with lower annotation density, our DADP shows stronger robustness than the depth predictor [36]. When the portion decreases to the extremely low 10%, *i.e.*, only with 0.185% pixels with valid annotations, our DADP still showcases 3.2% improvements on δ_1 over Midas [36] as the depth predictor.

Diffusion Steps. In this experiment, we ablate different diffusion steps on nuScenes [6] dataset. The experimental results are shown in Table 5. With larger diffusion steps, features become more abstract and structural while remaining less details. Structure-aware features of the noise predictor should achieve a balance of structural and detailed information. As a consequence, we choose the best-performing diffusion steps of (50, 100, 150) in our experiments.

5 CONCLUSION

Previous supervised depth methods lead to concave areas and artifacts on sparse driving scenes. In this paper, we propose a supervised framework termed Diffusion-Augmented Depth Prediction (DADP). With the structural information from our noise predictor and structural guidance from our object-guided integrality loss, DADP effectively alleviates those defects in a plug-and-play manner. Compared with the prevailing self-supervised methods, our DADP is more robust on challenging scenes without relying on pose estimation. Further analysis shows the efficacy of our design and proves that DADP can adapt to varied annotation density. Our work shows strong effectiveness in the field of autonomous driving.

6 ACKNOWLEDGMENTS

This work was funded by the National Natural Science Foundation of China under Grant No. U1913602 and was supported by Adobe.

REFERENCES

[1] Tomer Amit, Eliya Nachmani, Tal Shaharbany, and Lior Wolf. 2021. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390* (2021).

[2] Fan Bao, Chongxuan Li, Jun Zhu, and Bo Zhang. 2022. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. (2022).

[3] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khrulkov, and Artem Babenko. 2022. Label-Efficient Semantic Segmentation with Diffusion Models. In *International Conference on Learning Representations*.

[4] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. 2021. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 4009–4018.

[5] Andrew Brock, Jeff Donahue, and Karen Simonyan. 2019. Large scale GAN training for high fidelity natural image synthesis. (2019).

[6] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. 2020. nuScenes: A Multimodal Dataset for Autonomous Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 11618–11628.

[7] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. 2022. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 1290–1299.

[8] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* 34 (2021), 8780–8794.

[9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.

[10] David Eigen, Christian Puhrsch, and Rob Fergus. 2014. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, Vol. 27. 2366–2374.

[11] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. 2018. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2002–2011.

[12] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. 2016. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European Conference on Computer Vision (ECCV)*. Springer, 740–756.

[13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* 32, 11 (2013), 1231–1237.

[14] C. Godard, O. Aodha, M. Firman, and G. Brostow. 2019. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 3828–3838.

[15] Vitor Guizilini, Rares Ambrus, Wolfram Burgard, and Adrien Gaidon. 2021. Sparse Auxiliary Networks for Unified Monocular Depth Prediction and Completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 11078–11088.

[16] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 2020. 3D Packing for Self-Supervised Monocular Depth Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2482–2491.

[17] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 2020. 3D Packing for Self-Supervised Monocular Depth Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2482–2491.

[18] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 2020. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2485–2494.

[19] Vitor Guizilini, Igor Vasiljevic, Rares Ambrus, Greg Shakhnarovich, and Adrien Gaidon. 2022. Full Surround Monodepth From Multiple Cameras. *IEEE Robotics and Automation Letters* 7, 2 (2022), 5397–5404.

[20] Daechan Han, Jeongmin Shin, Namil Kim, Soonmin Hwang, and Yukyung Choi. 2022. Transdssl: Transformer based depth estimation via self-supervised learning. *IEEE Robotics and Automation Letters* 7, 4 (2022), 10969–10976.

[21] John A Hartigan and Manchek A Wong. 1979. Algorithm AS 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)* 28, 1 (1979), 100–108.

[22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778.

[23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* 33 (2020), 6840–6851.

[24] Jonathan Ho and Tim Salimans. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022).

[25] Jonathan Ho, Tim Salimans, Alexey A Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. [n. d.]. Video Diffusion Models. In *Advances in Neural Information Processing Systems*.

[26] Junjie Hu, Mete Ozay, Yan Zhang, and Takayuki Okatani. 2019. Revisiting Single Image Depth Estimation: Toward Higher Resolution Maps With Accurate Object Boundaries. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 1043–1051.

