

# Test-Time Adaptation for Depth Completion

Hyoungeob Park  
Yale Vision Lab

hyoungeob.park@yale.edu

Anjali Gupta  
Yale Vision Lab

anjali.gupta@yale.edu

Alex Wong  
Yale Vision Lab

alex.wong@yale.edu

## Abstract

*It is common to observe performance degradation when transferring models trained on some (source) datasets to target testing data due to a domain gap between them. Existing methods for bridging this gap, such as domain adaptation (DA), may require the source data on which the model was trained (often not available), while others, i.e., source-free DA, require many passes through the testing data. We propose an online test-time adaptation method for depth completion, the task of inferring a dense depth map from a single image and associated sparse depth map, that closes the performance gap in a single pass. We first present a study on how the domain shift in each data modality affects model performance. Based on our observations that the sparse depth modality exhibits a much smaller covariate shift than the image, we design an embedding module trained in the source domain that preserves a mapping from features encoding only sparse depth to those encoding image and sparse depth. During test time, sparse depth features are projected using this map as a proxy for source domain features and are used as guidance to train a set of auxiliary parameters (i.e., adaptation layer) to align image and sparse depth features from the target test domain to that of the source domain. We evaluate our method on indoor and outdoor scenarios and show that it improves over baselines by an average of 21.1%. Code available at [github.com/seobbro/TTA-depth-completion](https://github.com/seobbro/TTA-depth-completion).*

## 1. Introduction

Reconstructing the 3-dimensional (3D) structure of an environment can support a number of spatial tasks, from robotic navigation and manipulation to augmented and virtual reality. Most systems addressing these tasks are built for sensor platforms equipped with range (i.e., lidar or radar) or optics (i.e., camera or sensors). While range sensors can measure the 3D coordinates of the surrounding space, they often yield point clouds that are sparse. Likewise, these coordinates can also be estimated from images by means of Structure-from-Motion (SfM) or Visual Inertial Odometry (VIO). For the goal of dense mapping, depth completion is the task of

recovering the dense depth of a 3D scene as observed from a sparse point cloud, which is often post-processed into a sparse depth map by projecting the points onto the image plane, and guided by a synchronized calibrated image.

Training a depth completion model can be done in a supervised (using ground truth) or unsupervised (using SfM) manner. The former dominates in performance, but requires expensive annotations that are often unavailable; the latter uses unannotated images, but they must satisfy SfM assumptions between frames, i.e., motion, covisibility, etc. Like most learning-based methods, models trained under both paradigms typically experience a performance drop when tested on a new dataset due to a covariate shift, i.e., domain gap. As we can only assume that a single pair of image and sparse depth map is available in the target domain for the depth completion, models belonging to either learning paradigms cannot easily be trained or adapted to the new domain even when given the testing data. We focus on test-time adaptation (TTA) for depth completion, where one is given access to the test data in a stream, i.e., one batch at a time, without being able to revisit previously-seen examples. The goal is to learn causally and to quickly adapt a set of pre-existing weights trained on a source domain to a target test domain, so one can reduce the performance gap.

We begin with some motivating observations on the effects of the domain gap: (i) Errors in target domain tend to be higher when feed both the image and sparse depth as input rather than *sparse depth only*, as shown in Fig. 1. This implies that the depth modality exhibits a smaller covariate shift between the source and target domains than the image modality, to the extent that forgoing the image altogether often yields superior results than using either both sparse depth and image or the image alone. (ii) Yet, when operating in the source domain, we observe the opposite effect – forgoing the image is detrimental to performance. Naturally, this begs the question: How should one leverage data modalities that are less sensitive to the domain shift (e.g., sparse depth) to support alignment between source and target domains for modalities that are more sensitive (e.g., RGB image)?

To answer this question, we investigate a test-time adaptation approach that learns an embedding for guiding the

model parameter update by exploiting the data modality (sparse depth) that is less sensitive to the domain shift. The embedding module maps the latent features encoding sparse depth to the latent features encoding both image and sparse depth. The mapping is trained in the source domain and frozen when deployed to the target domain for adaptation. During test time, sparse depth is first fed through the encoder and mapped, through the embedding module, to yield a proxy for image and sparse depth embeddings from the source domain – we refer to the embedded sparse depth features as *proxy embeddings*. Note: As the mapping is learned in the source domain, the proxy embeddings will also follow the distribution of source image and sparse depth embeddings. Next, both image and sparse depth from the target test domain are fed as input to the encoder. By maximizing the similarity between test-time input embeddings and the proxy embeddings, we align the target distribution to that of the source to reduce the domain gap. In other words, our method exploits a proxy modality for guiding test-time adaptation and we call the approach, *ProxyTTA*. When used in conjunction with typical loss functions to penalize discrepancies between predictions and input sparse depth, and abrupt depth transitions, i.e., Total Variation, the embeddings serve as regularization to guide the model parameter update and prevent excessive drift from those trained on the source data.

Following test-time adaptation conventions, we assume limited computational resources, and that inputs arrive in a stream of small batches and must be processed within a time budget without access to the past data. To ensure fast model updates under these constraints, we deploy auxiliary parameters, or an adaptation layer, to be updated while freezing the rest of the network – thus achieving low-cost adaptation. We demonstrate our method in both indoor (VIO) and outdoor (lidar) settings across six datasets, where we not only target typical adaptation scenarios where the shift exists between real and synthetic data domains with similar scenes, i.e. from KITTI [39] to Virtual KITTI [11], but also between different scene layouts, i.e., from VOID [45] to NYUv2 [26] and SceneNet [25]. Our proxy embeddings consistently improve over baselines by an average of 21.09% across all methods and datasets. To the best of our knowledge, we are the first to introduce test-time adaptation for depth completion.

## 2. Related work

**Test Time Adaptation** (TTA) aims to adapt a given model, pretrained on source data, to test data without access to the source training data. Related fields along this vein include unsupervised domain adaptation [12, 28], which utilizes source domain data (in practice, this may not be available) for adaptation, and source-free domain adaptation [19], which does not assume access to source data, but allows access to test data on multiple passes. In contrast, we focus on test-time adaptation where we do not have access to source

data and must adapt to test data in a single pass.

