

D^3 -RSMDE: $40\times$ Faster and High-Fidelity Remote Sensing Monocular Depth Estimation

Ruizhi Wang^{1*}, Weihan Li^{1*}, Zunlei Feng^{1,2}, Haofei Zhang^{2,3 †}, Mingli Song^{1,2,3}, Jiayu Wang⁴, Jie Song¹, Li Sun⁵

¹School of Software Technology, Zhejiang University,

²State Key Laboratory of Blockchain and Data Security, Zhejiang University,

³ Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security,

⁴ College of Computer Science and Technology, Zhejiang University,

⁵ Ningbo Global Innovation Center, Zhejiang University

Abstract

Real-time, high-fidelity monocular depth estimation from remote sensing imagery is crucial for numerous applications, yet existing methods face a stark trade-off between accuracy and efficiency. Although using Vision Transformer (ViT) backbones for dense prediction is fast, they often exhibit poor perceptual quality. Conversely, diffusion models offer high fidelity but at a prohibitive computational cost. To overcome these limitations, we propose **Depth Detail Diffusion for Remote Sensing Monocular Depth Estimation (D^3 -RSMDE)**, an efficient framework designed to achieve an optimal balance between speed and quality. Our framework first leverages a ViT-based module to rapidly generate a high-quality preliminary depth map construction, which serves as a structural prior, effectively replacing the time-consuming initial structure generation stage of diffusion models. Based on this prior, we propose a **Progressive Linear Blending Refinement (PLBR)** strategy, which uses a lightweight U-Net to refine the details in only a few iterations. The entire refinement step operates efficiently in a compact latent space supported by a Variational Autoencoder (VAE). Extensive experiments demonstrate that D^3 -RSMDE achieves a notable 11.85% reduction in the Learned Perceptual Image Patch Similarity (LPIPS) perceptual metric over leading models like Marigold, while also achieving over a $40\times$ **speedup** in inference and maintaining VRAM usage comparable to lightweight ViT models.

Introduction

Real-time, high-fidelity monocular depth estimation from remote sensing images is a fundamental and critical task in computer vision, with profound implications across numerous domains such as autonomous UAV navigation and 3D terrain modeling. One prominent technical approach to this task involves dense prediction architectures employing ViT backbones (Dosovitskiy et al. 2021), such as DPT (Ranftl, Bochkovskiy, and Koltun 2021) and AdaBins (Bhat, Alhashim, and Wonka 2021). Although offering rapid in-

*These authors contributed equally.

†Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

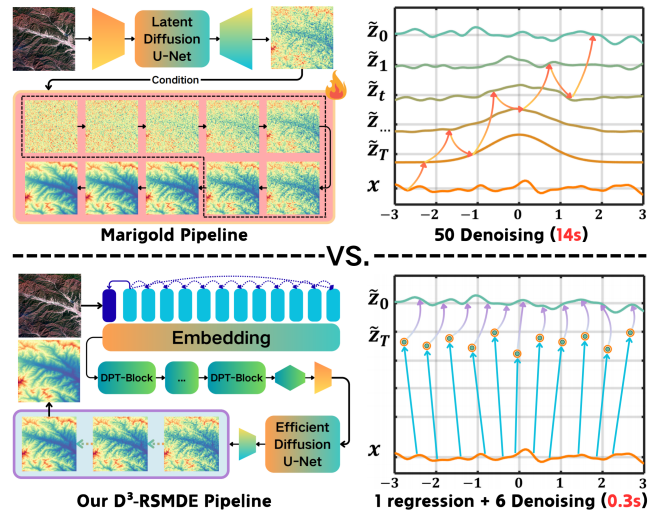


Figure 1: The difference between D^3 -RSMDE and Marigold. Compared to the multiple denoising reconstructions of Marigold, our D^3 -RSMDE firstly adopts efficient ViT to regression coarse depth map and then obtain fine-grained high-fidelity depth map with fewer denoising steps.

ference, these models have inherent limitations in capturing high-frequency details. Recent studies suggest that ViTs act as low-pass filters, showing a strong tendency to learn global, low-frequency signals while neglecting fine textures (Park and Kim 2022; Liu et al. 2022). This low-pass filtering characteristic of ViTs often leads to perceptually inferior depth maps with blurry details, a deficiency quantifiable by high LPIPS scores (Zhang et al. 2018).

Conversely, an alternative paradigm is diffusion-based generative frameworks, such as Marigold (Ke et al. 2024) and EcoDepth (Patni, Agarwal, and Arora 2024). These methods demonstrate remarkable fidelity, generating depth maps with fine-grained textures. This capability is especially useful for remote sensing applications, which often involve intricate surface details. However, their iterative refinement process is computationally intensive and ill-suited for real-

time requirements. While conventional acceleration strategies exist, such as optimizing samplers (Zhao et al. 2023; Lu et al. 2025; Xue et al. 2024) or employing model distillation (Liu et al. 2023; Luo et al. 2023; Wang et al. 2024), they typically require the pre-training of a large and resource-intensive base model or sacrifice generative quality for speed (Lu et al. 2022; Song et al. 2023), and the limited availability of large-scale training data in the remote sensing domain further constrains the applicability of such distillation-based methods. Furthermore, the iterative nature of diffusion models inherently dedicates the initial, computationally expensive steps to establishing low-frequency macrostructures before refining high-frequency details (Hertz et al. 2022; Karras et al. 2024). Consequently, these traditional acceleration methods, which speed up the entire process uniformly, fail to radically alter this inefficient workflow.