[27] Andrey Ignatov, Grigory Malivenko, David Plowman, Samarth Shukla, and Radu Timofte. 2021. Fast and Accurate Single-Image Depth Estimation on Mobile Devices, Mobile AI 2021 Challenge: Report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. 2545–2557.

[28] Andrey Ignatov, Grigory Malivenko, Radu Timofte, Lukasz Treszczotko, Xin Chang, Piotr Ksiazek, Michal Lopuszynski, Maciej Pioro, Rafal Rudnicki, Maciej Smyl, et al. 2022. Efficient single-image depth estimation on mobile devices, mobile AI & AIM 2022 challenge: report. In *European Conference on Computer Vision*. Springer, 71–91.

[29] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. 2019. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326* (2019).

[30] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. 2017. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 1925–1934.

[31] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2117–2125.

[32] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. (2019).

[33] Zhaoyang Lyu, Zhifeng Kong, Xudong Xu, Liang Pan, and Dahua Lin. 2021. Differentiable Diffusion for Dense Depth Estimation From Multi-View Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 8912–8921.

[34] Alexander Quinn Nichol and Prafulla Dhariwal. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*. PMLR, 8162–8171.

[35] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. 2021. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 12179–12188.

[36] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. 2020. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence* 44, 03 (2020), 1623–1637.

[37] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10684–10695.

[38] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. 2022. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*. 1–10.

[39] Tim Salimans and Jonathan Ho. 2022. Progressive Distillation for Fast Sampling of Diffusion Models. In *International Conference on Learning Representations*.

[40] Johannes Lutz Schönberger and Jan-Michael Frahm. 2016. Structure-from-Motion Revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 4104–4113.

[41] Chang Shu, Kun Yu, Zhixiang Duan, and Kuiyuan Yang. 2020. Feature-Metric Loss for Self-Supervised Learning of Depth and Egomotion. In *European Conference on Computer Vision (ECCV)*. Springer, 572–588.

[42] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*. PMLR, 2256–2265.

[43] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2020. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*.

[44] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2020. Score-based generative modeling through stochastic differential equations. (2020).

[45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, Vol. 30.

[46] Ruoyu Wang, Zehao Yu, and Shenghua Gao. 2022. PlaneDepth: Plane-Based Self-Supervised Monocular Depth Estimation. *arXiv preprint arXiv:2210.01612* (2022).

[47] Yiran Wang, Xingyi Li, Min Shi, Ke Xian, and Zhiguo Cao. 2021. Knowledge Distillation for Fast and Accurate Monocular Depth Estimation on Mobile Devices. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. 2457–2465.

[48] Yiran Wang, Zhiyu Pan, Xingyi Li, Zhiguo Cao, Ke Xian, and Jianming Zhang. 2022. Less is More: Consistent Video Depth Estimation with Masked Frames Modeling. In *Proceedings of the 30th ACM International Conference on Multimedia* (Lisboa, Portugal) *(MM '22)*. Association for Computing Machinery, New York, NY, USA, 6347–6358. https://doi.org/10.1145/3503161.3547978

[49] Yiran Wang, Min Shi, Jiaqi Li, Zihao Huang, Zhiguo Cao, Jianming Zhang, Ke Xian, and Guosheng Lin. 2023. Neural Video Depth Stabilizer. *arXiv preprint arXiv:2307.08695* (2023).

[50] Jamie Watson, Oisin Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. 2021. The Temporal Opportunist: Self-Supervised Multi-Frame Monocular Depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 1164–1174.

[51] Yi Wei, Linqing Zhao, Wenzhao Zheng, Zheng Zhu, Yongming Rao, Guan Huang, Jiwen Lu, and Jie Zhou. 2023. SurroundDepth: Entangling Surrounding Views for Self-Supervised Multi-Camera Depth Estimation. In *Proceedings of The 6th Conference on Robot Learning*, Vol. 205. 539–549.

[52] Julia Wolleb, Robin Sandkühler, Florentin Bieder, Philippe Valmaggia, and Philippe C Cattin. 2022. Diffusion models for implicit image segmentation ensembles. In *International Conference on Medical Imaging with Deep Learning*. PMLR, 1336–1348.

[53] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. 2020. Structure-Guided Ranking Loss for Single Image Depth Prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 608–617.

[54] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. 2017. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 1492–1500.

