

ETA: Energy-based Test-time Adaptation for Depth Completion

Younjoon Chung* Hyoungseob Park* Patrick Rim* Xiaoran Zhang Jihe He
Ziyao Zeng Safa Cicek† Byung-Woo Hong‡ James S. Duncan Alex Wong

Yale University †UCLA ‡Chung-Ang University

{younjoon.chung, hyoungseob.park, patrick.rim, xiaoran.zhang}@yale.edu
safacicek@ucla.edu, hong@cau.ac.kr, {james.duncan, alex.wong}@yale.edu

Abstract

We propose a method for test-time adaptation of pre-trained depth completion models. Depth completion models, trained on some “source” data, often predict erroneous outputs when transferred to “target” data captured in novel environmental conditions due to a covariate shift. The crux of our method lies in quantifying the likelihood of depth predictions belonging to the source data distribution. The challenge is in the lack of access to out-of-distribution (target) data prior to deployment. Hence, rather than making assumptions regarding the target distribution, we utilize adversarial perturbations as a mechanism to explore the data space. This enables us to train an energy model that scores local regions of depth predictions as in- or out-of-distribution. We update the parameters of pretrained depth completion models at test time to minimize energy, effectively aligning test-time predictions to those of the source distribution. We call our method “Energy-based Test-time Adaptation”, or ETA for short. We evaluate our method across three indoor and three outdoor datasets, where ETA improve over the previous state-of-the-art method by an average of 6.94% for outdoors and 10.23% for indoors. Project Page: <https://fuzzythecat.github.io/eta>.

1. Introduction

Reconstructing the 3-dimensional (3D) environment is essential for spatial tasks such as autonomous driving, physical AI agents, and extended reality. Many of these tasks rely on sensor platforms equipped with visual (e.g., camera) and range (e.g., LiDAR or radar) sensors. These sensors are complementary in that: Cameras capture RGB images with (dense) irradiance at each pixel, but 3D reconstruction from visual data is ill-posed due to the loss of a dimension in perspective projection. Range sensors measure coordinates or

point clouds of the 3D environment or scene; likewise these point clouds can be estimated from images via Structure-from-Motion (SfM) or Visual Inertial Odometry (VIO). But in both cases, the point clouds tend to be sparse. Depth completion is the multimodal 3D reconstruction task of estimating a dense depth map from an RGB image and its synchronized sparse point cloud, often projected onto the image plane as a sparse depth map.

As depth completion models, pretrained on “source” datasets, are deployed to support a range of spatial tasks, they often encounter novel environments or 3D scenes and conditions. Hence, there exists a covariate shift when transferring a depth completion model to “target” testing data. Such distributional shifts arise from nuisance variables, e.g., changes in illumination, occlusion, that cause changes in object appearance, which lead to model degradations.

Mitigating this covariate shift can be achieved by means of domain adaptation [13], which updates the model parameters to fit the target testing data, but often requires access to source and target datasets with ground truth. Source-free [31] and unsupervised [70] domain adaptation relax the assumption of having access to source data and ground truth, respectively. However, they still require access to the testing dataset and multiple passes through it for model updates. These requirements pose a hurdle for adoption in real-time spatial applications. Instead, we consider the problem of test-time adaptation (TTA), where one is given access to the testing data in a stream, i.e., one batch at a time, without access to ground truth. Hence, TTA methods must update the parameters of a pretrained model in an unsupervised manner, under low-computational budget, with time constraints, and without the ability to revisit previously-seen examples.

To this end, we aim to model the shift or domain gap between the source and target distributions in order to guide test-time updates to a pretrained model to yield predictions resembling those of the source distribution. This will be facilitated by implicitly modeling the data distributions through an energy function. The energy function

*Equal contribution

will quantify the likelihood of the data belonging to the source distribution by associating a scalar (energy) to the data, where low energy is assigned to more probable (source-like, in-distribution) data points and high energy to out-of-distribution data points. A key insight is that (sparse) point clouds or depth estimates are lower dimensional and tend to exhibit a smaller covariate shift than RGB images [44]. Hence, rather than quantifying the energy of the inputs, we will define an energy function over depth predictions, conditioned on its sparse depth map, allowing an energy model to validate the likelihood of predictions against measured coordinates of the 3D environment. Further, we estimate energy on a patch- or region-level to localize out-of-distribution predictions, which enables targeted updates.

However, it remains a challenge to determine what constitutes out-of-distribution data points without making assumptions on the target test-time data to be observed in the future. Instead, we posit that adversarial perturbations of the input RGB image and sparse depth map can serve as a means of exploring the data space of predicted depth maps. Thus, we train an energy model using source dataset and its out-of-distribution examples, synthesized by adversarial perturbations, in a preparation stage. Once trained, the energy model will be frozen and deployed with the pre-trained depth completion model to the testing environment. For each new example (or batch) observed at test time, we will update the parameters of the depth completion model to minimize energy, which effectively aligns predictions (and the model) back to those of the source distribution. We term our method “Energy-based Test-time Adaptation”, or ETA for short. We extensively evaluate ETA across three indoor and three outdoor datasets for three distinct adaptation scenarios, involving indoor-to-indoor, outdoor-to-outdoor, and outdoor-to-indoor, where we consistently outperform baseline methods and achieve the state of the art.

Our contributions: (1) We propose a novel test-time adaptation framework, ETA, that utilizes an energy model to quantify the likelihood of depth predictions belonging to the source distribution and adapts a pretrained model to the target distribution by minimizing the energy in predictions. (2) We utilize adversarial perturbations as a mechanism to explore the data space, which mitigates the challenge in the lack of out-of-distribution data for training the energy model. (3) We propose a patch-based energy model that localizes high energy regions, allowing further targeted updates to specific regions within the depth predictions.

2. Related Works

Depth completion is the multimodal 3D reconstruction task of inferring a dense depth map from a single image and its associated sparse depth map, unlike single image or monocular depth estimation [9, 16, 19, 29, 54, 64, 71, 86–88]. The sparse depth map is the projection of a point cloud

measured by LiDAR [78, 79], radar [51, 58], or visual inertial odometry systems [14, 73]. Depth completion can be trained in a supervised [5, 10, 12, 23, 24, 45, 63, 82] or unsupervised [2, 4, 36, 38, 39, 52, 57, 72–75, 77, 81] manner. We focus on supervised methods, where [24, 39, 82] introduce strategies to fuse image and depth features at both early and late network stages. In contrast, [22] advocates maintaining separate networks for each modality. To enhance the depth encoder with image cues, [30] proposes a multi-scale cascaded hourglass design, and [5] leverages convolutional spatial propagation layers; this is extended by [45] to non-local propagation for refining depth estimates based on confidence and learnable affinity, and by [35] to dynamically adaptive spatial propagation. Various studies also estimate depth uncertainty: [10, 11, 48, 49] model prediction variance, while [66] applies confidence maps to merge multiple depth predictions. Further, [47, 80, 91] utilize surface normals as an auxiliary cue for depth refinement. [25] incorporates cost volumes. [50] introduces transformer-based cross-modal attention, whereas [84] combines both convolutional and transformer operations. More recently, [62] presents a multi-scale approach with bilateral propagation and fusion of depth and image. Although these methods achieve high performance in controlled environments and on public benchmarks, distribution shifts arising from changes in illumination, sensor noise, or scene types can degrade performance, calling for robust adaptation strategies to new domains.