To investigate this performance bottleneck, we analyzed the depth estimation pipeline of an advanced diffusion model, Marigold, on remote sensing imagery (as illustrated in Fig. 1). We observed a critical phenomenon: during the entire inference process, which took nearly 14 seconds on an NVIDIA 3090 GPU, a majority of the timesteps (the early stage) were dedicated to establishing the macro-structure and coarse outline of the depth map. In contrast, only a few final steps were used for detail refinement. This insight suggests that the time-consuming initial structure-building phase could be effectively replaced by a more efficient, non-diffusion model, thereby dramatically improving efficiency while preserving high-fidelity details.

Motivated by this observation, we introduce D^3 -RSMDE, a novel framework designed to achieve a dual optimization of speed and accuracy. First, we leverage a fast ViT-based module, optimized with the Hierarchical Depth Normal (HDN) (Zhang et al. 2022) loss function, to efficiently predict a high-quality coarse depth map, thereby completely supplanting the time-consuming initial stage of the diffusion process. Second, we design a lightweight diffusion refinement module that performs coarse-to-fine detail enhancement over a significantly shorter trajectory. The core of this module is our innovative PLBR strategy, which ensures both accuracy and efficiency. In PLBR, at each refinement step, the model is conditioned on both the original coarse map and the output from the previous step, with their influences dynamically attenuated. This provides a stable global structure reference while preventing excessive interference with detail synthesis, enabling controllable and precise reconstruction. Finally, to further accelerate this process for large-scale remote sensing images, we incorporate a Variational Autoencoder (VAE) (Kingma, Welling et al. 2013), mapping the entire refinement operation into a compact latent space to drastically reduce computational overhead.

In summary, our main contributions are as follows:

- We propose D^3 -RSMDE, specifically designed for efficient and high-fidelity monocular depth estimation from remote sensing imagery, obtaining over $40\times$ speedups compared to Marigold.
- We introduce an innovative PLBR and leverage a VAE to operate in the latent space, significantly enhancing accu-

racy and computational efficiency.

- Through extensive experiments on five datasets, we demonstrate that our method achieves SOTA or second-best performance, while its efficiency is comparable to that of lightweight ViT-based models, effectively resolving the bottlenecks of existing technologies.

Related Works

ViT-based Monocular Depth Estimation

ViTs (Dosovitskiy et al. 2021) have been widely adopted for MDE, owing to their powerful global feature extraction capabilities. This line of research has produced a series of models prioritizing efficiency and global consistency.

One of the early explorations in this domain was AdaBins (2021). This work uses Transformer (Vaswani et al. 2017) to reformulate depth regression as a classification problem and significantly improved accuracy. Subsequently, DPT (2021) pioneered the use of the ViT architecture as an encoder backbone, demonstrating its superiority over traditional CNNs in capturing global image context and laying a solid foundation for subsequent research.

To enhance model generalization across diverse scenes, Eftekhari et al. introduced Omnidata (2021), an innovative parametric pipeline for sampling and rendering multi-task vision datasets from real-world 3D scans. Subsequent studies further proposed the HDN loss function (2022), which enhances geometric consistency by enforcing constraints on surface normals at multiple scales. These methods collectively form the foundation for the first component of our work: the rapid generation of a structural depth map prior.

More recently, Depth Anything (Yang et al. 2024), demonstrates exceptional zero-shot capabilities on general-purpose scenes. However, its performance fails to generalize to the domain of remote sensing imagery. The powerful depth priors learned from near-view perspectives are ill-suited for the unique top-down viewpoints, distinct geometric properties, and the absence of conventional depth cues found in remote sensing data. This domain gap underscores the necessity of developing specialized algorithms tailored for the unique challenges of monocular depth estimation in the remote sensing context.

Diffusion-based Monocular Depth Estimation

Diffusion models (Anderson and Akram 2024) have opened new frontiers in high-quality depth map synthesis, with their core strength lying in unparalleled detail recovery and realism. ECoDepth (2024) fuses pretrained ViT embeddings with a diffusion model for monocular depth estimation: global semantic features are first extracted via ViT, then conditioned alongside the input image in the diffusion denoising process to produce depth maps, achieving outstanding results on near-view scenes.

At the same time, Marigold (2024) ingeniously repurposes a pre-trained text-to-image diffusion model (like Stable Diffusion) for zero-shot monocular depth estimation. By incorporating depth information into Gaussian noise, it can generate depth maps with remarkable realism and fine detail. These works demonstrate the potent potential of diffu-

sion models for detail generation, they also highlight their inherent drawback of high computational cost, a core problem our work aims to solve.

Leveraging diffusion models for refinement is a nascent trend in computer vision. However, existing approaches are largely ill-suited for efficient, continuous depth map refinement: DDRM (Kawar et al. 2022) in image restoration aims to reverse physical degradation, not correct neural network prediction ambiguity; SegRefiner (Wang et al. 2023) in segmentation tackles discrete labels instead of continuous values; DiffFlow3D (Liu et al. 2024) predicts scene flow through DDIM-based iterative diffusion and achieves strong performance, yet it is designed for irregular point clouds and cannot be directly applied to monocular depth estimation in remote sensing; and DCTPose (Chen et al. 2024) in pose estimation refines sparse coordinates rather than depth maps. These fundamental differences in task definition and data structure highlight a critical gap in monocular depth estimation for a framework designed specifically for efficient, high-fidelity refinement of fine-grained remote-sensing depth maps, our proposed D^3 -RSMDE is precisely intended to fill this gap.