[55] Jialei Xu, Xianming Liu, Yuanchao Bai, Junjun Jiang, Kaixuan Wang, Xiaozhi Chen, and Xiangyang Ji. 2022. Multi-Camera Collaborative Depth Prediction via Consistent Structure Estimation. In *Proceedings of the 30th ACM International Conference on Multimedia* (Lisboa, Portugal) *(MM '22)*. Association for Computing Machinery, New York, NY, USA, 2730–2738. https://doi.org/10.1145/3503161.3548394

[56] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. 2019. Enforcing geometric constraints of virtual normal for depth prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 5684–5693.

[57] Weihao Yuan, Xiaodong Gu, Zuozhuo Dai, Siyu Zhu, and Ping Tan. 2022. NeWCRFs: Neural Window Fully-connected CRFs for Monocular Depth Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 3916–3925.

A MORE DETAILS ON DADP FRAMEWORK

A.1 The Noise Predictor

We adopt Unet as the noise predictor with multi-resolution attention [8] and BigGAN residual blocks [5]. Here we specify the details.

As shown in Fig. 9, the noise predictor downsamples the input noisy image to 1/32 of the original resolution. For each resolution, two Resblocks [22] are used for feature extraction. Linear layers are used in these blocks to embed the diffusion step t. We utilize multi-head self-attention [45] on resolutions of 32×32, 16×16, 8×8 to improve the representation ability. Besides, the simple up-sampling and down-sampling between different resolutions are replaced by the residual blocks in BigGAN [5], which can reduce information loss and improve the generation quality.

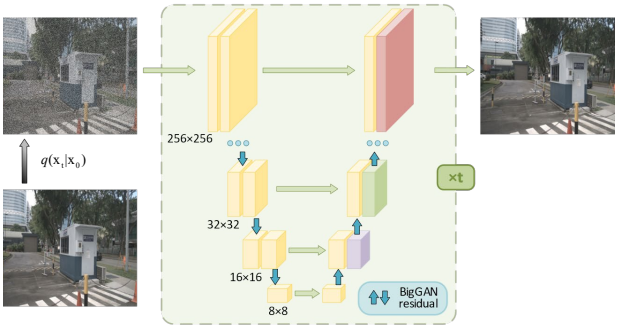

Figure 9: Architecture of the noise predictor. Some resolutions are omitted for simplicity. Blue arrows indicate the residual blocks in BigGAN [5] for up-/down-sampling. Two Resblocks [22] are used for each resolution.

A.2 Object-guided Integrality Loss

Here we explain different masks in the object-guided integrality loss. We visualize $\mathbf{M}^{obj}$, $\mathbf{M}^{ab}$, and $\mathbf{M}^{occ}$ for interpretation in Fig. 10. The indices i, j, and k represent different objects in the input image. $\mathbf{M}^{ab}$ can effectively mark the abnormal depth areas in the certain object $\mathbf{M}^{obj}$. For the regions with complex occlusions or segmentation errors, such as the region beneath the $\mathbf{M}^{obj}_k$ in the third row, we also exclude them as the $\mathbf{M}^{occ}_k$ by k-means clustering [21].

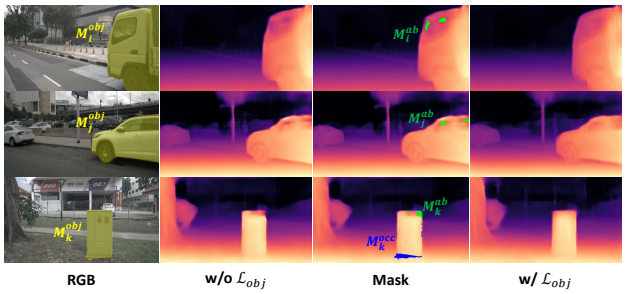

Figure 10: Masks in object-guided integrality loss. The yellow areas in the rgb images represent the object mask $\mathbf{M}^{obj}$ obtained by segmentation model [7]. The green areas $\mathbf{M}^{ab}$ in the third column indicate the regions with abnormal depth variations. The blue area $\mathbf{M}^{occ}$ represents the excluded pixels with segmentation errors by k-means clustering [21]. The indices i, j, k denote specific objects in the three images.