Domain Adaptation (DA) [13] aims to bridge the gap between data distributions in training (source) and testing (target). Unsupervised domain adaptation (UDA) [6, 17, 46, 70] typically assumes access to both source and target data, which frequently does not hold in practical depth completion scenarios. More recently, [27] introduces the source-free DA paradigm that rely solely on target data for adaptation [31], but requires multiple passes over the target. While these domain adaptation strategies have shown potential in tasks like classification or semantic segmentation [32, 56], few approaches specifically address depth completion [38], which often has to cope with not only domain shifts in image appearance, but also in sparse depth distribution (e.g., density changes in LiDAR points). Moreover, conventional DA methods typically require computationally expensive retraining or fine-tuning phases. Such constraints motivate test-time adaptation methods for depth completion that can adapt efficiently on-the-fly.

Test-Time Adaptation (TTA) [33] focuses on adapting models to new input distributions at test time, without access to the original source (training) data at test time. Unlike UDA, TTA assumes a single pass over the target dataset, making it well-suited for real-time robotic perception or autonomous driving. Early TTA works [28, 37, 61] used auxiliary self-supervised tasks (e.g., rotation prediction) to

guide adaptation. More recent methods update statistics in batch normalization layers [55], minimize entropy [68] for classification tasks. Follow-ups include strategies for stable parameter updates [43, 89], leveraging pseudo-label refinement [69] or contrastive objectives [3]. The work most closely related to ours are concurrent work [90] and TEA [85], which transforms pretrained classifiers into energy-based models to learn the underlying target distribution. However, most TTA approaches emphasize classification or segmentation; only a few extend to depth estimation or *completion* [44, 83], which deals with a continuous regression objective. A recent direction uses depth priors [26] derived from foundation models [53] for zero-shot generalization across different domains; however, while these generative methods achieve notable gains, they typically require non-trivial computational overhead at test time. Our approach offers a complementary alternative, focusing on a direct energy-based criterion that effectively bridges the gap across different domains for *depth completion*, while maintaining runtime efficiency.

Energy-Based Models (EBMs) are probabilistic models that define a probability distribution implicitly through an energy function, where lower value corresponds to higher likelihood. Unlike conventional probabilistic approaches, EBMs do not require an explicit computation of the normalization constant. Thus, density estimation with EBMs can be viewed as a nonlinear regression problem, allowing considerable flexibility in choosing any suitable nonlinear function as the energy function. Recent studies have demonstrated that combining EBMs with neural networks enables accurate modeling of complex data distributions, achieving significant success in various tasks such as image generation [8, 42, 65], discriminative learning [21], and density estimation [59, 67]. Despite these advancements, the use of EBMs for TTA remains relatively unexplored [85, 90], particularly in the context of depth completion, where no prior work currently exists. In this work, we propose a novel energy-based approach to explicitly bridge domain gaps by minimizing the estimated energy associated with erroneous predictions during TTA for depth completion.

3. Method Formulation

Depth Completion Models. Let $\Omega \subset \mathbb{R}^2$ denote the image space. An RGB image is defined as $I : \Omega \rightarrow \mathbb{R}^3$, a sparse depth map as $z : \Omega_z \subset \Omega \rightarrow \mathbb{R}_+$, and a dense depth map as $d : \Omega \rightarrow \mathbb{R}_+$. A depth completion model f_θ estimates a dense depth map from an RGB image and its corresponding sparse depth measurements. This is given by $\hat{d} = f_\theta(I, z)$, where $\hat{d} : \Omega \rightarrow \mathbb{R}_+$ is the estimated dense depth map.

We assume that a depth completion model f_θ is trained with supervision on a *source* dataset $\mathcal{D}_s = \{(I_s^{(n)}, z_s^{(n)}, d_s^{(n)})\}_{n=1}^{N_s}$, where d_s denotes the ground-truth

depth map. The model will be deployed to a *target* dataset $\mathcal{D}_t = \{(I_t^{(n)}, z_t^{(n)})\}_{n=1}^{N_t}$. We assume the underlying test-time data distribution is represented by \mathcal{D}_t and is out-of-distribution with respect to \mathcal{D}_s due to *covariate shifts*.

Energy-based Models. Let \mathcal{X} be the input (feature) space. Energy-based models (EBMs) define a function $E_\phi : \mathcal{X} \rightarrow \mathbb{R}$, which assigns a scalar value, *energy*, to each input sample $x \in \mathcal{X}$. Lower energy indicates higher likelihood that a sample x belongs to the distribution modeled by the EBM, which is represented by the Boltzmann distribution [34]:

$$p_\phi(x) = \frac{\exp(-E_\phi(x))}{Z(\phi)}, \quad (1)$$

where $Z(\phi) = \int_{\mathcal{X}} \exp(-E_\phi(x)) dx$ is the partition function, is parameterized by weight parameters ϕ .

3.1. Energy-based Model Formulation

Depth Energy-Based Models. We define an energy function E_ϕ , parameterized by ϕ , to quantify the likelihood of the depth predictions \hat{d} belonging to the source distribution. As the depth completion is an ill-posed task, we select for one of the infinitely many 3D scenes by conditioning on the given sparse depth input z ,

$$E_\phi : (\hat{d}; z) \rightarrow e, \quad e : \Omega \rightarrow [0, 1], \quad (2)$$

where e denotes energy. Instead of estimating a single energy value for the entire predicted depth map [85], we propose a region-based approach, where the energy model produces a lower-resolution energy map e , and each element of e corresponds to the likelihood of a local region in the full resolution depth map. Lower energy values indicate that the predicted depth \hat{d} is likely in-distribution (i.e., less erroneous), while higher energy values suggest potential errors (i.e., out-of-distribution predictions). This approach naturally imposes spatial regularity in the energy predicted, where we balance the ability to localize error but avoid spurious changes in the energy map. Additionally, by conditioning on the sparse depth map, we enable the energy model to validate the depth predictions against 3D coordinates measured within the environment, which aids in determining plausibility of the predicted 3D scene.