VAE for Diffusion Models

To mitigate the immense computational expense of operating directly in the high-dimensional pixel space, researchers have shifted to performing the diffusion process within a compact latent space. The theoretical foundation for this paradigm was laid by Kingma & Welling (Kingma, Welling et al. 2013) with the introduction of the VAE, which pioneered the use of variational Bayesian methods to learn a mapping from data to a probabilistic latent space.

Years later, Rombach et al. (Rombach et al. 2022) applied this concept to generative models, proposing Latent Diffusion Models and designing the crucial AutoencoderKL (AEKL). This VAE, widely adopted by models like Stable Diffusion, efficiently compresses images into a low-dimensional latent space, drastically reducing the computational cost of the diffusion process and establishing it as a mainstream approach.

However, the standard AEKL faces an optimization dilemma between “reconstruction” and “generation”: it must both ensure high-fidelity image reconstruction and provide a smooth, regularized latent space for the diffusion U-Net. To address this, various improvements have been proposed. For example, VA_VAE (Yao, Yang, and Wang 2025) demonstrates that by introducing a lightweight auxiliary decoder to exclusively handle the reconstruction loss, the primary decoder can be freed to focus on generative quality. This decoupled design, while retaining the AEKL backbone, significantly accelerates training and enhances final generation quality. The development of these VAE technologies is the key underpinning that enables our design of a computationally feasible, lightweight diffusion refinement module.

Method

Overview

The D^3 -RSMDE framework is a hybrid architecture that efficiently combines different paradigms for accurate monocular depth estimation from remote sensing images. It first employs a ViT-based module to quickly generate a structurally consistent coarse depth map, avoiding the slow contour construction of traditional diffusion methods. A lightweight diffusion module then refines this depth scene in a few steps in a compact latent space, producing a detailed depth output.

Preliminary Scene Structuring

The preliminary depth scene estimation module is designed to produce a globally consistent and structurally coherent initial depth map scene for subsequent refinement. It refers to the DPT model and employs a hybrid architecture that combines a ViT encoder with a convolution-based decoder. The encoder divides the input image into non-overlapping patches of size $p \times p$, yielding $N_p = \frac{HW}{p^2}$ flattened tokens. Each token is linearly projected into a D -dimensional embedding space, augmented with learnable positional encodings and a global readout token. These tokens $\{t_0, t_1, \dots, t_{N_p}\}$ are processed by L layers of multi-head self-attention, excelling at long-range dependency modeling to ensure global structural consistency.

The N_p tokens are then reassembled into a feature map of shape $\frac{H}{p} \times \frac{W}{p} \times D$ based on their original patch positions. A Resample layer adjusts the resolution using a 1×1 convolution to map channel dimension to \hat{D} , followed by either a strided 3×3 convolution for downsampling (if $s \geq p$) or a transposed 3×3 convolution with stride p/s for upsampling (if $s < p$). This reassembly and resampling process is performed at transformer layers $\{3, 6, 9, 12\}$ to extract multi-scale representations. These are subsequently fused in a top-down manner using a RefineNet-style decoder, where each stage doubles the spatial resolution and merges features hierarchically. The final feature map, at half the original input resolution, is passed to a task-specific output head to generate the coarse depth map.

To supervise the model, we employ the HDN loss, which balances global structural consistency with local detail preservation. The HDN loss is defined as:

$$L_{\text{HDN}} = \frac{1}{M} \sum_{i=1}^M \left(\frac{1}{|\mathcal{U}_i|} \sum_{u \in \mathcal{U}_i} |N_{u_i}(d_i) - N_{u_i}(\tilde{d}_i)| \right),$$

where the normalized representations is given by:

$$N_{u_i}(d_i) = \frac{d_i - \text{median}_{u_i}(\mathbf{d})}{\frac{1}{|\mathcal{U}_i|} \sum_{j=1}^{|\mathcal{U}_i|} |d_i - \text{median}_{u_i}(\mathbf{d})|},$$

median_{u_i} computes the median depth of locations, M is the total number of effective pixels, i represents each pixel, \mathcal{U}_i denotes the set of multi-scale contexts to which pixel i belongs, d is the ground truth and \tilde{d} is the predicted depth map. These contexts are constructed using three strategies: spatial grid partitioning, depth-range segmentation, and depth-quantile grouping. For each context, a shared SSI module