A.3 Affinity Invariant Loss

For the training of the depth predictor [35, 36, 53], we also use the affinity invariant loss [35, 36] in addition to the proposed object-guided integrality loss. Here we briefly present the details. If we denote the original predicted depth and ground truth as $\mathbf{d}$ and $\mathbf{d}^*$, the scale $s(\mathbf{d})$ and shift $t(\mathbf{d})$ of prediction $\mathbf{d}$ can be obtained as:

$$t(\mathbf{d}) = median(\mathbf{d}), s(\mathbf{d}) = \frac{1}{M}\sum_{i=1}^{M}|\mathbf{d}_i - t(\mathbf{d}_i)|, \tag{7}$$

where M denotes the number of pixels. $t(\mathbf{d}^*)$ and $s(\mathbf{d}^*)$ can also be calculated similarly. Then the prediction $\mathbf{d}$ and the ground truth $\mathbf{d}^*$ are aligned to zero translation and unit scale:

$$\hat{\mathbf{d}} = \frac{\mathbf{d} - t(\mathbf{d})}{s(\mathbf{d})}, \hat{\mathbf{d}}^* = \frac{\mathbf{d}^* - t(\mathbf{d}^*)}{s(\mathbf{d}^*)}. \tag{8}$$

$\hat{\mathbf{d}}$ and $\hat{\mathbf{d}}^*$ represent the aligned prediction and ground truth respectively. The affinity invariant loss $\mathcal{L}_{af}$ can be formulated as:

$$\mathcal{L}_{af}\left(\hat{\mathbf{d}}, \hat{\mathbf{d}}^*\right) = \frac{1}{2M}\sum_{i=1}^{M}\left|\hat{\mathbf{d}}_i - \hat{\mathbf{d}}^*_i\right|, \tag{9}$$

B DEPTH ESTIMATION METRICS

We adopt the commonly-applied depth metrics defined as follows:

- **Absolute relative error (Abs Rel):** $\frac{1}{|M|}\sum_{d\in M}|d - d^*|/d^*$;
- **Square relative error (Sq Rel):** $\frac{1}{|M|}\sum_{d\in M}\|d - d^*\|^2/d^*$
- **Root mean square error (RMSE):** $\sqrt{\frac{1}{|M|}\sum_{d\in M}\|d - d^*\|^2}$;
- **Accuracy with threshold t:** Percentage of d_i such that $max(\frac{d_i}{d_i^*}, \frac{d_i^*}{d_i}) = \delta < t \in \left[1.25, 1.25^2, 1.25^3\right]$,

where M denotes numbers of pixels with valid depth annotation, d_i and d_i^* are estimated and ground truth depth of pixel i respectively.

C MORE RESULTS FOR MOTIVATION REVIEW.

Data Augmentation and Weight Decay. In Sec. 4.2 of main paper, we illustrate that data augmentations and weight decay (*e.g.*, regularization) cannot solve the incomplete objects, concave areas, and artifacts for overfitting to sparse valid pixels. The related visual comparisons are given in the Fig. 3 of the main paper. Here we provide the corresponding quantitative metrics in Table 6.

We find that data augmentations reduce the depth accuracy for the reason that those augmentations change the data distribution of nuScenes dataset [6]. The random cropping sometimes leads to even sparser depth supervision. Our method improves the depth accuracy and enhances the spatial structural integrity of the predicted depth maps by fusing structure-aware features from the noise predictor.

Features of CNN Encoders. In Sec. 4.2 of the main paper, the noise predictor is replaced by widely-used CNN encoders [22, 54]. We fuse the extracted features into the depth predictor. The visualization of those features in Fig. 4 of the main paper demonstrates the better representations of spatial structures for our noise predictor. Here we report visual and quantitative results of this experiment.

Table 8: Comparisons with state-of-the-art methods on KITTI dataset [13]. The first five rows contain results of self-supervised methods, while the last seven rows contain supervised approaches. Best performance is in boldface.

Method	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1↑	δ_2↑	δ_3↑
Monodepth2 [14]	0.115	0.882	4.701	0.879	0.961	0.982
Packnet-SFM [17]	0.107	0.802	4.538	0.889	0.962	0.981
Shu *et al.* [41]	0.088	0.712	4.137	0.915	0.965	0.982
PlaneDepth [46]	0.083	0.533	3.919	0.913	0.969	0.985
ManyDepth [50]	0.087	0.685	4.142	0.920	0.968	0.983
DORN [11]	0.072	–	2.626	0.932	0.984	0.994
VNL [56]	0.072	–	3.258	0.938	0.990	0.998
MiDaS [36]	0.069	0.280	3.006	0.949	0.991	0.998
Ours(MiDaS)	0.065	0.253	2.832	0.955	0.992	0.998
BTS [29]	0.059	0.241	2.756	0.956	0.993	0.998
DPT [35]	0.062	**0.222**	**2.573**	0.959	0.995	0.999
Ours(DPT)	**0.059**	0.230	2.661	**0.965**	**0.995**	**0.999**

Table 6: Comparisons of data augmentations (Aug) and weight decay (Reg) on nuScenes dataset [6].