Mapping Distance to Energy. To train the energy model E_ϕ , we begin by modeling the target energy as proportional to the prediction error with respect to ground truth depth, for both the depth predictions \hat{d}_s on the original source data and \tilde{d}_s on the (synthesized) out-of-distribution data. For each patch $\Omega_p \subset \Omega$, we compute the mean squared error (MSE):

$$\Delta(\Omega_p) = \frac{1}{|\Omega_p|} \sum_{x \in \Omega_p} (\hat{d}_s(x) - d_s(x))^2. \quad (3)$$

for \hat{d}_s and likewise for \tilde{d}_s . We formulate the target energy y as the complement of the Gibbs probabilistic distribution:




Figure 1. *Overview.* Our energy-based test-time adaptation framework. The depth completion model predicts a dense depth map from an RGB image and sparse depth inputs. The energy model estimates a region-wise energy map of the dense depth map, which indicates the likelihood of the prediction associated with source data distribution. The adaptation layer is placed within the RGB encoder of the depth completion model to adapt the RGB features to the unseen target domain. During adaptation, only the adaptation layer (and optionally BatchNorm) is updated, while all other parameters remain frozen. The adaptation is guided by conventional unsupervised losses and the proposed energy minimization, which updates the adaptation layer to generate plausible 3D scene geometry in the prediction that align with scenes observed in the source data. Solid arrows denote data flow, while dashed arrows represent adaptation steps.

$$y = 1 - \exp(-\Delta/\tau), \quad (4)$$

where τ is a temperature term. Under this definition, small differences (minimal perturbations) yield energy values close to 0, while larger differences produce higher energy values near 1, marking unreliable or out-of-distribution predictions. This negative exponential mapping bounds energies within $[0, 1]$, matching the output range of E_ϕ . Note that the mapped energy is a distribution conditioned by depth completion parameters θ .

Generating out-of-distribution samples. While the energy model E_ϕ must learn to distinguish between in-distribution from out-of-distribution depth predictions, we only have access to the source dataset \mathcal{D}_s during training, or the preparation stage. Therefore, out-of-distribution examples need to be generated from source data. However, it is difficult to model test-time variations without making strong assumptions regarding the target distribution, which would limit the generalizability of the method. Instead, we repurpose *adversarial perturbations* [1, 20, 76] of the input RGB image I_s and sparse depth map z_s as a mechanism to explore the space of depth predictions $\hat{d}_s = f_\theta(I_s, z_s)$. This enables us to simulate the out-of-distribution examples that induce errors to the depth prediction. Adversarial perturbations adjust the inputs to statistical outliers in the blind high-density regions of the data space [20], which will cause out-of-distribution errors. This encompasses a broad range of variations that covers more than a single “target domain”. Hence, this enables the reuse of a single energy model across different testing dataset or distributions.

The adversarial perturbations correspond directly to the model’s failure modes, which presents the energy model E_ϕ

with out-of-distribution data that are assigned high energy as formulated by Eq. 4. We choose Fast Gradient Sign Method (FGSM) [20, 76] for perturbing the inputs (I_s, z_s) :

$$\tilde{I}_s = I_s + \epsilon_I \cdot \text{sign}(\nabla_{I_s} \mathcal{L}_{\text{sup}}(\hat{d}_s, d_s)) \quad (5)$$

$$\tilde{z}_s = z_s + \epsilon_z \cdot \text{sign}(\nabla_{z_s} \mathcal{L}_{\text{sup}}(\hat{d}_s, d_s)) \quad (6)$$

where ϵ_I and ϵ_z are parameters controlling the perturbation magnitudes, and \mathcal{L}_{sup} is the supervised loss between the predicted depth \hat{d}_s and ground-truth depth d_s . This is accomplished by a forward pass through the depth completion model f_θ with the original inputs (I_s, z_s) and computing the gradient with respect to I_s and z_s . The perturbations are constrained within an l_∞ ball of radius ϵ . After adding the perturbations to yield $(\tilde{I}_s, \tilde{z}_s)$, we make a second forward pass to obtain the synthesized out-of-distribution example $\tilde{d}_s = f_\theta(\tilde{I}_s, \tilde{z}_s)$, from which we compute the target energy y (see Eqs. (3) and (4)) for training our energy model.

Training Energy-based Model. Since marginalizing over $Z(\phi)$ is intractable, we instead pose the learning of an energy model as a discriminative task. Therefore, we optimize the energy model E_ϕ by minimizing cross entropy over all patches between the predicted energy map e and the target energy map y , which enables us to model the probability of in- (low energy) and out-of-distribution (high energy) data:

$$\mathcal{L}_{\text{energy}} = -\frac{1}{|\Omega_p|} \sum_{x \in \Omega_p} y(x) \log(y(x)/e(x)) \quad (7)$$

Once trained, the energy-based model will assign lower energy values to predictions aligned with the source data distribution and higher energy to erroneous or out-of-distribution predictions. Note: the energy model trained

on one depth completion model is not applicable for other models since the trained distribution is conditioned by the depth completion model parameters.

3.2. Test-Time Adaptation via Energy Guidance

The energy model E_ϕ is deployed in conjunction with the depth completion model f_θ to infer depth on the target testing data \mathcal{D}_t . Both the parameters of E_ϕ and f_θ are fixed or frozen. Following [44], we insert a lightweight adaptation module m_ψ in the pretrained image encoder of f_θ . Only the adaptation parameters m_ψ and BatchNorm statistics of f_θ are updated at inference. Fig. 1 shows an overview.

Energy-based Alignment. As high-energy test-time predictions $\hat{d}_t = f_{\theta,\psi}(I_t, z_t)$ correspond to out-of-distribution predictions, we posit that minimizing their energy will improve the fidelity of predictions by aligning them back to those of the source data distribution. Hence, minimizing the energy e predicted from E_ϕ is analogous to minimizing the likelihood of error learned from the source data. This motivates an energy-based adaptation loss that aims to reduce the energy of $\hat{d}_t(x)$ conditioned on $z_t(x)$:

$$\ell_{\text{energy}} = -\frac{1}{|\Omega_p|} \sum_{x \in \Omega_p} \log(1 - E_\phi(\hat{d}_t, z_t)(x)). \quad (8)$$

Sparse Depth Consistency. Sparse depth measurements provide partial yet accurate information about the 3D scene structure. To ensure the predicted depth map \hat{d}_t aligns accurately with these sparse measurements, we minimize the L_1 distance between the sparse depth points z_t and their corresponding predictions \hat{d}_t :

$$\ell_{\text{sparse}} = \frac{1}{|\Omega_z|} \sum_{x \in \Omega_z} |\hat{d}_t(x) - z_t(x)|, \quad (9)$$

where Ω_z denotes the set of pixels with valid sparse depth measurements.

Local Smoothness. Assuming that 3D scenes exhibit local smoothness and connectivity, we apply a similar constraint to the predicted depth map \hat{d}_t . To achieve this, we use an L_1 penalty on the gradients of \hat{d}_t along the horizontal and vertical directions, represented by ∂_X and ∂_Y . To maintain discontinuities at the object boundaries, we discount these penalties with spatially varying weights λ_X and λ_Y , which are defined based on image gradients. Specifically, we set $\lambda_X(x) = e^{-|\partial_X I_t(x)|}$ and $\lambda_Y(x) = e^{-|\partial_Y I_t(x)|}$.

$$\ell_{\text{smooth}} = \frac{1}{|\Omega|} \sum_{x \in \Omega} \lambda_X(x) |\partial_X \hat{d}_t(x)| + \lambda_Y(x) |\partial_Y \hat{d}_t(x)|. \quad (10)$$

Adaptation Objective. Finally, we integrate the three aforementioned loss components into a unified adaptation objective. Our adaptation loss is a linear combination:

$$\mathcal{L}_{\text{adapt}} = w_e \ell_{\text{energy}} + w_z \ell_{\text{sparse}} + w_s \ell_{\text{smooth}}, \quad (11)$$

where the weights w balance the contributions of each loss term during test-time optimization. We will update the adaptation parameters ψ by minimizing Eq. (11) on each sequentially observed example or batch of testing data D_t .