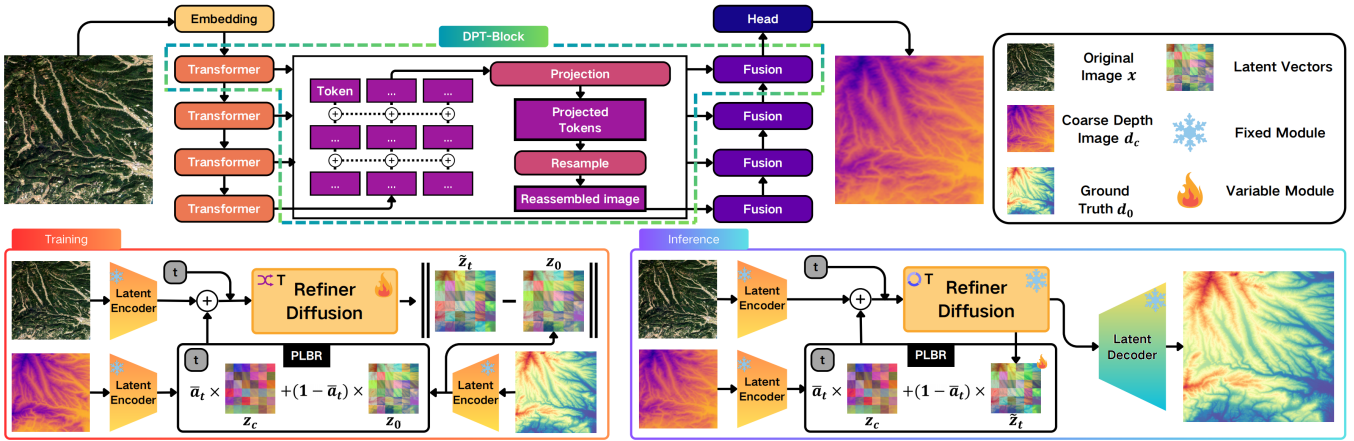


Figure 2: **The framework of our D^3 -RSMDE.** During the training process, ViT first performs regression on the input original remote sensing images x to obtain the coarse depth map construction d_c , and then together with Ground Truth d_0 and the x , obtains the samples for training Refiner Diffusion through PLBR. In the inference process, the d_0 is replaced by the output of each step of Refiner Diffusion to obtain a refined and high-fidelity remote sensing depth estimation map.

(Zhang et al. 2022) computes the MAE, and the errors are aggregated across contexts and scales.

Progressive Detail Refinement

Although the coarse depth map generated by the initial prediction module exhibits a low MAE, it scores poorly on LPIPS, appearing visually blurry and lacking in high-frequency details. To address this, we designed a Dynamic Guided Diffusion Refiner. This design departs from the conventional paradigm of Markovian based diffusion models (Benton et al. 2024) that reconstruct data from pure noise. Inspired by SegRefiner (Wang et al. 2023), we instead formulate a non-Markovian based coarse-to-fine refinement process. This mechanism ensures that the globally consistent structural information from the coarse depth map continuously guides the entire refinement process. This module greatly shortens the process of diffusion to reconstruct depth map details from pure noise, which can efficiently recover clear and realistic depth details with only a few iterations.

Efficient Diffusion backbone. The core is a conditional diffusion denoising model f , tasked with predicting the latent representation of the refined depth map at a given timestep t . To strike a balance between computational efficiency and expressive power, our model operates entirely within a compact latent space defined by a pre-trained VAE.

Compared to a conventional Stable Diffusion U-Net, our model excises modules that are superfluous for our depth map refinement task, such as text cross-attention and multi-source conditioning frameworks. This results in a significantly more lightweight architecture specialized for refining depth details from image and timestep information. The detailed Unet architecture can be found in Appendix A.4.

Progressive Linear Blending Refinement. Unlike traditional diffusion models that gradually transform data into pure Gaussian noise, we introduces a refiner-specific strategy called PLBR. This process is different from the tradi-

tional diffusion strategy based on Markov. PLBR is based on a non-Markovian process, linearly interpolates between the high-quality ground truth depth map and a coarse depth map during training. During inference, this process is reversed through a Progressive Refinement procedure.

The goal of the forward process is to generate training samples of varying “levels of noise” for the model f . We define two key inputs: the ground truth depth map d_0 and a coarse depth map d_c generated by a DPT module. These inputs are first encoded by the VAE into their respective latent representations, z_0 and z_c .

We design a diffusion schedule coefficient:

$$\bar{\alpha}_t = \frac{\epsilon}{T-1}(T-t-1),$$

where $t \in [0, \dots, T-1]$, ϵ is a positive constant that close to 1 but not equal to 1, ensuring that the coarse depth map’s contribution is always present (In our experiment, $\epsilon = 0.8$). The blended latent representation z_t at any timestep t is generated via the following linear interpolation formula:

$$z_t = \bar{\alpha}_t z_0 + (1 - \bar{\alpha}_t) z_c,$$

This process simulates a continuous transition from fine to coarse. When t is small, $\bar{\alpha}_t$ is close to 1, and the primary component of z_t is the ground truth z_0 . As t increases, $\bar{\alpha}_t$ decreases, and z_t progressively approaches the coarse representation z_c . At the same time, the original remote sensing image x will also be used as the input of additional information concat into the model to provide more basic information. This strategy enables the model to learn how to recover fine depth structures from inputs of varying coarseness across the entire refinement trajectory.

The inference is an iterative refinement process that aims to progressively recover a high-quality depth map d_0 , starting from the coarse map d_c . The process begins at $t = T-1$ and proceeds backward to $t = 0$.

Model	MAE ↓					δ^3 ↑					PSNR ↑					LPIPS ↓				
	J&K	SA	Med	Swi	Ast	J&K	SA	Med	Swi	Ast	J&K	SA	Med	Swi	Ast	J&K	SA	Med	Swi	Ast
Adabins	16.7	28.4	29.0	19.6	44.3	79.9	68.9	<u>86.3</u>	93.1	<u>73.2</u>	22.3	17.4	17.7	21.2	14.7	0.181	0.405	0.367	0.127	0.528
DPT	17.3	34.2	29.7	30.9	43.5	77.3	62.5	81.6	84.6	72.7	22.2	16.7	17.8	17.6	14.9	0.313	0.604	0.520	0.204	0.579
Omnidata	20.1	30.7	28.0	19.2	42.6	61.9	67.8	80.9	90.8	72.1	21.2	18.5	18.2	21.6	15.0	0.354	0.482	0.479	0.135	0.553
Pix2pix	24.5	39.3	39.4	38.9	44.3	68.9	55.1	72.0	76.8	69.3	18.6	15.2	15.1	15.5	14.3	0.450	0.485	0.434	0.775	0.937
Marigold	14.2	23.7	24.7	21.3	<u>40.0</u>	<u>83.1</u>	<u>71.7</u>	85.8	89.6	72.8	<u>24.3</u>	19.6	19.3	21.4	<u>15.7</u>	0.162	<u>0.326</u>	0.329	0.144	0.488
EcoDepth	26.4	49.0	49.0	37.4	43.3	65.7	49.4	67.4	77.9	69.9	17.9	13.3	13.2	16.0	14.5	0.461	0.428	0.563	0.265	0.702
D^3 -RSMDE (VA_VAE)	<u>13.6</u>	<u>21.7</u>	<u>23.4</u>	<u>14.1</u>	41.7	79.1	70.0	85.9	<u>93.3</u>	67.8	23.7	<u>20.1</u>	<u>20.0</u>	<u>24.2</u>	15.1	0.203	0.366	<u>0.301</u>	<u>0.107</u>	0.574
D^3 -RSMDE (AEKL)	12.7	20.5	22.1	13.4	36.1	83.3	73.9	88.1	94.3	76.1	24.5	20.6	20.4	24.8	16.2	<u>0.180</u>	0.318	0.290	0.104	<u>0.511</u>