Previous studies have proposed strategies to select the source model’s component to be preserved, such as the class prototypes extracted from the source data [5, 22, 33], the subset of source model parameters [16, 42], and the discriminative feature from the self-supervised learning (SSL) [5]. For instance, [42] proposes TENT, a simple but effective batch-norm layer adaptation with entropy minimization for fully test-time adaptation. TTT [38] performs classification-layer adaptation by updating the last linear layer of the source model; [22] extends this with TTT++ and utilizes joint task-specific and model-specific information based on self-supervised learning. [16] presents T3A, an optimization-free classifier adjustment module. [5] uses shift-agnostic weight regularization (SWR) to prevent an effect from the erroneous signal in test time, jointly with the nearest source prototype classifier and a self-supervised proxy task. [2] proposes a contrastive learning with an online pseudo-label refinement while [33] proposes pseudo-label refinement and momentum update for 3D point cloud segmentation. [43] proposes continual test-time adaptation based on stochastic restoration and weight-averaged pseudo-labels. [35] uses efficient residual modules to realign the pretrained weights. The above methods largely focus on single-image-based tasks, i.e., classification [2, 5, 22] and semantic segmentation [33], and rely on entropy constraints from [42].

Unlike prior work on classification [22, 38] and segmentation [33], depth completion is a regression problem; hence, existing methods using entropy-based objectives [42], which operate on logits, are not applicable in this task. Instead, we propose to minimize sparse depth reconstruction and local smoothness objectives – similar to that of some existing unsupervised methods [44, 45] – and to maximize cosine similarity between the proxy embeddings and the test time image and sparse depth embeddings.

**Depth Completion** aims to output dense depth from a single image and synchronized point cloud, i.e., from lidar or tracked by VIO, projected onto a sparse depth map, by multimodal fusion [8, 45, 50–52].

*Unsupervised depth completion* approaches rely on Structure-from-Motion (SfM) for training and require access to an auxiliary training dataset containing stereo image pairs [34] or monocular videos [24, 45–48] with synchronized sparse depth maps. Typically, they minimize a linear combination of photometric reprojection consistency, sparse depth reconstruction error, and local smoothness [24, 45–47]. These methods can support online training, but are limited by the need for stereo or monocular videos with sufficient parallax and co-visibility. In contrast, our approach does not rely on SfM and can be used in more general scenarios.

*Supervised depth completion* trains the model by minimizing a loss with respect to ground truth. [3, 15] focus on network operations and designs to effectively deal with

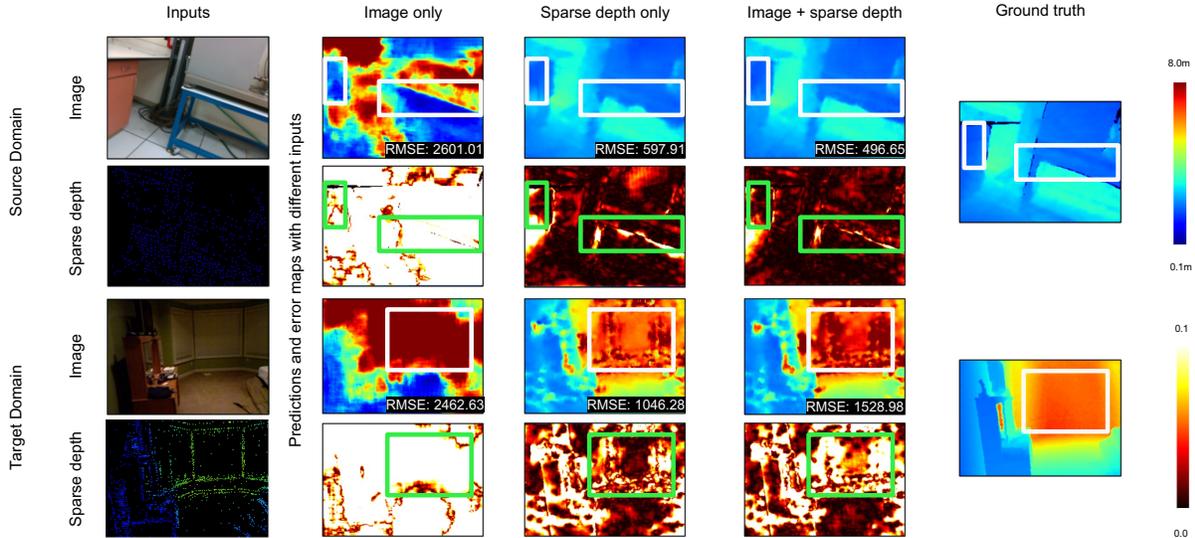


Figure 1. *Model sensitivity to input modalities.* While utilizing both sparse depth and image as input, the best performance is achieved in the source domain (VOID). Yet, forgoing the image in the test domain (NYUv2) often yields lower error than using both as input.

sparse inputs. [17, 24, 53] propose early and late fusion of image and depth encoder features while [14] uses separate networks for each. [20] proposes a multi-scale cascaded hourglass network to enhance the depth encoder with image features. [4] proposes convolutional spatial propagation network; [27] extends it to non-local spatial propagation to refine an initial depth map based on confidence and learnable affinity; [21] further extends it to dynamic spatial propagation. [9, 10, 30, 31] learn uncertainty of the depth estimates. [41] utilizes confidence maps to combine depth predictions while [29, 49, 55] use the surface normals to guide depth prediction. [18] incorporates cost volume for depth prediction. [32] devises transformer architecture with cross-modal attention, and [54] proposes a hybrid convolution-transformer architecture for depth completion.

While both unsupervised and supervised methods have demonstrated strong performance on benchmarks, they often fail to generalize to test datasets with large domain discrepancies. Moreover, obtaining ground truth is unrealistic for real-time applications, and accumulating sufficient parallax incurs large latencies – presenting significant challenges for online adaptation. Unlike past works, we do not assume access to ground truth nor data outside of the input.

**Unsupervised Domain Adaptation** (UDA) addresses the discrepancy between labeled source data and unlabeled target data [6, 12, 28, 37]. The only existing UDA depth completion method [23] models the domain gap as the noise in sparse points and the appearance in images. Unlike most UDA approaches that require source data during adaptation, we are only given the inputs necessary for inference in a stream without the ability to revisit past data, and must update the model online under a limited computational budget.

**Our Contributions.** We present (i) a study on how the

domain shift in each data modality (e.g., image and sparse depth) affects model performance when transferring it from source to target test domain. This study motivates (ii) our approach to learn an embedding of sparse depth features (which are less sensitive to the domain shift) that serves as proxy to source features for guiding test-time adaptation. (iii) To the best of our knowledge, we are the first to propose test-time adaptation for the depth completion task, and (iv) will release code, models, and dataset benchmarking setup to make development accessible for the research community.