4. Experiments

4.1. Experimental Setup

Datasets. We evaluate our method on a diverse range of datasets, spanning both real-world and synthetic domains. For adaptations in *indoor* settings, we use: **VOID** [73], a real-world dataset captured using XIVO [14] (e.g., classrooms, laboratory, etc.); **SceneNet** [40], a synthetic dataset with diverse room layouts and furniture arrangements; **ScanNet** [7], a large-scale indoor dataset captured with Structure Sensor; **NYUv2** [41], providing household, office, and commercial environments captured with a Microsoft Kinect. For *outdoor* driving scenarios, we use: **KITTI** [18], a widely-used autonomous driving benchmark captured in daytime conditions using Velodyne LiDAR sensors; **VKITTI** [15], a synthetic dataset designed to replicate and augment KITTI; **Waymo** [60], providing large-scale autonomous driving data across diverse driving scenarios. Further details on the dataset are provided in Supp. Mat.

Models. For evaluation, we employ ETA on four widely-used architectures for depth completion: MSG-CHN [30], NLSPN [45], CostDCNet [25] and BP-Net [62]. For each depth completion model, an energy-based model is trained on the same source dataset as the depth completion model. We evaluate performance using standard metrics in depth completion: Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE), both reported in *meters*. Exact equations of metrics are provided in Supp. Mat.

Baseline Methods. We compare ETA against four baseline methods for test-time adaptation: TENT [68], CoTTA [69], TEA [85], and ProxyTTA [44]. Since the original implementations of TENT and TEA are not directly applicable to our task, we employ adapted versions. Specifically, *BN Adapt* denotes a variant of TENT, which minimizes Eqs. (9) and (10) instead of entropy. Likewise for TEA, we interpret the sum of Eqs. (9) and (10) as “energy”, since the log-sum-exp of classifier logits cannot be computed in depth completion tasks. Further details can be found in Supp. Mat.

4.2. Main Results

We report the performance of ETA in both *outdoor* and *indoor* settings. For each scenario, the models are pre-trained on KITTI [18] and VOID [73], respectively.

Main Result. As shown in Tab. 1, ETA consistently achieves superior performance, significantly lowering both Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) across all tested domains. In the outdoor setting (e.g., VKITTI-FOG, nuScenes, and Waymo), ETA sur-

Model	Method	Outdoor (KITTI)						Indoor (VOID)					
		VKITTI-FOG		nuScenes		Waymo		NYUv2		SceneNet		ScanNet	
		MAE↓	RMSE↓	MAE↓	RMSE↓	MAE↓	RMSE↓	MAE↓	RMSE↓	MAE↓	RMSE↓	MAE↓	RMSE↓
MSG-CHN [30]	Pre-trained	2.842	6.557	3.331	6.449	1.107	2.962	1.040	1.528	0.281	0.645	0.687	1.201
	CoTTA [69]	0.730	3.330	3.157	6.434	0.655	2.213	0.876	1.148	0.223	0.498	0.619	1.141
	ProxyTTA-fast [44]	0.728	3.087	2.834	6.096	0.608	1.921	0.699	1.120	0.192	0.424	0.302	0.480
	ETA (Ours)	0.703	2.996	2.666	6.073	0.583	1.907	0.561	0.850	0.187	0.401	0.298	0.460
NLSPN [45]	Pre-trained	1.309	7.423	2.656	6.146	1.175	3.078	0.388	0.702	0.167	0.438	0.233	0.431
	BN Adapt [68]	0.775	3.465	2.928	8.209	0.494	1.921	0.147	0.271	0.120	0.345	0.082	0.181
	CoTTA [69]	0.767	3.799	2.650	6.242	0.933	2.763	0.390	0.704	0.205	0.540	0.234	0.496
	TEA [85]	0.735	3.417	2.841	6.667	1.055	2.942	0.256	0.411	0.206	0.415	0.132	0.226
	ProxyTTA [44]	0.686	2.666	2.589	6.006	0.477	1.598	0.124	0.240	0.113	0.333	0.074	0.166
	ETA (Ours)	0.545	2.617	2.359	5.927	0.472	1.568	0.105	0.204	0.109	0.311	0.067	0.154
CostDCNet [25]	Pre-trained	1.042	6.301	3.064	6.630	1.093	2.798	0.189	0.446	0.173	0.443	0.144	0.458
	BN Adapt [68]	0.729	3.413	2.288	6.338	0.469	1.572	0.136	0.338	0.134	0.385	0.068	0.164
	CoTTA [69]	0.756	3.686	2.676	6.099	0.689	2.140	0.147	0.376	0.136	0.405	0.101	0.322
	TEA [85]	0.786	3.697	2.841	6.667	0.738	1.748	0.156	0.278	0.206	0.415	0.933	0.301
	ProxyTTA [44]	0.512	2.735	2.062	5.509	0.466	1.580	0.095	0.203	0.125	0.357	0.068	0.162
	ETA (Ours)	0.508	2.629	2.048	5.439	0.463	1.569	0.089	0.189	0.117	0.364	0.059	0.152
BP-Net [62]	Pre-trained	0.893	3.926	2.787	6.500	0.692	2.534	0.234	0.405	0.271	0.488	0.123	0.196
	BN Adapt [68]	0.605	3.426	2.618	5.850	0.540	1.838	0.178	0.257	0.236	0.373	0.116	0.177
	CoTTA [69]	0.611	3.090	2.432	5.490	0.538	1.874	0.212	0.364	0.293	0.421	0.140	0.199
	TEA [85]	0.631	3.311	2.581	5.677	0.673	1.942	0.171	0.253	0.256	0.404	0.115	0.165
	ProxyTTA [44]	0.571	2.844	2.373	5.413	0.489	1.520	0.174	0.248	0.231	0.368	0.102	0.155
	ETA (Ours)	0.544	2.729	2.281	5.278	0.451	1.433	0.161	0.231	0.221	0.352	0.093	0.148

Table 1. *Quantitative results.* For outdoors, we adapt from KITTI to VKITTI-FOG, nuScenes, and Waymo; for indoors, from VOID to NYUv2, SceneNet, and ScanNet. We report MAE and RMSE in *meters*. red denotes best, orange second-best, and yellow third-best.

passes baseline methods, including the previous state-of-the-art ProxyTTA [44] by an average of 5.36% and 1.97% in MAE and RMSE, respectively. ProxyTTA specifically targets depth completion by learning a proxy embedding from the sparse depth modality, which is less sensitive to domain shifts, to guide adaptation of image features towards source-domain distributions. However, it uniformly adapts the adaptation layer parameters across the entire predicted depth map, and solely relies on the proxy embedding from sparse depth inputs. This does not account for local errors in the depth map prediction and can potentially mislead the adaptation of the regions not covered by sparse depth. In contrast, the proposed energy model predicts the likelihood of the depth predictions from dense depth map prediction conditioned on the sparse point cloud as an energy map. Hence, even on the regions missing sparse depth, the energy model can still quantify energy.