Table 1: Quantitative analysis of SOTA methods. The best result is highlighted in **bold** and the second best result is underlined.

- First, we encode the initial coarse depth map d_c into its latent representation z_c and set it as the input for the first timestep, i.e., $z_{T-1} = z_c$.
- At each timestep t , the f model receives the current input z_t , the remote sensing image latent z_x , and the timestep embedding e_t , and predicts the final refined latent representation $\tilde{z}_{0|t}$:

$$\tilde{z}_{0|t} = f([z_x, z_t], e_t),$$

- Then, we employ a novel progressive refinement strategy to generate the input for the next step z_{t-1} . Unlike conventional DDPMs, which rely on the previous state z_t and the current prediction $\tilde{z}_{0|t}$ to estimate z_{t-1} , our method directly blends the new prediction with the original coarse representation z_c . This ensures that each refinement step is anchored to the initial coarse structure, preventing error accumulation, follows the update rule:

$$z_{t-1} = \bar{\alpha}_{t-1} \tilde{z}_{0|t} + (1 - \bar{\alpha}_{t-1}) z_c,$$

This iterative process continues until $t = 0$, yielding the final latent representation $\tilde{z}_{0|1}$. Finally, the VAE decoder transforms $\tilde{z}_{0|1}$ back into the pixel space to produce the final, refined depth map \tilde{d}_0 .

Experiment

Experimental Settings

Benchmark and Metrics. In order to comprehensively evaluate the performance of the model in different terrain, resolution, and dataset sizes, we chose 5 datasets from RS3DBench(Wang et al. 2025): Japan + Korea (2,650 pairs, coastal mountainous terrain, 30 m resolution, J&K), Southeast Asia (7,000 pairs, plains and hills, 30 m resolution, SA), Mediterranean (29,225 pairs, desert and plateau, 30 m resolution, Med), Australia (1,249 pairs, plain, 5m resolution, Ast), Switzerland (4,827 pairs, mountain, 2m resolution, Swi). In order to better reflect human perception of depth map quality, in addition to traditional MAE, δ^3 and PSNR, etc., we also introduce LPIPS metric. LPIPS is a trained perceptual loss metric that uses the pre-trained

AlexNet network to calculate the structural and texture similarity between images to evaluate the perceptual effect of image reconstruction or generation quality. A detailed description of the metrics can be found in Appendix A.1.

Implementation Details. We release two model versions using AEKL (2022) and VA_VAE (2025), respectively. For reproducibility, the random seed for all experiments was fixed to 42. The initial prediction module was trained using the HDN loss with an initial learning rate (LR) of 5×10^{-5} and a weight decay of 10^{-4} . We employed a LR scheduler that reduces the LR by a factor of 0.6 upon a validation loss plateau of 5 epochs. To generate unbiased training and testing sets for the diffusion refiner, we performed 5-fold cross-validation on the outputs of the initial prediction module. The refiner was subsequently trained on these outputs using an L1 loss. The initial LR was set to 1×10^{-4} with $T = 6$, $\epsilon = 0.8$ and a similar LR scheduler (patience=5, factor=0.5). Additional hyperparameter details are provided in the Appendix A.2.

Baselines for Comparison To comprehensively benchmark our D^3 -RSMDE model against existing technologies, we selected a series of baseline models that have achieved SOTA performance and are based on different mainstream architectural paradigms.

- **ViT Models:** We include DPT, Omnidata and AdaBins. These three models are representative works that leverage ViT for efficient monocular depth estimation.
- **Diffusion Models:** We compare against Marigold and EcoDepth. They represent the cutting edge in high-fidelity depth synthesis using diffusion processes.
- **GAN-based Model:** We also include Pix2pix (Panagiotou et al. 2020), a GAN-based model specifically optimized for remote sensing monocular depth estimation.