### 3. Method Formulation

For ease of use, we assume access to a (source) pretrained depth completion model  $f_\theta$  that infers a dense depth map  $\hat{d}$  from a calibrated RGB image  $I \in \mathbb{R}^{H \times W \times 3}$  and its associated sparse point cloud projected onto the image plane as a sparse depth map  $z \in \mathbb{R}_+^{H \times W}$ , i.e.,  $f_\theta(I, z) \rightarrow \hat{d} \in \mathbb{R}^{H \times W}$ . For simplicity, we assume that the model was trained to minimize a supervised loss between prediction and ground truth  $d \in \mathbb{R}_+^{H \times W}$  on a source dataset  $\mathcal{D}_s = \{I_s^{(n)}, z_s^{(n)}, d^{(n)}\}_{n=1}^{N_s}$ , where  $N_s$  indicates the number of data samples. Following conventions in TTA, we assume access to the source domain dataset prior to deployment.

During test-time adaptation, we follow the protocol of [22, 38], where we only have access to the target domain data  $\mathcal{D}_t = \{I_t^{(n)}, z_t^{(n)}\}_{n=1}^{N_t}$  and utilize an online procedure to adapt to unseen  $N_t$  target data samples. Note that we make no assumptions about supervision during test-time; hence, while we present results on supervised methods for controlled experiments, we see our method being applicable towards unsupervised methods as well.

Our method, ProxyTTA, is split into three stages (Fig. 3):

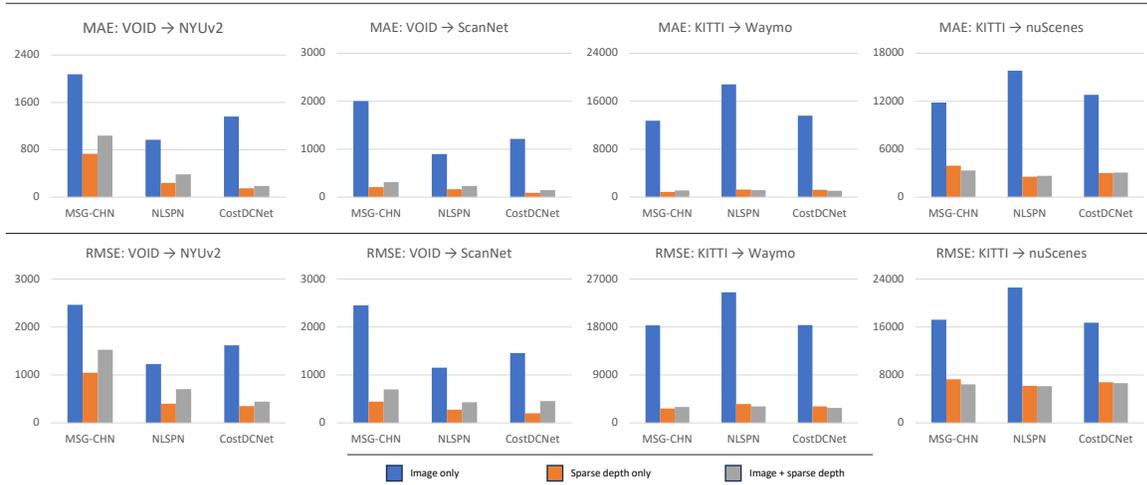


Figure 2. *Model sensitivity to input modalities.* Depth completion networks have a high reliance on sparse depth modality. Performing inference in a novel domain without the RGB image, i.e., using just sparse depth as input, can improve over using both data modalities.

(a) During an initialization stage, we augment the network encoder with an adaptation layer and train it using source domain data. (b) In the preparation stage, we learn a mapping from sparse depth features to image and sparse depth (proxy) embeddings. (c) During test time, we do not need the source dataset; we freeze the mapping and use its proxy embeddings for updating the adaptation layer parameters in test domain.

### 3.1. Sensitivity Study on Data Modalities

To motivate our approach, we begin with a sensitivity study of depth completion networks to input modalities, e.g. image, sparse depth, and the effect of domain shift on them. To this end, we alter the inputs by zeroing out either  $I$  or  $z$  to yield  $(I, z)$ ,  $(I_0, z)$ , and  $(I, z_0)$ , where  $I_0$  and  $z_0$  indicate the zero matrices with identical size to  $I$  and  $z$ , respectively. We evaluate the pretrained models using  $(I, z)$ ,  $(I_0, z)$ , and  $(I, z_0)$  to highlight their dependence on each input modality and to gauge their sensitivity when one modality gives no useful information at all. Fig. 1 and Fig. 2 show qualitative (error maps) and quantitative results (bar graphs), respectively, of pretrained depth completion models when fed the different inputs on the source dataset  $\mathcal{D}_s$  and the target dataset  $\mathcal{D}_t$ .

In the source domain, inference using both image and sparse depth as inputs, i.e.,  $\hat{d}_s(I, z)$ , shows the best performance. Surprisingly, the inference using sparse depth alone (i.e., with null-image)  $\hat{d}_s(I_0, z)$  is comparable to  $\hat{d}_s(I, z)$ . This shows the first intuition behind our approach: (i) Even though depth inputs are sparse, they are sufficient to support the reconstruction of the scene. Additionally, inference with image alone (i.e., null-depth)  $\hat{d}_s(I, z_0)$  is worse than  $\hat{d}_s(I_0, z)$  and  $\hat{d}_s(I, z)$ , which suggests that a depth completion network relies heavily on sparse depth modality for inference, and the image for guiding recovery of finer details.

In the target test domain, expectedly, performance degrades for inference using both image and sparse depth due to a covariate shift. Remarkably, we observe that predic-

tions from sparse depth alone  $\hat{d}_t(I_0, z)$  remain consistent in performance to those using both inputs  $\hat{d}_t(I, z)$ . Moreover, we observe that in most cases  $\hat{d}_t(I_0, z)$ , in fact, outperforms  $\hat{d}_t(I, z)$  across several methods and datasets, i.e., inference without image information in the test domain is better than with it. Conversely, the performance gap between inference with both inputs,  $\hat{d}_t(I, z)$ , and just the image,  $\hat{d}_t(I, z_0)$ , becomes more evident under the domain shift. This observation illustrates another intuition: (ii) The domain shift largely affects the image modality, and less so depth.

The two intuitions above motivate our approach. As object shapes tend to persist across domains, and the measured sparse points being a coarse representation of them, we aim to leverage sparse depth modality to bridge the domain gap. To this end, we exploit the observation that depth completion networks are able to recover (coarse) 3D scenes from sparse points alone and that the image serves to propagate and refine depth for regions lacking points. This is done by learning to map features encoding sparse depth inputs to features encoding both modalities in the source domain and, during test-time, recover the source domain features compatible with target domain sparse depth to guide model adaptation. Specifically, as observed, the covariate shift is largely photometric, so we propose to adapt the RGB image encoder branch by introducing a *adaptation layer*: a single convolutional layer designed to align target domain RGB embeddings to those of the source domain. As the rest of the network is frozen, adapting just the adaptation layer allows for low-cost model updates.