We also compare our method against TEA [85], which is the closest existing work to our method, using an energy-based objective originally designed for classification tasks. TEA does not explicitly address continuous regression tasks or localized correction of errors. ETA gains notable average improvement of 24.9% and 19.0% on MAE and RMSE against TEA. Similar improvements are observed

on nuScenes and Waymo, underscoring ETA’s adaptation ability to challenging outdoor conditions such as fog and varying illumination. The improvement comes from ETA is even more prominent when adapting models that initially exhibit a larger errors after being transferred to the target testing data (e.g., KITTI→VKITTI-FOG on MSG-CHN). This demonstrates ETA’s ability to align source and target distributions, and enable less performant models to produce high-fidelity outputs in the target domain. The latest depth completion models, such as BP-Net, see smaller but still consistent improvements, which aligns with our expectation that the energy-based model identifies fewer erroneous regions when the initial error is already low.

In indoor environments, including NYUv2, SceneNet, and ScanNet, ETA also demonstrates clear advantages, achieving an average of 10.13% and 10.33% improvement over the previous state-of-the-art in MAE and RMSE. Consistent with results from outdoors, we again improve by and 58.2% and 27.13% over TEA. These results demonstrate that our method can successfully guide model adaptation when transferred to a novel testing data distribution.

Generalization to Adverse Conditions. We further assess the robustness of ETA on a challenging subset of the Waymo dataset, which we denote as *Waymo-Hard*. It

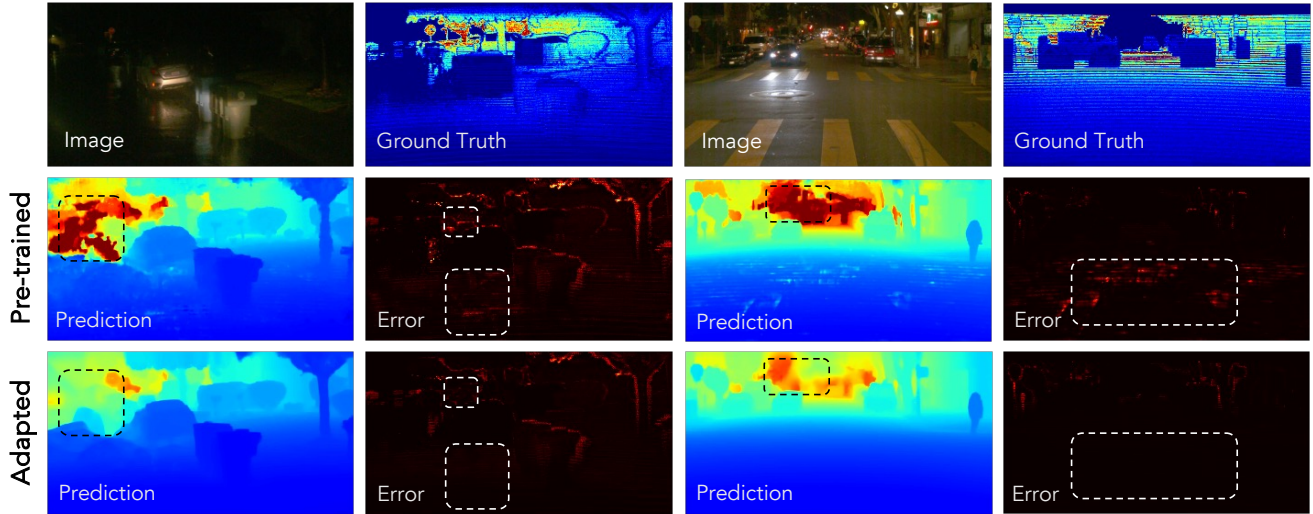


Figure 2. *Qualitative results on Waymo-Hard.* We adapt CostDCNet from KITTI→Waymo-Hard. ETA can effectively adapt to adverse weather conditions. Pre-trained denotes predictions using KITTI trained model, and adapted denotes results after TTA.

Method		MAE	RMSE
CostDCNet [25]	Pre-trained	0.776	2.072
	ProxyTTA	0.595	1.728
	ETA (Ours)	0.582	1.699

Table 2. *Quantitative results on Waymo-Hard.* We adapt CostDCNet from KITTI → Waymo-Hard. The best result in red.

Method		MAE	RMSE
MSG-CHN [30]	Pre-trained	2.842	6.557
	Global (Image-based)	1.406	4.226
	Local (Region-based)	0.703	2.996

Table 3. *Comparison of different ETA update strategies.* We adapt MSG-CHN from KITTI → VKITTI-FOG. The best result in red.

specifically contains night-time scenes, adverse weather, and highly dynamic scenarios with blurred moving vehicles and pedestrians. As shown in Tab. 2 and Fig. 2, ETA outperforms ProxyTTA in both MAE and RMSE, effectively bridging the domain gap by severe adverse conditions.

Generalization Across Large Domain Shifts. ETA is evaluated on a particularly challenging domain shift that we designed: Adapting a model from an outdoor driving environment to indoor scenes. Specifically, BP-Net model trained on KITTI is adapted to the NYUv2, SceneNet, and ScanNet datasets. As shown in Tab. 4, the baseline methods suffer from a significant performance degradation coming from the large gap in the prediction range and the appearance of the 3D scene. Nevertheless, ETA consistently achieves the best performance across all three indoor target domains, demonstrating its robustness even when the domain discrepancy is larger than the previously tested (e.g., indoor-to-indoor and outdoor-to-outdoor).

Model	Method	NYUv2		SceneNet		ScanNet	
		MAE↓	RMSE↓	MAE↓	RMSE↓	MAE↓	RMSE↓
BP-Net [62]	Pre-trained	1.987	2.578	1.432	1.872	2.657	3.416
	BN Adapt	1.343	1.839	0.423	0.634	0.325	0.480
	CoTTA	1.401	1.881	0.411	0.591	0.381	0.515
	TEA	1.392	1.856	0.409	0.617	0.348	0.501
	ProxyTTA	1.380	1.874	0.401	0.583	0.311	0.464
	ETA (Ours)	1.322	1.807	0.340	0.521	0.272	0.406

Table 4. *Quantitative results on outdoor-to-indoor adaptation.* We adapt BP-Net from KITTI to NYUv2, SceneNet and ScanNet. ETA consistently achieves the best results (red).

4.3. Ablation Studies

Energy Model with Varying Region size. Energy values allow us to make localized adjustments, and the area affected by the adjustment depends on the fidelity of energy predictions. An extreme case is to have a single energy value accounting the errors in the entire prediction, which ambiguates the regions with error. Our region-based energy formulation mitigates this issue by providing updates by minimizing the energy with areas in the granular level, as demonstrated in Fig. 4. Further, this region-based energy model provides a coarse level of interpretability in the adaptation process by highlighting the contribution of each prediction region in energy, shown in Fig. 3.