Quantitative Analysis

The quantitative evaluation results for all models are summarized in Table 1, with visual comparisons against representative models shown in Fig. 4. The results comprehensively demonstrate that our proposed D^3 -RSMDE framework achieves SOTA or second-best performance across

Configuration	MAE ↓			RMSE ↓			δ^3 ↑			PSNR ↑			LPIPS ↓		
	J&K	SA	Med	J&K	SA	Med	J&K	SA	Med	J&K	SA	Med	J&K	SA	Med
DPT (Baseline)	17.3	34.2	29.7	22.1	42.3	36.4	77.3	62.5	81.6	22.2	16.7	17.8	0.313	0.604	0.520
ViT Module Only	12.8	23.3	23.2	17.2	30.4	28.5	79.2	70.3	84.8	<u>24.4</u>	19.8	20.2	0.305	0.503	0.439
Full Model (w/o VAE, T=3)	13.0	22.1	22.9	17.2	28.6	28.0	83.1	72.2	86.2	24.3	20.1	<u>20.3</u>	0.222	0.361	0.335
Full Model (w/o VAE, T=6)	12.4	<u>21.6</u>	<u>22.6</u>	16.7	<u>27.9</u>	<u>27.5</u>	84.9	<u>72.4</u>	<u>86.6</u>	24.5	<u>20.3</u>	20.4	0.218	<u>0.343</u>	0.344
Full Model (w/o VAE, T=10)	13.5	23.0	24.5	17.6	29.2	29.5	82.5	69.3	85.0	24.1	19.8	19.9	0.239	0.365	0.365
Full Model (AEKL, T=3)	13.3	21.8	24.3	17.4	28.3	29.8	82.3	71.5	83.0	24.2	20.1	19.7	0.199	0.350	0.323
Full Model (AEKL, T=6)	<u>12.7</u>	20.5	22.1	<u>16.9</u>	26.6	27.1	<u>83.3</u>	73.9	88.1	24.5	20.6	20.4	0.180	0.318	0.290
Full Model (AEKL, T=10)	13.3	21.9	24.1	17.3	28.3	29.7	80.8	72.0	84.5	24.1	20.0	19.6	<u>0.187</u>	0.349	0.334
Full Model (VA_VAE, T=3)	14.4	21.9	24.3	19.3	28.3	29.8	78.2	69.0	83.0	23.2	20.1	19.7	0.260	0.388	0.323
Full Model (VA_VAE, T=6)	13.6	21.7	23.4	18.2	28.1	28.8	79.1	70.0	85.9	23.7	20.1	20.0	0.203	0.366	<u>0.301</u>
Full Model (VA_VAE, T=10)	14.9	21.9	24.1	20.1	28.6	29.7	77.8	70.4	84.5	22.7	20.0	19.6	0.324	0.374	0.334

Table 2: Ablation study of D^3 -RSMDE. The best result is highlighted in **bold** and the second best result is underlined.

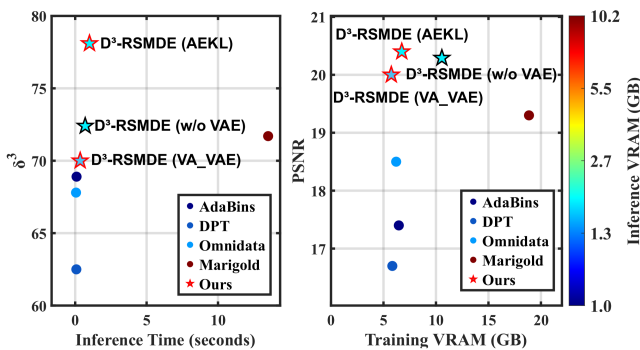


Figure 3: Comparison of model efficiency.

the majority of metrics on all test datasets. Our framework significantly outperforms the GAN-based Pix2pix and the ViT-based Adabins, DPT, and Omnidata. More remarkably, even when compared against retrained Marigold, a model renowned for its high-fidelity synthesis, our D^3 -RSMDE achieves a substantial relative improvement of up to 13.50% in MAE and 11.85% in LPIPS, showcasing its exceptional overall performance. The general observation is that all diffusion-based models, except the retrained EcoDepth, outperform the LPIPS perception metric. This highlights the distinct advantage of the diffusion architecture for generating photorealistic RSMDE. A detailed diagnostic analysis for the underperformance of EcoDepth, which is also a diffusion model, is provided in Appendix A.3.

Efficiency Analysis

To comprehensively evaluate the computational efficiency of our D^3 -RSMDE framework, we conducted a systematic analysis of the average inference time per image, training time per epoch while training Japan + Korea Dataset, maximum inference VRAM usage, and maximum training VRAM usage (with a batch size of 2). All experiments were performed on a consistent hardware/software platform (Ubuntu 16.04, Intel CPU E5-2699 v4, NVIDIA 3090 GPU, 125G memory, python3.10.6), with the results presented in

Fig. 3, where D^3 -RSMDE shows its diffusion module.

Experimental data reveals that our framework demonstrates a decisive advantage across all efficiency metrics when compared to SOTA diffusion methods like Marigold and EcoDepth. Most notably, D^3 -RSMDE achieves an inference speed over **40 times** faster than Marigold while also significantly reducing training time. In terms of resource consumption, both the training and inference memory footprints of our model are substantially lower than other diffusion models. Furthermore, it is noteworthy that the VRAM of our model is on par with that of lightweight ViT-based model such as DPT and Omnidata at inference and training time. This result provides strong evidence that D^3 -RSMDE successfully reduces computational overhead to a level comparable to non-generative models, all while retaining the high-quality synthesis capabilities of diffusion models.

Ablation Study

In this section, we conduct a series of detailed ablation studies to individually validate the effectiveness and efficiency of the key components within our D^3 -RSMDE framework. The comprehensive results are summarized in Table 2.

Effectiveness of the ViT Module. First, we compared the performance of our ViT Module against a standard DPT. As shown in Table 2, although our module is based on the DPT architecture, its optimization with the HDN loss function leads to a significant improvement in the quality of the initial depth map. This provides a superior structural prior for the subsequent diffusion refinement, thereby effectively enhancing the final output accuracy.

