Uncertainty Diffusion: Parameter-Efficient Depth Refinement via Uncertainty-Guided Diffusion Models

Jeng-Huo Tzeng, Chuan-Yuan Huang, and Kuan-Wen Chen

Department of Computer Science, National Yang Ming Chiao Tung University

Challenge

Monocular depth estimation struggles with occlusions, textureless surfaces, and novel environments, leading to unreliable predictions in ambiguous regions.

Solution

We propose **Uncertainty Diffusion**, a model-agnostic framework that integrates pixel-wise uncertainty into diffusion-based refinement, adaptively focusing on low-confidence regions for targeted depth enhancement.

Key Contributions

- **Uncertainty-guided diffusion** that dynamically focuses refinement on uncertain regions
- **Model-agnostic framework** compatible with any depth estimation baseline
- **Parameter-efficient adaptation**: only fine-tune lightweight refinement network (~14M params)
- **Up to 61% reduction** in log10 error for domain adaptation

Method

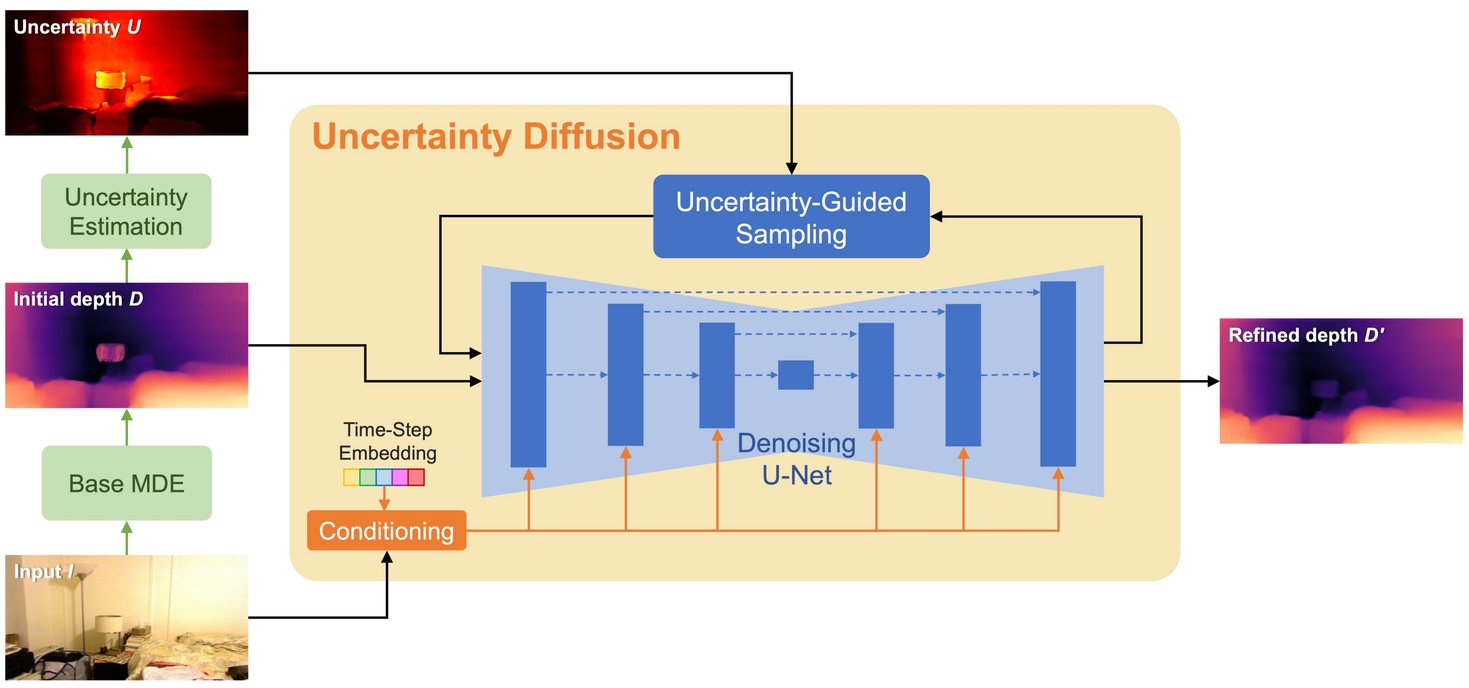

Algorithm 1: Uncertainty Diffusion

Input: Initial depth map D, Uncertainty map U, RGB image I

Output: Refined depth map D'

$\mathbf{x}_{T'} = D$

for $t = T', \ldots, 1$ **do**

$\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}} \varepsilon_\theta(\mathbf{x}_t, t, I) \right) + \sigma_t \mathbf{z} U$$

end

$D' = \mathbf{x}_0$

return D'

Key Insight: Uncertainty U modulates sampling variance—narrow range for confident regions, broad range for uncertain areas.