**Intuition for integrating adaptation layer.** Guided by our observations, the adaptation layer should be (i) placed in the image encoder branch prior to the fusion of image and depth features, and (ii) located within later layers to modulate higher level representations (i.e., object shapes, as opposed to low-level edges). (iii) connected as a skip connection to decoder to more directly affect the output.

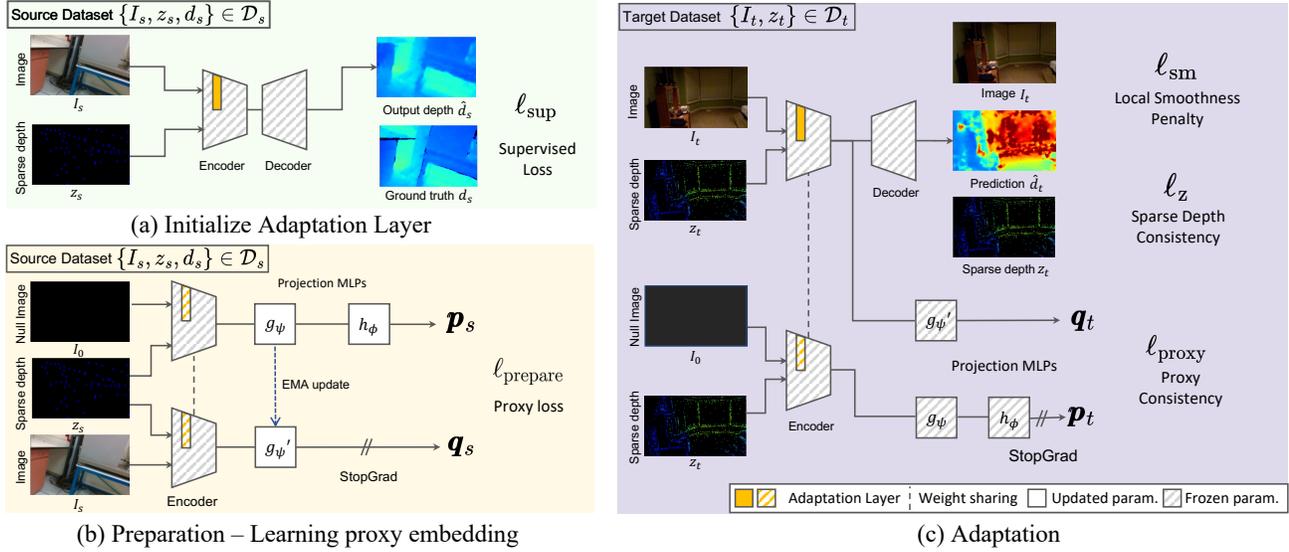


Figure 3. *Overview.* (a) The pretraining stage integrates an adaptation layer into a pretrained encoder and pretrains the adaptation layer on the source dataset. (b) The preparation stage learns the proxy mapping of features encoding sparse depth to those encoding both inputs. (c) The adaptation stage deploys the model to the target domain and updates the adaptation layer by leveraging proxy embeddings as guidance.

### 3.2. Preparation Stage - Source Domain

**Initialize adaptation layer from source domain.** Updating the entire network is largely infeasible in test-time adaptation scenarios. For the sake of speed and efficiency, we implement an adaptation layer  $m_\phi$ , i.e., a convolutional layer, within the encoder of a pretrained network. Note that the entire network will be frozen during all stages of our method with the exception of the adaptation layer and proxy mapping, where both will be initialized during preparation stage in the source domain; the proxy mapping will then be frozen and used to adapt  $m_\phi$  in the target domain. To ease the adaptation process, we initialize the  $m_\phi$  by minimizing a supervised loss over the source dataset (Fig. 3-(a)). We denote the pretrained encoder integrated with  $m_\phi$  as  $e_\phi$ .

**Learning proxy mapping from source domain.** As observed in Fig. 1, the best results in the source domain are achieved by feeding in both the image and sparse depth modalities for inference. However, the image is susceptible to domain shift which degrades performance when the model is transferred to an unseen test domain. Conversely, sparse depth is more resilient to the domain shift than RGB images, i.e., the shape of a car (or another object) remains similar regardless of (synthetic or real) domain. Our method aims to leverage the sparse depth modality, which is less sensitive to the domain shift, in the downstream adaptation process. To this end, we employ a soft mapping [13] from just the encoded sparse depth features to sparse depth *and image features* to learn the photometric information that is captured from the same scene as the sparse point cloud. This strategy allows us to learn the mapping that projects the sparse depth features to “proxy” embeddings close to those that also en-

code the image. In other words, it fills in what is missing in the image encoder branch by predicting the residual latent image encoding that is compatible with the input sparse depth, i.e., 3D scene. As this is trained in the source domain, the mapping naturally yields proxy embeddings that encode the source domain image (and sparse depth), which can be later used to guide the adaptation layer  $m_\phi$  to transform test domain RGB features close to those of the source domain.

This mapping by MLPs can be denoted as  $g_\psi(\cdot)$ ,  $g_{\psi'}(\cdot)$  and  $h_\omega(\cdot)$ ; to learn them, we get two embeddings  $\mathbf{p}_s$  and  $\mathbf{q}_s$ ,

$$\begin{aligned} \mathbf{p}_s &= h_\omega(g_\psi(\text{StopGrad}(e_\phi(I_0, z_s)))) \\ \mathbf{q}_s &= \text{StopGrad}(g_{\psi'}(e_\phi(I_s, z_s))) \end{aligned} \quad (1)$$

where  $e_\phi$  denotes the encoder augmented with the adaptation layer trained on source dataset,  $I_s, z_s$ , the image and sparse depth from source domain, and  $I_0$  the null-image. The embedding modules  $g_\psi$  and  $h_\omega$  are updated to maximize the similarity between  $\mathbf{p}_s$  and  $\mathbf{q}_s$ . To learn them, we minimize:

$$\ell_{\text{prepare}} = 1 - \left( \frac{\mathbf{p}_s \cdot \mathbf{q}_s}{\|\mathbf{p}_s\| \cdot \|\mathbf{q}_s\|} \right), \quad (2)$$

where  $\|\cdot\|$  is  $L2$ -norm, and  $(\mathbf{a} \cdot \mathbf{b})$  indicates the dot product of the vectors  $\mathbf{a}$  and  $\mathbf{b}$ . To this end, we first train the MLP heads  $g_\psi, h_\omega$  by minimizing Eqn. 2. Note that the MLP head  $g_{\psi'}$  is updated with EMA update following BYOL [13] to avoid collapse:  $g_{\psi'} \leftarrow \tau \cdot g_{\psi'} + (1 - \tau) \cdot g_\psi$ .