Sensitivity on Iterations. Increasing the number of adaptation iterations adds the computational budget but may offer better adaptation result. Fig. 5 demonstrates this sensitivity for NLSPN adapting from VOID to NYUv2. While optimal performance is achieved at three iterations in this case, we note that in most scenarios, a single iteration—the lower bound—is sufficient to attain state-of-the-art results.

Generation Quality. We assess the quality of our synthetic out-of-distribution (OOD) samples in modeling do-




Figure 3. *Qualitative evaluation of ETA across iterations.* We adapt from KITTI \rightarrow Waymo. The initial predictions were made with the pre-trained model. The estimated energy effectively highlights the regions to be corrected, guiding the model to make localized changes.




Figure 4. *Sensitivity to output energy map size.* We adapt MSG-CHN from KITTI \rightarrow VKITTI-FOG. Smaller patch sizes (e.g., 1×1) correspond to global updates, and larger sizes (e.g., 8×38) imply highly localized updates.

main gaps. For visualization, we adversarially perturb source data (VOID) to synthesize target (NYUv2) samples. The UMAP plot in Fig. 6 places synthetic OOD samples closer to target (NYUv2) than source (VOID) clusters, showing that our synthetic OOD data can effectively serve as a proxy for the target distribution.

5. Discussion

Our experiments show that ETA achieves superior adaptation performance in diverse indoor and outdoor scenarios, consistently outperforming test-time adaptation baselines by a noticeable margin. We attribute these gains to three main factors. First, the energy-based model offers a flexible way to quantify how well the predictions align with the training distribution. Second, the region-based formulation localizes errors and focuses on specific regions in need of correction. Third, restricting updates to a lightweight adaptation layer allows the rest of the network to remain stable, mitigating the risk of catastrophic forgetting that can arise during iterative test-time updates.

Notably, ETA demonstrates strong robustness in challenging environments such as foggy or nighttime outdoor scenes, and highly cluttered indoor scenarios. These conditions often impose large domain shifts that degrade depth completion accuracy. By detecting out-of-distribution regions and adjusting only a small set of parameters, our approach consistently brings improvement in both quantitative metrics (MAE, RMSE) and qualitative visualizations. Beyond depth completion, the energy-based test-time adaptation principle can be extended to other dense regression




Figure 5. *Sensitivity to Adaptation Iterations.* We adapt NLSPN from VOID \rightarrow NYUv2. Solid and dashed red lines represent baseline MAE and RMSE, respectively.




Figure 6. *UMAP plot of RGB features.* Adversarially perturbed examples from source data set (VOID) closely approximate the target domain (NYUv2), to bridge the domain gap.

tasks like optical flow, surface normal prediction, or medical image segmentation, where continuous outputs must be adapted to various lighting, weather, and sensor conditions.

Limitations. Although our method demonstrates notable improvements, certain limitations remain. First, our region energy approach avoids large-scale, uniform corrections across the entire image, yet some domain shifts may require broader scene-level updates. If the target environment exhibits drastically different global characteristics, the localized corrections alone may be insufficient to account for large variations in scale or depth ranges. Second, our energy-based model relies on adversarial perturbations to approximate out-of-distribution samples in the source domain. This design may be less effective if the real target domain features extreme conditions, sensor anomalies, or semantic shifts not captured by the perturbations.

Acknowledgments This work is supported by the NSF 2112562 Athena AI Institute, NIH/NHLBI R01HL121226, IITP-RS-2021-II211341 and NRF-RS-2023-00251366.

References

- [1] Zachary Berger, Parth Agrawal, Tian Yu Liu, Stefano Soatto, and Alex Wong. Stereoscopic universal perturbations across different architectures and datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15180–15190, 2022. 4
- [2] Marvin Chancán, Alex Wong, and Ian Abraham. 3d reprojection-driven robot navigation improves depth sensing. In *2025 International Conference on Advanced Robotics and Mechatronics (ICARM)*. IEEE, 2025. 2
- [3] Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *CVPR*, pages 295–305, 2022. 3
- [4] Xien Chen, Suchisrit Gangopadhyay, Michael Chu, Patrick Rim, Hyungseob Park, and Alex Wong. Uncle: Unsupervised continual learning of depth completion. *arXiv preprint arXiv:2410.18074*, 2024. 2
- [5] Xinjing Cheng, Peng Wang, Chenye Guan, and Ruigang Yang. Cspn++: Learning context and resource aware convolutional spatial propagation networks for depth completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10615–10622, 2020. 2
- [6] Safa Cicek and Stefano Soatto. Unsupervised domain adaptation via regularized conditional alignment. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1416–1425, 2019. 2
- [7] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 5
- [8] Yilun Du and Igor Mordatch. Implicit generation and modeling with energy-based models. In *NIPS*, 2019. 3
- [9] David Eigen, Christian Puhusch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014. 2
- [10] Abdelrahman Eldesokey, Michael Felsberg, and Fahad Shahbaz Khan. Propagating confidences through cnns for sparse data regression. *arXiv preprint arXiv:1805.11913*, 2018. 2
- [11] Abdelrahman Eldesokey, Michael Felsberg, Karl Holmquist, and Michael Persson. Uncertainty-aware cnns for depth completion: Uncertainty from beginning to end. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12014–12023, 2020. 2
- [12] Vadim Ezhov, Hyungseob Park, Zhaoyang Zhang, Rishi Upadhyay, Howard Zhang, Chethan Chinder Chandrappa, Achuta Kadambi, Yunhao Ba, Julie Dorsey, and Alex Wong. All-day depth completion. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024. 2
- [13] Abolfazl Farahani, Sahar Voghoei, Khaled Rasheed, and Hamid R Arabnia. A brief review of domain adaptation. *Advances in data science and information engineering: proceedings from ICDATA 2020 and IKE 2020*, pages 877–894, 2021. 1, 2
- [14] Xiaohan Fei, Alex Wong, and Stefano Soatto. Geosupervised visual depth prediction. *IEEE Robotics and Automation Letters*, 4(2):1661–1668, 2019. 2, 5
- [15] Adrien Gaidon, Qiao Wang, Yohann Cabon, and Eleonora Vig. Virtual worlds as proxy for multi-object tracking analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4340–4349, 2016. 5
- [16] Suchisrit Gangopadhyay, Jung-Hee Kim, Xien Chen, Patrick Rim, Hyungseob Park, and Alex Wong. Extending foundational monocular depth estimators to fisheye cameras with calibration tokens. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. 2
- [17] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189. PMLR, 2015. 2
- [18] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32:1231 – 1237, 2013. 5
- [19] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3828–3838, 2019. 2
- [20] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *CoRR*, abs/1412.6572, 2014. 4
- [21] Fredrik K. Gustafsson, Martin Danelljan, Goutam Bhat, and Thomas B. Schön. Energy-based models for deep probabilistic regression. In *ECCV*, 2020. 3
- [22] Mu Hu, Shuling Wang, Bin Li, Shiyu Ning, Li Fan, and Xiaojin Gong. Penet: Towards precise and efficient image guided depth completion. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13656–13662. IEEE, 2021. 2
- [23] Zixuan Huang, Junming Fan, Shenggan Cheng, Shuai Yi, Xiaogang Wang, and Hongsheng Li. Hms-net: Hierarchical multi-scale sparsity-invariant network for sparse depth completion. *IEEE Transactions on Image Processing*, 29: 3429–3441, 2019. 2
- [24] Maximilian Jaritz, Raoul De Charette, Emilie Wirbel, Xavier Perrotton, and Fawzi Nashashibi. Sparse and dense data with cnns: Depth completion and semantic segmentation. In *2018 International Conference on 3D Vision (3DV)*, pages 52–60. IEEE, 2018. 2
- [25] Jaewon Kam, Jungeon Kim, Soongjin Kim, Jaesik Park, and Seungyong Lee. Costdncnet: Cost volume based depth completion for a single rgb-d image. In *European Conference on Computer Vision*, pages 257–274. Springer, 2022. 2, 5, 6, 7
- [26] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *CVPR*, 2024. 3
- [27] Youngeun Kim, Donghyeon Cho, Kyeongtak Han, Priyadarshini Panda, and Sungeun Hong. Domain