Efficiency of using VAE. Next, we evaluated the role of performing refinement in the latent space using a VAE. The results indicate that, with the same number of denoising steps, the model variant with a VAE achieves comparable accuracy to a variant that performs refinement directly in the pixel space. However, considering the efficiency data from Fig. 3, we found that incorporating a VAE improves training speed by 54.91% and reduces training VRAM by 36.17%. This provides strong evidence that latent space diffusion can

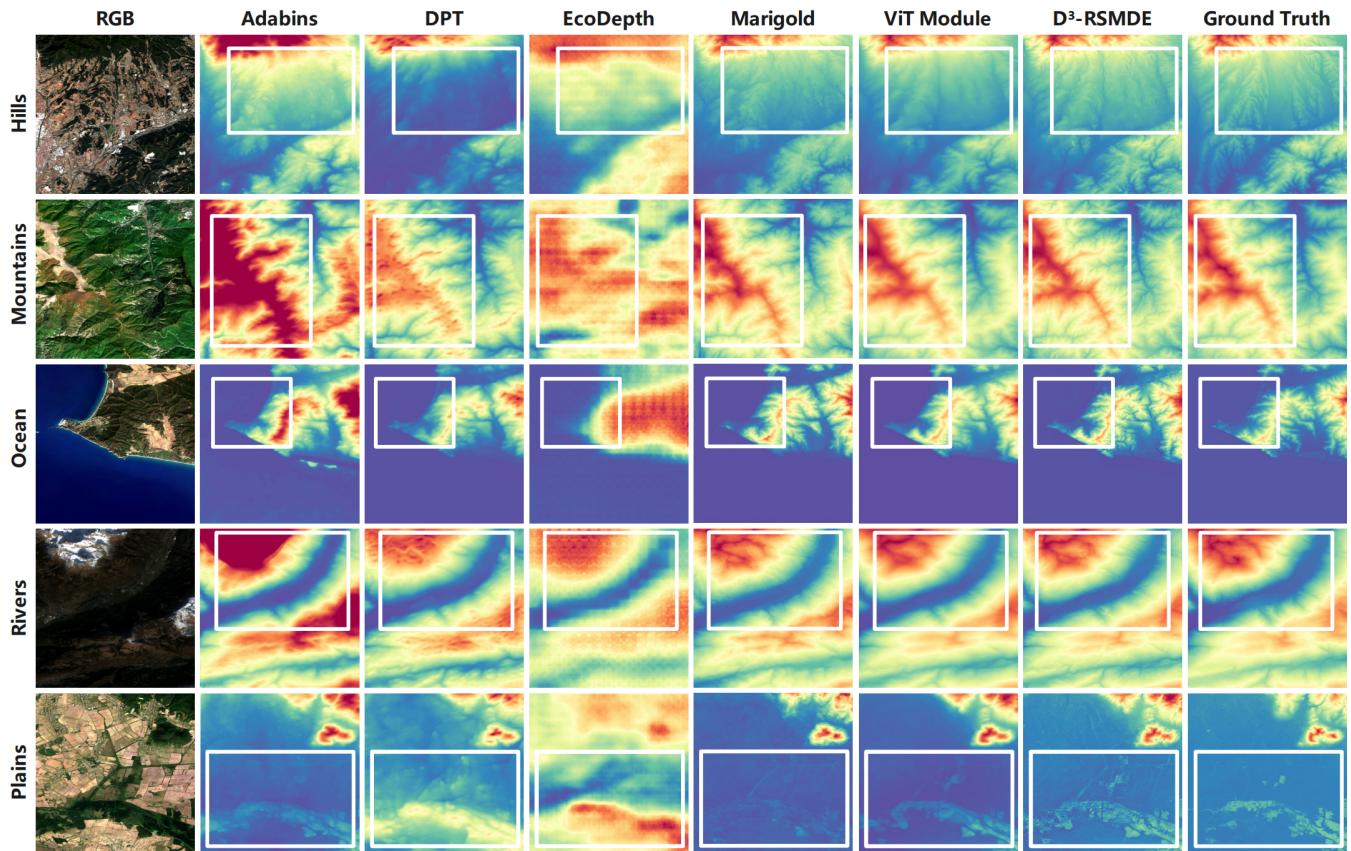


Figure 4: Comparison of our D^3 -RSMDE and some SOTA methods in different categories of remote sensing images.

dramatically enhance training efficiency and lower the resource threshold without compromising model performance.

Impact of Denoising Steps. We investigated the impact of varying the number of denoising steps (T) on the final result. As shown in Table 2, the vast majority of those achieving the SOTA metric are concentrated in the models with step=6 without VAE and with AEKL. The performance of the model improves markedly as the number of steps increases from 3 to 6, suggesting that $T=3$ allows for an insufficient refinement process that does not fully leverage the model’s detail recovery capabilities. However, when the steps are further increased to 10, performance slightly degrades. We hypothesize that this is due to an “over-refinement” phenomenon. After several iterations, the intermediate result is already quite detailed. Excessive additional steps may cause the model to amplify minor noise or artifacts from the initial prediction, or even to hallucinate spurious textures that are plausible according to its generative prior but not faithful to the source image. Therefore, $T=6$ represents the optimal trade-off between performance and efficiency.

Effectiveness of Diffusion Refinement. Finally, to validate the effectiveness of the diffusion module in our framework, we compared the full model against the results from the initial prediction module alone (ViT Module Only). The data reveals that after diffusion refinement, the LPIPS score

significantly decreased by 40.98%, 36.78%, and 33.94% on the J&K, SA, and Med datasets, respectively. Meanwhile, it can also be seen from the Fig. 4 that D^3 -RSMDE outputs more texture information than ViT Module. These evidences strongly prove that our lightweight diffusion module dramatically enhances the perceptual quality and visual clarity of the generated depth maps.