Once the mapping is learned, we can freeze the embedding module and deploy it for test-time adaptation where we update the adaptation layer weights  $\phi$  to maximize the similarity between the embeddings of a test domain image and sparse depth, and its proxy from the source domain. Naturally, due to the domain shift, the embeddings will yield

low similarity scores; hence, maximizing the scores through our proxy embedding implicitly aligns the target RGB distribution to that of the source distribution, i.e., minimizing the cosine similarity between the source and target distributions.

### 3.3. Deploying Proxy Mapping to Target Domain

**Adaptation stage** aims to update the adaptation layer parameters by minimizing a test-time loss function over the target test domain data  $\{I_t, z_t\} \in \mathcal{D}_t$ . To do so, we deploy the learned proxy mapping module (MLP heads  $\{g_{\psi}^*(\cdot), g_{\psi'}^*(\cdot), \text{and } h_{\omega}^*(\cdot)\}$ ) along with the adaptation layer  $m_{\phi}$  integrated into the frozen encoder as  $e_{\phi}$ .

**Adaptation loss.** For adaptation, our loss is composed of a linear combination of three loss terms:

$$\mathcal{L}_{\text{adapt}} = w_z \ell_z + w_{\text{sm}} \ell_{\text{sm}} + w_{\text{proxy}} \ell_{\text{proxy}}, \quad (3)$$

where  $\ell_z, \ell_{\text{sm}}$  denote sparse depth consistency loss and local smoothness loss, respectively,  $\ell_{\text{proxy}}$  is proxy mapping consistency loss, and  $w$  indicates a weight of each loss term.

**Sparse Depth Consistency.** Sparse point clouds capture a coarse structure of the 3D scene. To obtain metric scale predictions consistent with the scene structure, we minimize L1 error between the sparse depth  $z_t$  and the prediction  $\hat{d}_t$ :

$$\ell_z = \frac{1}{|\Omega(z_t)|} \sum_{x \in \Omega(z_t)} |\hat{d}_t(x) - z_t(x)|, \quad (4)$$

where  $x \in \Omega(z_t)$  are the pixel locations where sparse points were projected onto the image plane.

**Local Smoothness.** Based on the assumption of local smoothness and connectivity in a 3D scene, we impose the same in the predicted depth map  $\hat{d}_t$ . Specifically, we apply an L1 penalty to its gradients in both the x- and y-directions (i.e.,  $\partial_X$  and  $\partial_Y$ ). We balance the weight of each term with  $\lambda_X$  and  $\lambda_Y$ , to allow discontinuities over object boundaries based on the image gradients, where  $\lambda_X(x) = e^{-|\partial_X I_t(x)|}$ ,  $\lambda_Y(x) = e^{-|\partial_Y I_t(x)|}$ , and  $\Omega$  denotes the image domain.

$$\ell_{\text{sm}} = \frac{1}{|\Omega|} \sum_{x \in \Omega} \lambda_X(x) |\partial_X \hat{d}_t(x)| + \lambda_Y(x) |\partial_Y \hat{d}_t(x)|. \quad (5)$$

**Proxy Consistency.** In order to regularize the adaptation with the learned mapping from the previous stage, we freeze the weight parameters of MLP heads  $\{g_{\psi}^*(\cdot), h_{\omega}^*(\cdot)\}$ , and update the parameters of the adaptation layer  $m_{\phi}$ . First, we obtain the features  $\mathbf{p}_t$  and  $\mathbf{q}_t$  using the null-image  $I_0$  in one and the given target test domain image  $I_t$  in the other:

$$\begin{aligned} \mathbf{p}_t &= \text{StopGrad}(h_{\omega}^*(g_{\psi}^*(e_{\phi}(I_0, z_t)))), \\ \mathbf{q}_t &= g_{\psi'}^*(e_{\phi}(I_t, z_t)). \end{aligned} \quad (6)$$

We maximize the cosine similarity between the feature  $\mathbf{q}_t$  and  $\mathbf{p}_t$  via a proxy loss  $\ell_{\text{proxy}}$  to update adaptation layer  $m_{\phi}$ :

$$\ell_{\text{proxy}} = 1 - \left( \frac{\mathbf{p}_t}{\|\mathbf{p}_t\|} \cdot \frac{\mathbf{q}_t}{\|\mathbf{q}_t\|} \right). \quad (7)$$

## 4. Experiments

We demonstrate the effectiveness of our approach on a mix of both real and synthetic datasets including indoor SLAM/VIO scenarios (VOID [45], NYUv2 [26], SceneNet [25], and ScanNet [7]) and outdoor driving scenarios using lidar sensor (KITTI [39], Virtual KITTI (VKITTI) [11], nuScenes [1], and Waymo Open Dataset [36]). We chose three representative architectures of current depth completion methods to test our method: MSG-CHN [20] (CNN-based), NLSPN [27] (SPN-based) and CostDCNet [18] (cost volume-based). All reported results are averaged over 5 independent trials. We describe implementation details, hyper-parameters used, hardware requirements, evaluation metrics as well as additional experimental results in the Supp. Mat.

**Main Result.** We use pretrained models (MSG-CHN, NLSPN, and CostDCNet) from the two source datasets, VOID for indoor, and KITTI for outdoor. For indoor, we adapt models pretrained on VOID to NYUv2, SceneNet, and ScanNet; for outdoors, we adapt from KITTI to VKITTI (with fog), nuScenes, and Waymo. BN Adapt denotes updating the batch statistics (i.e., running mean and variance). BN Adapt,  $\ell_z, \ell_{\text{sm}}$  is a variation of TENT [42] which minimizes Eqn. 4, 5 instead of entropy by updating learnable scale factors. CoTTA denotes replacing proxy loss with L1 consistency loss w.r.t. the pretrained prediction [43]. ProxyTTA-fast denotes our method without batch norm update, which improves adaptation runtime by 25.32%.