Method	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1↑	δ_2↑	δ_3↑
Midas [36]	0.122	1.106	5.485	0.844	0.933	0.964
w/ Reg	0.126	1.154	5.709	0.840	0.930	0.963
w/ Aug	0.141	1.334	5.696	0.816	0.917	0.957
w/ Aug&Reg	0.141	1.339	5.705	0.815	0.919	0.958
Ours(MiDaS)	**0.117**	**1.084**	**5.370**	**0.856**	**0.938**	**0.967**

Table 7: Comparisons with features of different CNN encoders on nuScenes [6]. Best performance is in boldface.

Method	Abs Rel↓	Sq Rel↓	RMSE↓	δ_1↑	δ_2↑	δ_3↑
Xian *et al.* [53]	0.183	1.375	6.266	0.799	0.913	0.957
w/ ResNet50	0.142	1.269	6.054	0.808	0.917	0.960
w/ ResNeXt101	0.140	1.277	**5.918**	0.813	0.920	0.960
w/ noise predictor	**0.132**	**1.141**	5.951	**0.824**	**0.925**	**0.963**

As shown in Fig. 11 and Table 7, along with the Fig. 4 in the main paper, we showcase the effectiveness of introducing structure-aware features from the noise predictor and diffusion model. The results demonstrate that the structure-aware features can significantly improve the spatial structural integrality of output depth.

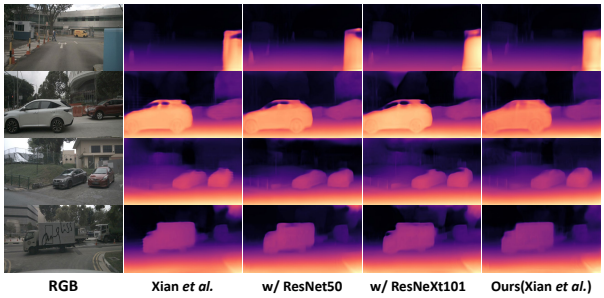

Figure 11: Visual comparisons with CNN encoders [22, 54]. Our DADP effectively augments spatial structural integrality of depth predictions by noise predictor and diffusion model.

D EXPERIMENTS ON KITTI DATASET

In Sec. 4.1 of the main paper, we mentioned that we also train our DADP on KITTI dataset [13] for sufficient comparisons and evaluations, following the train/test split as Eigen *et al.* [10, 29].

Quantitative comparisons in Table 8 show that our DADP improves the depth accuracy than the depth predictors Midas [36] and DPT [35] on KITTI dataset [13].

E MORE QUALITATIVE DEPTH RESULTS

Due to the sparsity of the depth annotations on autonomous driving scenarios [6, 13, 16], the quantitative metrics cannot fully reflect our improvements. Here we show more visual comparisons to demonstrate the effectiveness of our approach, especially for better spatial structural completeness of the predicted depth maps.

The visual results on nuScenes [6], DDAD [16], and KITTI [13] datasets are shown in Fig. 12, Fig. 13, and Fig. 14 respectively. In Fig. 12, we compare our DADP with different depth predictors [35, 36] on nuScenes dataset [6] to demonstrate our plug-and-play manner. Comparisons with state-of-the-art supervised or self-supervised methods on DDAD dataset [16] are shown in Fig. 13. The proposed DADP framework shows significant improvements than the previous state-of-the-art supervised framework PackNet-SAN [15] on DDAD dataset [16]. Visual comparisons with prior arts on KITTI [13] dataset are also shown in Fig. 14 for sufficient evaluations.