Conclusion

In conclusion, even though inferring depth information from a single remote sensing image remains difficult for both models and human eyes, our comprehensive experiments strongly demonstrate the effectiveness and exceptional efficiency of the proposed D^3 -RSMDE framework. Our model achieves accuracy comparable to, and in some cases surpassing, the SOTA Marigold, a model renowned for its high-fidelity synthesis. More significantly, D^3 -RSMDE represents a breakthrough in computational efficiency. It uses PLBR strategy that provides a substantially accelerated training process and achieves an inference speedup of up to 40 times, all while maintaining a VRAM footprint on par with lightweight ViT-based architectures. Therefore, our work successfully addresses the critical trade-off between accuracy and efficiency, resolving the prohibitive computational bottleneck that has hindered the practical deployment of high-fidelity diffusion-based models in the field of RSMDE.

Acknowledgments

This work was sponsored by National Natural Science Foundation of China (62576305), and the Fundamental Research Funds for the Central Universities (No. 226-2025-00057).

References

- Anderson, J.; and Akram, N. 2024. Denoising Diffusion Probabilistic Models (DDPM) Dynamics: Unraveling Change Detection in Evolving Environments. *Innovative Computer Sciences Journal*, 10(1): 1–10.
- Benton, J.; Shi, Y.; De Bortoli, V.; Deligiannidis, G.; and Doucet, A. 2024. From denoising diffusions to denoising markov models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(2): 286–301.
- Bhat, S. F.; Alhashim, I.; and Wonka, P. 2021. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4009–4018.
- Chen, Z.; Dai, J.; Pan, J.; and Zhou, F. 2024. Diffusion model with temporal constraint for 3D human pose estimation. *The Visual Computer*, 1–17.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR*.
- Eftekhari, A.; Sax, A.; Malik, J.; and Zamir, A. 2021. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10786–10796.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Karras, T.; Aittala, M.; Lehtinen, J.; Hellsten, J.; Aila, T.; and Laine, S. 2024. Analyzing and Improving the Training Dynamics of Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 24174–24184.
- Kawar, B.; Elad, M.; Ermon, S.; and Song, J. 2022. Denoising diffusion restoration models. *Advances in neural information processing systems*, 35: 23593–23606.
- Ke, B.; Obukhov, A.; Huang, S.; Metzger, N.; Daut, R. C.; and Schindler, K. 2024. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9492–9502.
- Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.
- Liu, J.; Wang, G.; Ye, W.; Jiang, C.; Han, J.; Liu, Z.; Zhang, G.; Du, D.; and Wang, H. 2024. DifFlow3D: Toward Robust Uncertainty-Aware Scene Flow Estimation with Iterative Diffusion-Based Refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15109–15119.
- Liu, X.; Zhang, X.; Ma, J.; Peng, J.; et al. 2023. Instaflow: One step is enough for high-quality diffusion-based text-to-image generation. In *The Twelfth International Conference on Learning Representations*.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11976–11986.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural information processing systems*, 35: 5775–5787.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2025. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, 1–22.
- Luo, S.; Tan, Y.; Huang, L.; Li, J.; and Zhao, H. 2023. Latent Consistency Models: Synthesizing High-Resolution Images with Few-Step Inference. *arXiv:2310.04378*.
- Panagiotou, E.; Chochlakis, G.; Grammatikopoulos, L.; and Charou, E. 2020. Generating elevation surface from a single RGB remotely sensed image using deep learning. *Remote Sensing*, 12(12): 2002.
- Park, N.; and Kim, S. 2022. How Do Vision Transformers Work?
- Patni, S.; Agarwal, A.; and Arora, C. 2024. Ecodepth: Effective conditioning of diffusion models for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28285–28295.
- Ranftl, R.; Bochkovskiy, A.; and Koltun, V. 2021. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, 12179–12188.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency models.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wang, J.; Wang, R.; Song, J.; Zhang, H.; Song, M.; Feng, Z.; and Sun, L. 2025. RS3DBench: A Comprehensive Benchmark for 3D Spatial Perception in Remote Sensing. *arXiv preprint arXiv:2509.18897*.
- Wang, M.; Ding, H.; Liew, J. H.; Liu, J.; Zhao, Y.; and Wei, Y. 2023. SegRefiner: Towards Model-Agnostic Segmentation Refinement with Discrete Diffusion Process. In *NeurIPS*.
- Wang, Y.; Cheng, L.; Duan, M.; Wang, Y.; Feng, Z.; and Kong, S. 2024. Improving knowledge distillation via regularizing feature direction and norm. In *European Conference on Computer Vision*, 20–37. Springer.

- Xue, S.; Liu, Z.; Chen, F.; Zhang, S.; Hu, T.; Xie, E.; and Li, Z. 2024. Accelerating diffusion sampling with optimized time steps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8292–8301.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10371–10381.
- Yao, J.; Yang, B.; and Wang, X. 2025. Reconstruction vs. generation: Taming optimization dilemma in latent diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 15703–15712.
- Zhang, C.; Yin, W.; Wang, Z.; Yu, G.; Fu, B.; and Shen, C. 2022. Hierarchical Normalization for Robust Monocular Depth Estimation. *NeurIPS*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhao, W.; Bai, L.; Rao, Y.; Zhou, J.; and Lu, J. 2023. Unipc: A unified predictor-corrector framework for fast sampling of diffusion models. *Advances in Neural Information Processing Systems*, 36: 49842–49869.