Our method consistently improves over baselines and variants of BN Adapt (Table 1). Specifically, we improve over BN Adapt,  $\ell_z, \ell_{\text{sm}}$  by 11.60% on average across all methods for indoor, 19.73% on outdoors, and 15.67% overall to achieve state-of-the-art performance. Qualitatively, Fig. 4 and Fig. 5 show that our method performs better in boundary regions and homogeneous regions, thus exhibiting less oversmoothing on curtains in Fig. 4-(a) and car in Fig. 5-(b), and undersmoothing on blackboard in Fig. 4-(d) and road in Fig. 5-(a), respectively, during adaptation. This trend is due to the proxy loss and the adaptation layer, which allows us to adapt with minimum weight adjustments while preserving high-level features (object shapes) learned from the source domain by mapping the target RGB modality to that of the source domain. Notably, ProxyTTA-fast still improves over BN Adapt even though we only adapt our adaptation layer, which demonstrates the effectiveness of our design choice as well as our proposed proxy embeddings. We visualize image and sparse depth features from the source and target domains along with proxy embeddings in the target domain using t-SNE [40] in Fig. 1 of Supp. Mat.; we observe that proxy embeddings are close to source domain features.

**Comparison to BN adaptation<sup>1</sup> and CoTTA.** To assess the impact of our adaptation layer, we compare to batch norm

<sup>1</sup>MSG-CHN lacks Batch Norm (BN) layer so we cannot use BN adapt.

Method	MAE	RMSE	MAE	RMSE	MAE	RMSE	
	KITTI → VKITTI-FOG		KITTI → nuScenes		KITTI → Waymo		
MSG-CHN	Pretrained	2842.88	6557.38	3331.821	6449.094	1107.22	2962.45
	CoTTA	<i>730.6±11.67</i>	<i>3330.23±44.83</i>	<i>3157.69</i>	<i>6434.14</i>	<i>655.77±30.98</i>	<i>2213.27±98.80</i>
	ProxyTTA-fast (Ours)	<b>728.24±3.73</b>	<b>3087.36±15.92</b>	<b>2834.08±17.64</b>	<b>6096.56±21.08</b>	<b>608.91±1.74</b>	<b>1921.83±2.54</b>
NLSPN	Pretrained	1309.99	7423.48	2656.609	6146.590	1175.83	3078.377
	BN Adapt	1140.21±35.89	4592.86±198.21	11291.57±21.32	16670.87±52.56	7283.33±104.58	9670.36±250.22
	BN Adapt, $\ell_z, \ell_{sm}$	775.20±5.65	3465.05±32.73	2928.51±75.89	8209.24±164.31	494.94±3.08	1921.17±338.06
	CoTTA	767.93±5.47	3799.88±17.29	2650.45±15.04	6242.52±33.14	933.41±4.31	2763.88±143.48
	ProxyTTA-fast	<i>732.61±29.57</i>	<i>3002.19±52.29</i>	2733.96±34.32	<i>6099.48±82.32</i>	875.01±15.8	2400.17±21.44
	ProxyTTA (Ours)	<b>686.91±22.14</b>	<b>2666.70±56.64</b>	<b>2589.25±59.03</b>	<b>6006.18±90.66</b>	<b>477.28±3.32</b>	<b>1598.64±18.95</b>
CostDCNet	Pretrained	1042.98	6301.60	3064.724	6630.649	1093.79	2798.25
	BN Adapt	1476.57±1.38	5428.20±8.15	2306.04±28.86	6391.98±48.97	596.08±5.55	1877.91±45.56
	BN Adapt, $\ell_z, \ell_{sm}$	<i>729.67±3.14</i>	3413.76±14.59	2288.85±14.02	6338.38±31.31	<i>469.97±2.47</i>	<b>1572.95±10.63</b>
	CoTTA	756.32±3.59	3686.69±14.75	2676.83±68.92	<i>6099.49±66.79</i>	689.94±1.95	2140.23±16.12
	ProxyTTA-fast	<i>756.98±31.07</i>	<i>3091.78±105.42</i>	2595.81±12.13	6373.01±7.74	606.10±11.10	1817.79±19.14
	ProxyTTA (Ours)	<b>512.72±0.74</b>	<b>2735.01±3.53</b>	<b>2062.28±11.24</b>	<b>5509.96±23.41</b>	<b>466.44±1.63</b>	<i>1580.38±11.48</i>
	VOID → NYUv2		VOID → SceneNet		VOID → ScanNet		
MSG-CHN	Pretrained	1040.934	1528.983	281.28	645.01	687.988	1201.747
	CoTTA	<i>876.93±146.95</i>	<i>1148.62±173.53</i>	<i>223.19±14.77</i>	<i>498.46±28.21</i>	<i>619.37±4.14</i>	<i>1141.04±7.35</i>
	ProxyTTA-fast (Ours)	<b>699.60±6.00</b>	<b>1120.37±9.76</b>	<b>192.74±1.72</b>	<b>424.49±4.58</b>	<b>302.21±4.10</b>	<b>480.08±8.03</b>
NLSPN	Pretrained	388.87	702.80	167.250	438.71	233.33	431.20
	BN Adapt	250.13±5.23	447.18±10.32	143.61±6.34	385.56±9.84	207.00±0.57	401.41±2.84
	BN Adapt, $\ell_z, \ell_{sm}$	<i>147.55±1.36</i>	<i>271.10±2.17</i>	<i>120.48±1.94</i>	<i>345.91±7.14</i>	<i>82.76±0.47</i>	<i>181.97±1.21</i>
	CoTTA	390.50±8.29	704.72±16.74	205.02±1.79	540.01±4.08	234.77±1.52	496.18±2.75
	ProxyTTA-fast	168.43±3.46	309.48±6.92	124.67±1.33	357.56±2.59	104.06±11.03	232.84±20.46
	ProxyTTA (Ours)	<b>124.41±2.27</b>	<b>240.73±5.72</b>	<b>113.93±1.49</b>	<b>333.41±4.32</b>	<b>74.77±0.31</b>	<b>166.61±0.45</b>
CostDCNet	Pretrained	189.10	446.71	173.37	443.22	144.31	458.69
	BN Adapt	160.31±2.7	410.55±10.70	176.62±0.72	446.32±8.52	159.65±4.63	399.14±13.92
	BN Adapt, $\ell_z, \ell_{sm}$	136.80±5.35	338.59±22.36	134.22±2.33	385.9±6.68	<i>68.44±0.46</i>	<i>164.59±2.82</i>
	CoTTA	147.69±5.3	376.87±21.25	136.42±3.41	405.38±11.63	101.98±1.53	322.63±5.04
	ProxyTTA-fast	<i>131.93±2.58</i>	<i>269.02±5.61</i>	<i>129.99±3.88</i>	<b>353.86±7.91</b>	128.12±3.41	244.62±7.53
	ProxyTTA (Ours)	<b>95.87±2.16</b>	<b>203.83±4.72</b>	<b>125.75±1.93</b>	<i>357.12±4.13</i>	<b>68.17±0.44</b>	<b>162.35±1.12</b>

Table 1. *Qualitative results.* For indoors, we adapt from VOID to NYUv2, SceneNet, and ScanNet; for outdoors, from KITTI to VKITTI with fog, nuScenes, and Waymo. **Bold** denotes best and *Italics* second-best. ProxyTTA-fast denotes our method without updating BatchNorm.