- adaptation without source data. *IEEE Transactions on Artificial Intelligence*, 2(6):508–518, 2021. 2
- [28] Dong Lao, Yangchao Wu, Tian Yu Liu, Alex Wong, and Stefano Soatto. Sub-token vit embedding via stochastic resonance transformers. In *International Conference on Machine Learning*. PMLR, 2024. 2
- [29] Dong Lao, Fengyu Yang, Daniel Wang, Hyungseob Park, Samuel Lu, Alex Wong, and Stefano Soatto. On the viability of monocular depth pre-training for semantic segmentation. In *European Conference on Computer Vision*. Springer, 2024. 2
- [30] Ang Li, Zejian Yuan, Yonggen Ling, Wanchao Chi, Chong Zhang, et al. A multi-scale guided cascade hourglass network for depth completion. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 32–40, 2020. 2, 5, 6, 7
- [31] Jingjing Li, Zhiqi Yu, Zhekai Du, Lei Zhu, and Heng Tao Shen. A comprehensive survey on source-free domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8):5743–5762, 2024. 1, 2
- [32] Rui Li, Qianfen Jiao, Wenming Cao, Hau-San Wong, and Si Wu. Model adaptation: Unsupervised domain adaptation without source data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [33] Jian Liang, Ran He, and Tieniu Tan. A comprehensive survey on test-time adaptation under distribution shifts. *IJCV*, 133: 31–64, 2024. 2
- [34] EM Lifshitz and Lev Davidovich Landau. Statistical physics, course of theoretical physics. In *Part 2: Theory of the Condensed State*. Butterworth-Heinemann Pergamon, London, 1980. 3
- [35] Yuankai Lin, Tao Cheng, Qi Zhong, Wending Zhou, and Hua Yang. Dynamic spatial propagation network for depth completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1638–1646, 2022. 2
- [36] Tian Yu Liu, Parth Agrawal, Allison Chen, Byung-Woo Hong, and Alex Wong. Monitored distillation for positive congruent depth completion. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 35–53. Springer, 2022. 2
- [37] Yuejiang Liu, Parth Kothari, Bastien van Delft, Baptiste Bellot-Gurlet, Taylor Mordan, and Alexandre Alahi. TTT++: When Does Self-Supervised Test-Time Training Fail or Thrive? In *Advances in Neural Information Processing Systems*, pages 21808–21820. Curran Associates, Inc., 2021. 2
- [38] Adrian Lopez-Rodriguez, Benjamin Busam, and Krystian Mikolajczyk. Project to adapt: Domain adaptation for depth completion from noisy and sparse sensor data. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 2
- [39] Fangchang Ma, Guilherme Venturilli Cavalheiro, and Sertac Karaman. Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3288–3295. IEEE, 2019. 2
- [40] John McCormac, Ankur Handa, Stefan Leutenegger, and Andrew J Davison. Scenenet rgb-d: 5m photorealistic images of synthetic indoor trajectories with ground truth. *arXiv preprint arXiv:1612.05079*, 2016. 5
- [41] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, 2012. 5
- [42] Jiquan Ngiam, Zhenghao Chen, Pang Wei Koh, and Andrew Y. Ng. Learning deep energy models. In *ICML*, 2011. 3
- [43] Shuaicheng Niu, Jiayang Wu, Yifan Zhang, Yaofu Chen, Shijian Zheng, Peilin Zhao, and Mingkui Tan. Efficient Test-Time Model Adaptation without Forgetting. In *Proceedings of the 39th International Conference on Machine Learning*, pages 16888–16905. PMLR, 2022. ISSN: 2640-3498. 3
- [44] Hyungseob Park, Anjali Gupta, and Alex Wong. Test-time adaptation for depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20519–20529, 2024. 2, 3, 5, 6
- [45] Jinsun Park, Kyungdon Joo, Zhe Hu, Chi-Kuei Liu, and In So Kweon. Non-local spatial propagation network for depth completion. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pages 120–136. Springer, 2020. 2, 5, 6
- [46] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, pages 1406–1415, 2019. 2
- [47] Jiaxiong Qiu, Zhaopeng Cui, Yinda Zhang, Xingdi Zhang, Shuaicheng Liu, Bing Zeng, and Marc Pollefeys. DeepLidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3313–3322, 2019. 2
- [48] Chao Qu, Ty Nguyen, and Camillo Taylor. Depth completion via deep basis fitting. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 71–80, 2020. 2
- [49] Chao Qu, Wenxin Liu, and Camillo J Taylor. Bayesian deep basis fitting for depth completion with uncertainty. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16147–16157, 2021. 2
- [50] Kyeongha Rho, Jinsung Ha, and Youngjung Kim. Guideformer: Transformers for image guided depth completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6250–6259, 2022. 2
- [51] Patrick Rim, Hyungseob Park, Vadim Ezhov, Jeffrey Moon, and Alex Wong. Radar-guided polynomial fitting for metric depth estimation. *arXiv preprint arXiv:2503.17182*, 2025. 2
- [52] Patrick Rim, Hyungseob Park, Ziyao Zeng, Younjoon Chung, and Alex Wong. Protodepth: Unsupervised continual depth completion with prototypes. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6304–6316, 2025. 2
- [53] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 3