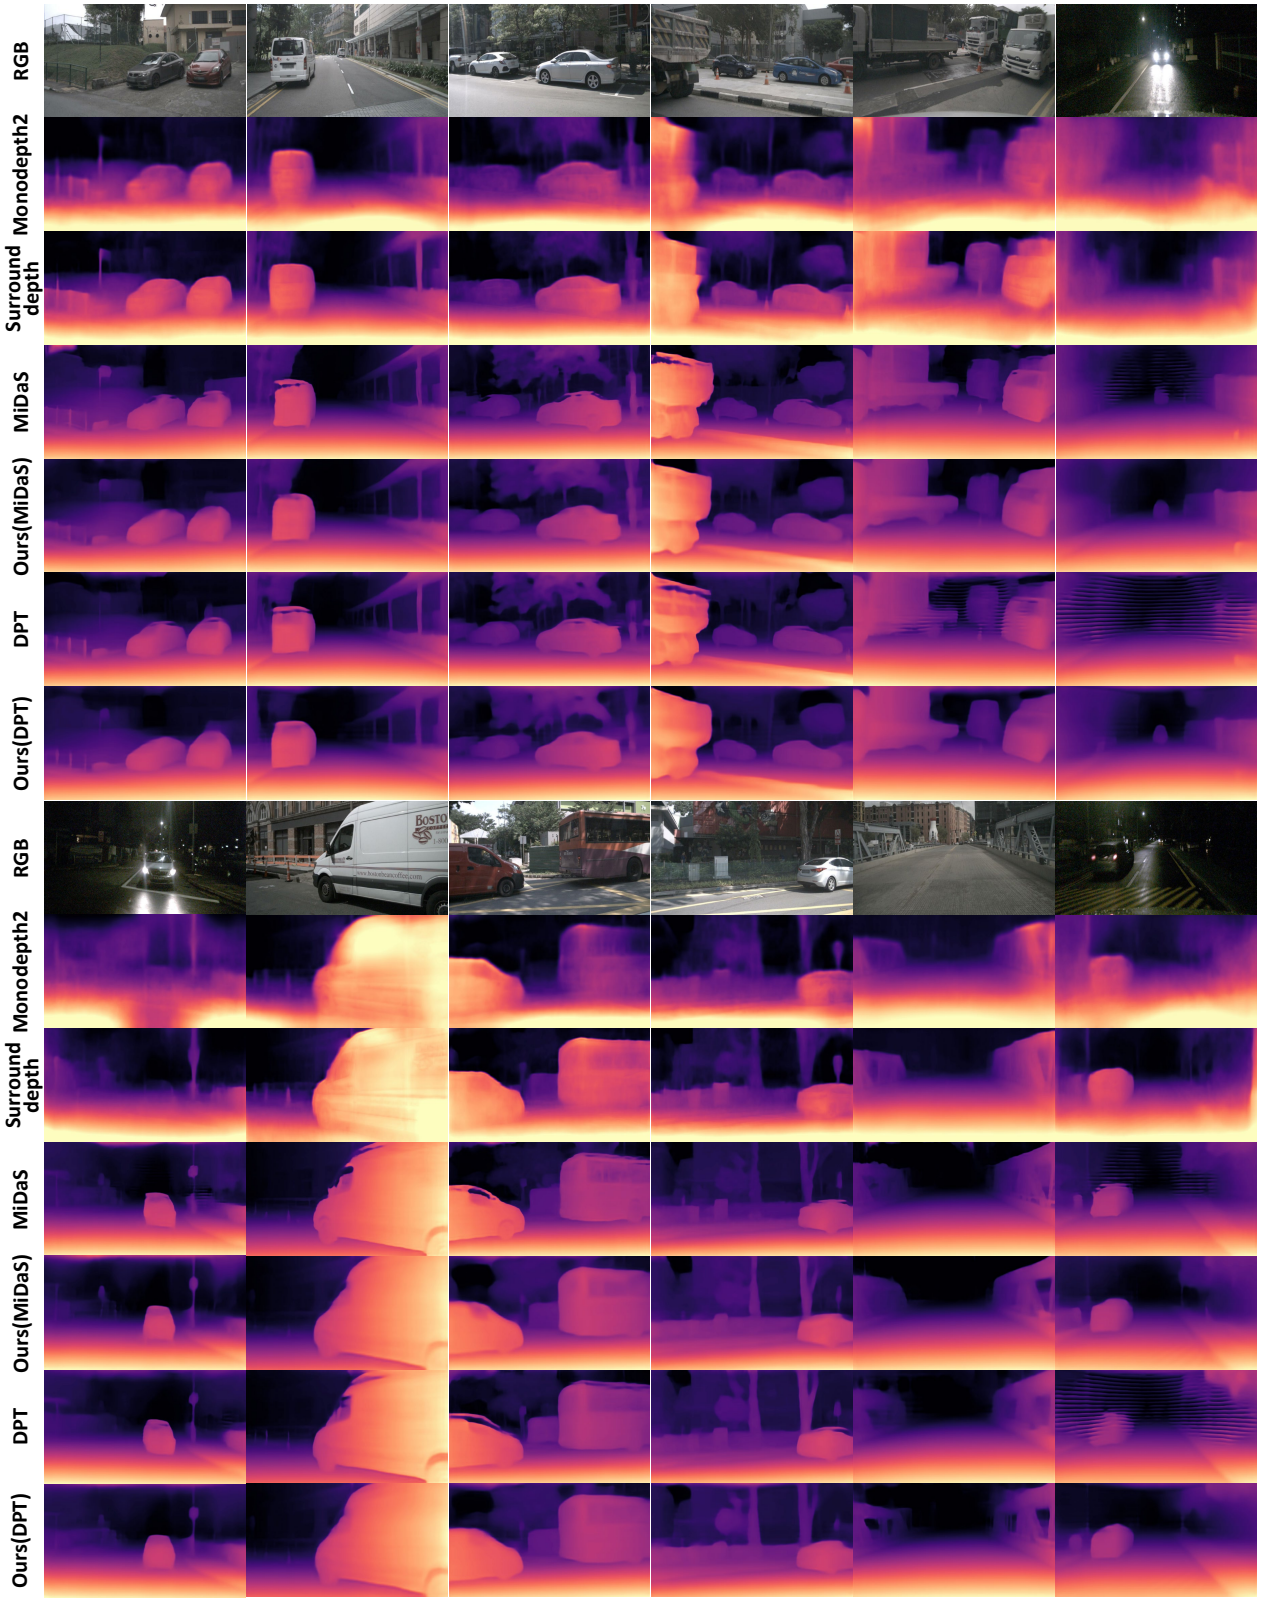

Figure 12: Visual comparisons on nuScenes dataset [6]. we compare our DADP framework with different depth predictors [35, 36] and state-of-the-art approaches [14, 51]. The results demonstrate the effectiveness of our plug-and-play manner.