(BN) adaptation from TENT [42]. In BN adaptation, we only update the batch norm layer’s scale and shift factor based on the loss function. On average, BN Adapt with  $\ell_z, \ell_{sm}$  improves the pretrained model by 32.77%; whereas, updating just our adaptation layer (ProxyTTA-fast) improves it more by 34.07% (Table 1). The improvement of ProxyTTA-fast over BN adapt demonstrates the efficacy of updating adaptation layer, which directly adjusts the high-level features from the RGB branch guided by proxy loss, where BN adapt realigns the learned source features from both RGB and range sensors by updating feature statistics.

Nonetheless, the best results are achieved when we include batch norm update, which improves the pretrained model by 44.53%, but at the cost  $\approx 33.2\%$  of total extra time. The improvement of ProxyTTA over BN adapt implies that the large domain discrepancy may not be addressed by adapting only BN parameters (i.e., scaling and shifting); ProxyTTA explicitly adjusts RGB features by updating the adaptation layer with proxy embeddings as guidance.

We also compared our approach to CoTTA [43], which adapts *the whole model parameters* using the prediction from the teacher model updated by exponential moving average of

pretrained weight and the model prediction. We combined additional loss  $\ell_z, \ell_{sm}$  on top of CoTTA loss, since we observed that the models cannot be adapted with CoTTA alone. Specifically, our method without proxy shows a 25.26% average improvement on the CoTTA method. CoTTA updates the whole parameters including RGB and sparse depth branch, which causes a drift from the learned model parameters. On the other side, our method only updates additional layer at RGB branch, based on the study from the most domain discrepancy comes from RGB modality as studied in Sec. 3.1, and this prevents the model from a drift from learned domain. Also, CoTTA assumes the test-time augmentation can mitigate the domain shift. However, the results shows test-time augmentation on RGB image, causing a small distributional shift, may not solve a large domain discrepancy.

Also, our method with batch normalization layer update shows 26.52% average improvement, while using 25.05% less adaptation time. Note: CoTTA costs not only *additional memory for teacher model* but also *inference time to get the teacher model prediction*, even if CoTTA does not require any preparation process. Overall, our method shows 21.09% average improvement over BN adapt and CoTTA methods.

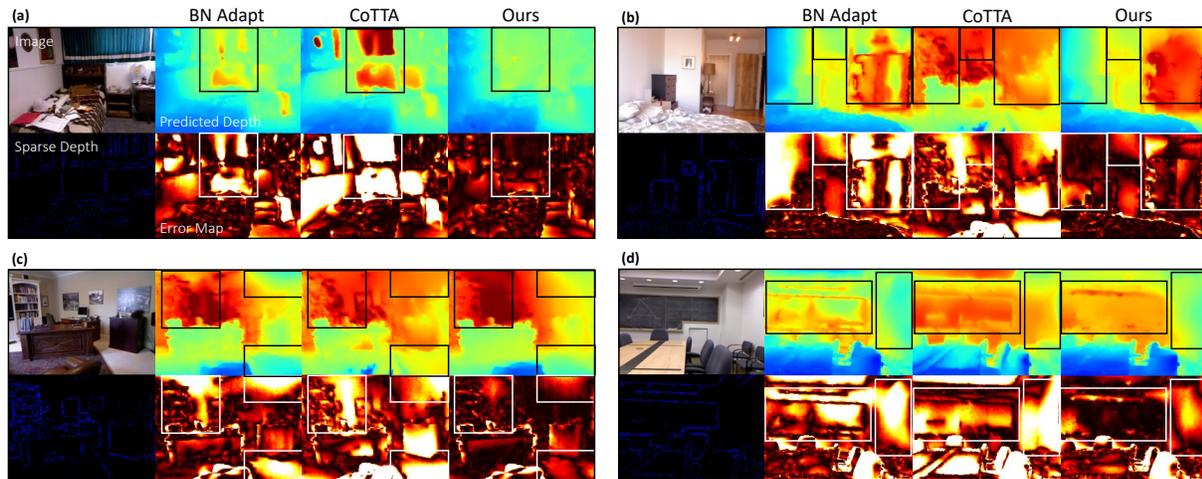


Figure 4. *Qualitative results on NYUv2*. For indoors scenarios, ProxyTTA performs better in boundary regions displaying the discontinuity in depth (e.g., curtains, (a)), as well as homogeneous regions (e.g., blackboard, (d)). Boxes highlight detailed comparisons.

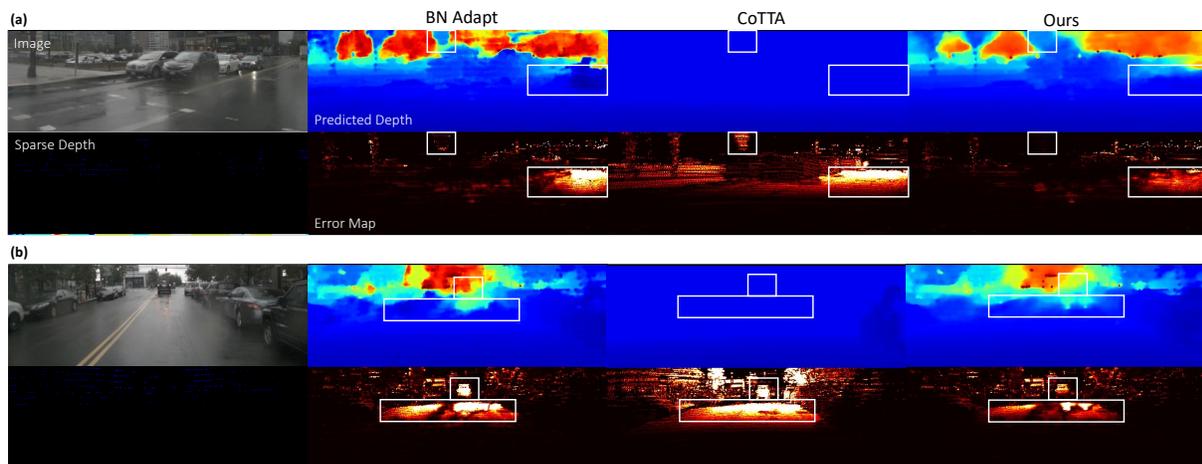


Figure 5. *Qualitative results on NuScenes*. For outdoor adaptation scenarios, ProxyTTA improves over BN Adapt and CoTTA, notably in both depth-discontinuous regions (e.g., car in (b)) and homogeneous regions (e.g., road in (a) and (b)). Boxes highlight detailed comparisons.