- [54] Ashutosh Saxena, Sung Chung, and Andrew Ng. Learning depth from single monocular images. *Advances in neural information processing systems*, 18, 2005. 2
- [55] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation, 2020. arXiv:2006.16971. 3
- [56] Inkyu Shin, Yi-Hsuan Tsai, Bingbing Zhuang, Samuel Schulter, Buyu Liu, Sparsh Garg, In So Kweon, and Kuk-Jin Yoon. Mm-tta: multi-modal test-time adaptation for 3d semantic segmentation. In *CVPR*, pages 16928–16937, 2022. 2
- [57] Shreyas S Shivakumar, Ty Nguyen, Ian D Miller, Steven W Chen, Vijay Kumar, and Camillo J Taylor. Dfuset: Deep fusion of rgb and sparse depth information for image guided dense depth completion. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 13–20. IEEE, 2019. 2
- [58] Akash Deep Singh, Yunhao Ba, Ankur Sarker, Howard Zhang, Achuta Kadambi, Stefano Soatto, Mani Srivastava, and Alex Wong. Depth estimation from camera image and mmwave radar point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9275–9285, 2023. 2
- [59] Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. In *Conference on Uncertainty in Artificial Intelligence*, 2019. 3
- [60] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 5
- [61] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-Time Training with Self-Supervision for Generalization under Distribution Shifts. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9229–9248. PMLR, 2020. ISSN: 2640-3498. 2
- [62] Jie Tang, Fei-Peng Tian, Boshi An, Jian Li, and Ping Tan. Bilateral propagation network for depth completion. In *CVPR*, 2024. 2, 5, 6, 7
- [63] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant cnns. In *2017 international conference on 3D Vision (3DV)*, pages 11–20. IEEE, 2017. 2
- [64] Rishi Upadhyay, Howard Zhang, Yunhao Ba, Ethan Yang, Blake Gella, Sicheng Jiang, Alex Wong, and Achuta Kadambi. Enhancing diffusion models with 3d perspective geometry constraints. *ACM Transactions on Graphics (TOG)*, 42(6):1–15, 2023. 2
- [65] Aäron Van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016. 3
- [66] Wouter Van Gansbeke, Davy Neven, Bert De Brabandere, and Luc Van Gool. Sparse and noisy lidar completion with rgb guidance and uncertainty. In *2019 16th international conference on machine vision applications (MVA)*, pages 1–6. IEEE, 2019. 2
- [67] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011. 3
- [68] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *ICLR*, 2021. 3, 5, 6
- [69] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continuous test-time domain adaptation. In *CVPR*, 2022. 3, 5, 6
- [70] Garrett Wilson and Diane J Cook. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5):1–46, 2020. 1, 2
- [71] Alex Wong and Stefano Soatto. Bilateral cyclic constraint and adaptive regularization for unsupervised monocular depth prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5644–5653, 2019. 2
- [72] Alex Wong and Stefano Soatto. Unsupervised depth completion with calibrated backprojection layers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12747–12756, 2021. 2
- [73] Alex Wong, Xiaohan Fei, Stephanie Tsuei, and Stefano Soatto. Unsupervised depth completion from visual inertial odometry. *IEEE Robotics and Automation Letters*, 5(2): 1899–1906, 2020. 2, 5
- [74] Alex Wong, Safa Cicek, and Stefano Soatto. Learning topology from synthetic data for unsupervised depth completion. *IEEE Robotics and Automation Letters*, 6(2):1495–1502, 2021.
- [75] Alex Wong, Xiaohan Fei, Byung-Woo Hong, and Stefano Soatto. An adaptive framework for learning unsupervised depth completion. *IEEE Robotics and Automation Letters*, 6(2):3120–3127, 2021. 2
- [76] Alex Wong, Mukund Mundhra, and Stefano Soatto. Stereopagnosia: Fooling stereo networks with adversarial perturbations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2879–2888, 2021. 4
- [77] Yangchao Wu, Tian Yu Liu, Hyoungseob Park, Stefano Soatto, Dong Lao, and Alex Wong. Augundo: Scaling up augmentations for monocular depth completion and estimation. In *European Conference on Computer Vision*, pages 274–293. Springer, 2024. 2
- [78] Chao Xia, Chenfeng Xu, Patrick Rim, Mingyu Ding, Nanning Zheng, Kurt Keutzer, Masayoshi Tomizuka, and Wei Zhan. Quadric representations for lidar odometry, mapping and localization. *IEEE Robotics and Automation Letters*, 8(8):5023–5030, 2023. 2
- [79] Yichen Xie, Chenfeng Xu, Marie-Julie Rakotosaona, Patrick Rim, Federico Tombari, Kurt Keutzer, Masayoshi Tomizuka, and Wei Zhan. Sparsefusion: Fusing multi-modal sparse representations for multi-sensor 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17591–17602, 2023. 2
- [80] Yan Xu, Xinge Zhu, Jianping Shi, Guofeng Zhang, Hujun Bao, and Hongsheng Li. Depth completion from sparse lidar data with depth-normal constraints. In *ICCV*, 2019. 2

- [81] Zhiqiang Yan, Kun Wang, Xiang Li, Zhenyu Zhang, Jun Li, and Jian Yang. Desnet: Decomposed scale-consistent network for unsupervised depth completion. In *Proceedings of the AAAI conference on artificial intelligence*, pages 3109–3117, 2023. 2
- [82] Yanchao Yang, Alex Wong, and Stefano Soatto. Dense depth posterior (ddp) from single image and sparse range. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3353–3362, 2019. 2
- [83] Teresa Yeo, Oğuzhan Fatih Kar, Zahra Sodagar, and Amir Zamir. Rapid network adaptation: Learning to adapt neural networks using test-time feedback. In *ICCV*, 2023. 3
- [84] Zhang Youmin, Guo Xianda, Poggi Matteo, Zhu Zheng, Huang Guan, and Mattoccia Stefano. Completionformer: Depth completion with convolutions and vision transformers. *arXiv preprint arXiv:2304.13030*, 2023. 2
- [85] Yige Yuan, Bingbing Xu, Liang Hou, Fei Sun, Huawei Shen, and Xueqi Cheng. Tea: Test-time energy adaptation. In *CVPR*, 2024. 3, 5, 6
- [86] Ziyao Zeng, Jingcheng Ni, Daniel Wang, Patrick Rim, Youn-joon Chung, Fengyu Yang, Byung-Woo Hong, and Alex Wong. Priordiffusion: Leverage language prior in diffusion models for monocular depth estimation. *arXiv preprint arXiv:2411.16750*, 2024. 2
- [87] Ziyao Zeng, Daniel Wang, Fengyu Yang, Hyungseob Park, Stefano Soatto, Dong Lao, and Alex Wong. Worddepth: Variational language prior for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9708–9719, 2024.
- [88] Ziyao Zeng, Yangchao Wu, Hyungseob Park, Daniel Wang, Fengyu Yang, Stefano Soatto, Dong Lao, Byung-Woo Hong, and Alex Wong. Rsa: Resolving scale ambiguities in monocular depth estimators through language descriptions. *Advances in neural information processing systems*, 37, 2024. 2
- [89] Marvin Zhang, Sergey Levine, and Chelsea Finn. MEMO: Test Time Robustness via Adaptation and Augmentation. *Advances in Neural Information Processing Systems*, 35: 38629–38642, 2022. 3
- [90] Xiaoran Zhang, Byung-Woo Hong, Hyungseob Park, Daniel H. Pak, Anne-Marie Rickmann, Lawrence H. Staib, James S. Duncan, and Alex Wong. Progressive test time energy adaptation for medical image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025. 3
- [91] Yinda Zhang and Thomas Funkhouser. Deep depth completion of a single rgb-d image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 175–185, 2018. 2