Experiments

Results on NYU Depth V2

	RMSE(↓)	AbsRel(↓)	log10(↓)
SC-DepthV2	0.532	0.138	0.059
BTS	0.392	0.110	0.047
AdaBins	0.364	0.103	0.044
BinsFormer	0.330	0.094	0.040
DINOv2	0.279	0.090	0.037
Monodepth2(*)	0.4554	0.1365	0.0550
Monodepth2 with U.D.	**0.4469**	**0.1285**	**0.0528**
ZoeDepth	0.2700	0.0750	0.0320
ZoeDepth with U.D.	**0.2665**	**0.0730**	**0.0310**
Depth-Anything	0.206	0.0560	0.0240
Depth-Anything with U.D.	**0.2041**	**0.0556**	**0.0238**

Consistent improvements across all baseline models

Parameter-Efficient Domain Adaptation

	DIODE			SUN RGB-D		
	RMSE(↓)	AbsRel(↓)	log10(↓)	RMSE(↓)	AbsRel(↓)	log10(↓)
Monodepth2	3.817	0.911	1.142	2.009	0.876	0.924
Monodepth2 with U.D.	**2.975 (-22.06%)**	**0.583 (-36.00%)**	**0.444 (-61.12%)**	**1.533 (-23.69%)**	**0.626 (-28.53%)**	**0.431 (-53.35%)**
ZoeDepth	2.092	0.394	0.241	0.750	0.291	0.120
ZoeDepth with U.D.	**1.923 (-8.07%)**	**0.371 (-5.83%)**	**0.202 (-16.18%)**	**0.685 (-8.67%)**	**0.286 (-1.71%)**	**0.115 (-4.17%)**
Depth-Anything	1.943	0.391	0.230	0.491	0.223	0.068
Depth-Anything with U.D.	**1.729 (-11.01%)**	**0.329 (-15.85%)**	**0.176 (-23.47%)**	**0.455 (-7.33%)**	**0.219 (-1.79%)**	**0.066 (-2.94%)**

No base model retraining required

Qualitative Results

❋Green arrows highlight improvements

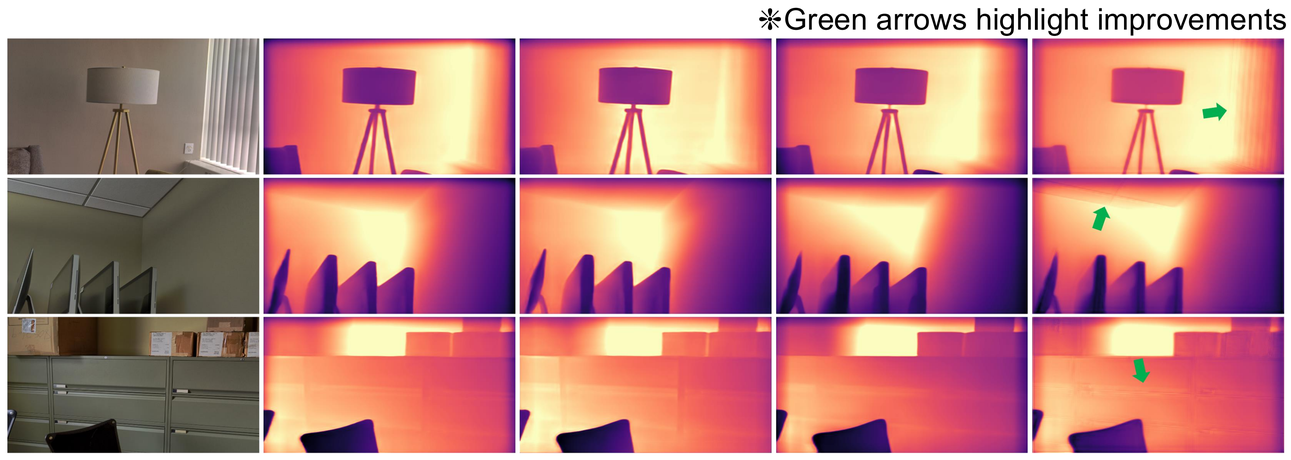

Conclusion

Uncertainty Diffusion is a versatile, model-agnostic framework that:

- Leverages pixel-wise uncertainty to guide diffusion-based depth refinement
- Achieves consistent improvements across diverse baseline models
- Enables highly efficient domain adaptation (only ~14M params fine-tuned)