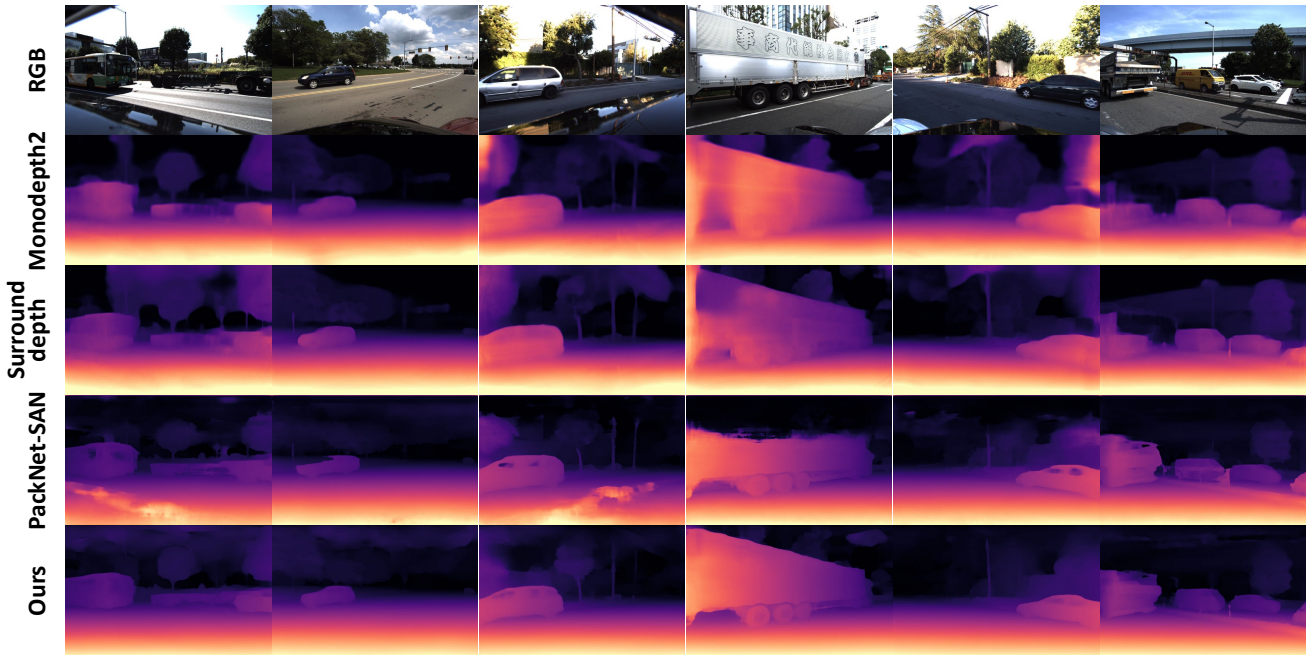

Figure 13: Visual comparisons on DDAD [16]. We compare our DADP with state-of-the-art self-supervised [14, 51] and supervised [15] approaches. Our method shows significant improvements than previous state-of-the-art supervised PackNet-SAN [15].

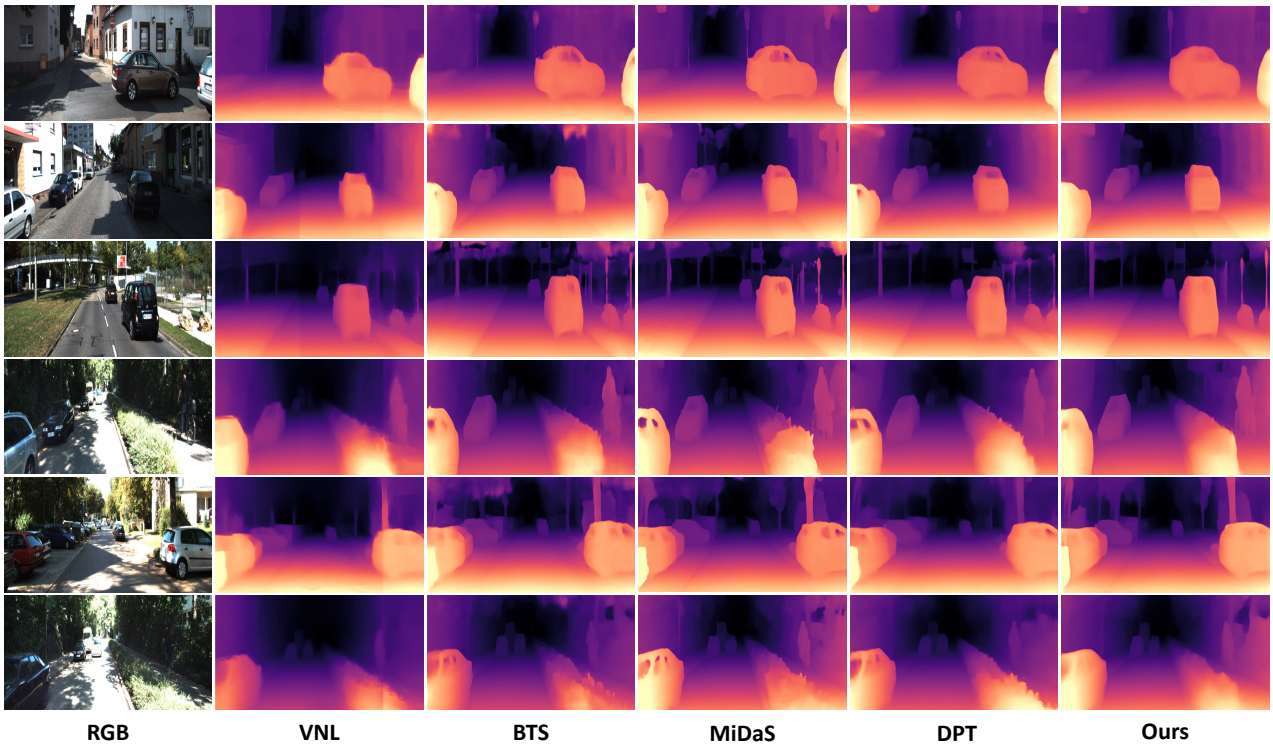

Figure 14: Visual comparisons on KITTI [13] dataset. We compare DADP with state-of-the-art supervised [29, 35, 36, 56] approaches.