## 5. Discussion

We have proposed a method for test-time adaptation for depth completion that leverages the strength of complementary multi-sensor setup in the presence of domain shift. By studying model sensitivity to each input modality as well as the data under domain shift, we designed a way to exploit the modality (sparse depth) that is less sensitive to guide adaptation. We do so through a proxy embedding that learns the photometric information from the source domain that is compatible with the sparse depth depicting a 3D scene. Our proxy embedding works well as a regularizer for scenarios where there exists covariate shifts in photometry (i.e., KITTI to VKITTI) as well as scene layouts (i.e., VOID to NYUv2 and SceneNet). While one may surmise that the application of the embeddings are specific to scene distributions, we show otherwise. VOID (classrooms, laboratories, and gardens), NYUv2 (households and shopping centers), and SceneNet (randomly arranged synthetic rooms) all differ in layouts. The proxy embedding captures latent photometric features of the object shapes populating them; the same

proxy embedding can be transferred across domains even when scene differ, but share objects within them.

This leads to possible limitations in the scenarios where the source dataset is sampled from scenes that do not share any objects with the target test dataset; in this case, the proxy embeddings should give little to no gain and one must rely on generic regularizers like local smoothness. Additionally, while we follow the conventions in TTA and assume access to the source dataset prior to deployment, in reality, many models are trained on private datasets, so adapting “off-the-shelf” models remains a challenge. In such cases, one must incorporate our preparation pipeline into their model training and release the adaptation layer and proxy embedding module together with network weights. Nonetheless, this is the first test-time adaptation work in depth completion; in addition to our findings, we release models, dataset, adaptation, and evaluation code, and hope to further motivate interest in TTA for multi-modal tasks like depth completion.

### Acknowledgements.

This work was supported by NSF-2112562.

## References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 6
- [2] Dian Chen, Dequan Wang, Trevor Darrell, and Sayna Ebrahimi. Contrastive test-time adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 295–305, 2022. 2
- [3] Yun Chen, Bin Yang, Ming Liang, and Raquel Urtasun. Learning joint 2d-3d representations for depth completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10023–10032, 2019. 2
- [4] 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. 3
- [5] Sungha Choi, Seunghan Yang, Seokeon Choi, and Sungrack Yun. Improving test-time adaptation via shift-agnostic weight regularization and nearest source prototypes. In *ECCV*, pages 440–458. Springer, 2022. 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. 3
- [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. 6
- [8] Yiming Dou, Fengyu Yang, Yi Liu, Antonio Loquercio, and Andrew Owens. Tactile-augmented radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2
- [9] Abdelrahman Eldesokey, Michael Felsberg, and Fahad Shahbaz Khan. Propagating confidences through cnns for sparse data regression. *arXiv preprint arXiv:1805.11913*, 2018. 3
- [10] 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. 3
- [11] 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. 2, 6
- [12] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 2, 3
- [13] Jean-Bastien Grill, Florian Strub, Florent Alché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 5
- [14] 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. 3
- [15] 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
- [16] Yusuke Iwasawa and Yutaka Matsuo. Test-time classifier adjustment module for model-agnostic domain generalization. 34:2427–2440, 2021. 2
- [17] 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. 3
- [18] Jaewon Kam, Jungeon Kim, Soongjin Kim, Jaesik Park, and Seungyong Lee. Costdnet: Cost volume based depth completion for a single rgb-d image. In *European Conference on Computer Vision*, pages 257–274. Springer, 2022. 3, 6
- [19] 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
- [20] 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. 3, 6
- [21] 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. 3
- [22] 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? 34:21808–21820, 2021. 2, 3
- [23] 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. 3
- [24] Fangchang Ma, Guilherme Venturéli 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, 3
- [25] 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. 2, 6
- [26] Pushmeet Kohli, Nathan Silberman, Derek Hoiem, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012. 2, 6

- [27] 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. 3, 6
- [28] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019. 2, 3
- [29] 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. 3
- [30] 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. 3
- [31] 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. 3
- [32] 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. 3
- [33] 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
- [34] Shreyas S Shivakumar, Ty Nguyen, Ian D Miller, Steven W Chen, Vijay Kumar, and Camillo J Taylor. Dfusenet: 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
- [35] Junha Song, Jungsoo Lee, In So Kweon, and Sungha Choi. Ecotta: Memory-efficient continual test-time adaptation via self-distilled regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11920–11929, 2023. 2
- [36] Pei Sun, Henrik Kretschmar, 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. 6
- [37] Yu Sun, Eric Tzeng, Trevor Darrell, and Alexei A Efros. Unsupervised domain adaptation through self-supervision. *arXiv preprint arXiv:1909.11825*, 2019. 3
- [38] 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 *ICML*, pages 9229–9248. PMLR, 2020. 2, 3
- [39] 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, 6
- [40] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 6
- [41] 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. 3
- [42] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*, 2020. 2, 6, 7
- [43] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7201–7211, 2022. 2, 6, 7
- [44] 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
- [45] 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, 6
- [46] 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.
- [47] 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
- [48] Yangchao Wu, Tian Yu Liu, Hyoungseob Park, Stefano Soatto, Dong Lao, and Alex Wong. Augundo: Scaling up augmentations for unsupervised depth completion. *arXiv preprint arXiv:2310.09739*, 2023. 2
- [49] Yan Xu, Xinge Zhu, Jianping Shi, Guofeng Zhang, Hujun Bao, and Hongsheng Li. Depth completion from sparse lidar data with depth-normal constraints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2811–2820, 2019. 3
- [50] Fengyu Yang, Chenyang Ma, Jiacheng Zhang, Jing Zhu, Wenzhen Yuan, and Andrew Owens. Touch and go: Learning from human-collected vision and touch. *Neural Information Processing Systems (NeurIPS) - Datasets and Benchmarks Track*, 2022. 2
- [51] Fengyu Yang, Jiacheng Zhang, and Andrew Owens. Generating visual scenes from touch. *International Conference on Computer Vision (ICCV)*, 2023.
- [52] Fengyu Yang, Chao Feng, Ziyang Chen, Hyoungseob Park, Daniel Wang, Yiming Dou, Ziyao Zeng, Xien Chen, Rit Gan-gopadhyay, Andrew Owens, and Alex Wong. Binding touch to everything: Learning unified multimodal tactile representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2

- [53] 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. 3
- [54] 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. 3
- [55] 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. 3