

AugUndo: Scaling Up Augmentations for Monocular Depth Completion and Estimation

Yangchao Wu¹, Tian Yu Liu¹, Hyoungseob Park²,
Stefano Soatto¹, Dong Lao¹, and Alex Wong²

¹ UCLA Vision Lab, Los Angeles, CA 90095, USA
`{wuyangchao1997, tianyu, soatto, lao}@cs.ucla.edu`

² Yale Vision Lab, New Haven, CT 06511, USA
`{hyoungseob.park, alex.wong}@yale.edu`

Abstract. Unsupervised depth completion and estimation methods are trained by minimizing reconstruction error. Block artifacts from resampling, intensity saturation, and occlusions are amongst the many undesirable by-products of common data augmentation schemes that affect image reconstruction quality, and thus the training signal. Hence, typical augmentations on images viewed as essential to training pipelines in other vision tasks have seen limited use beyond small image intensity changes and flipping. The sparse depth modality in depth completion have seen even less use as intensity transformations alter the scale of the 3D scene, and geometric transformations may decimate the sparse points during resampling. We propose a method that unlocks a wide range of previously-infeasible geometric augmentations for unsupervised depth completion and estimation. This is achieved by reversing, or “undo”-ing, geometric transformations to the coordinates of the output depth, warping the depth map back to the original reference frame. This enables computing the reconstruction losses using the original images and sparse depth maps, eliminating the pitfalls of naive loss computation on the augmented inputs and allowing us to scale up augmentations to boost performance. We demonstrate our method on indoor (VOID) and outdoor (KITTI) datasets, where we consistently improve upon recent methods across both datasets as well as generalization to four other datasets. Code available at: <https://github.com/alexklwong/ augundo>.

Keywords: depth completion · monocular depth · augmentations

1 Introduction

Data augmentation is essential to training machine learning models; it plays a role in performance and generalization [52, 64, 71]. One common axiom of choosing augmentations is that the task output should remain invariant to the augmentation. For example, image flipping is a viable augmentation for classifying animals, since it does not alter the label. Conversely, flipping road signs can alter their meanings; hence, such augmentation can be detrimental to tasks involving road sign recognition. For geometric tasks, the range of augmentations

is more restricted due to constraints of problem formulation: Stereo assumes pairs of frontoparallel rectified images; hence, in-plane rotations are not viable. Image-guided sparse depth completion relies on sparse points to ground estimates to metric scale; therefore, intensity transformations on sparse depth maps that alter the scale of the 3-dimensional (3D) scene are infeasible. Unsupervised learning of depth completion and estimation further limit the use of augmentations as the supervision signal comes from reconstructing the inputs, where augmenting the input introduces artifacts that impact reconstruction quality and therefore the supervision (see Sec. D, Fig. 1, Tab. 8, 9 in the Supp. Mat. for examples of artifacts and extended discussion on their effect on learning). Moreover, video-based unsupervised training assumes rigid motion, so augmentations that introduce padding (e.g., translation, rotation) will yield constant or edge extended borders across images (i.e., no motion), preventing the model from properly learning depth and pose – leaving few augmentations viable. While simulating nuisances is desirable, naively applying augmentations may do more harm than good; thus, it is unsurprising that existing work in unsupervised depth completion [45, 48, 63, 82–84, 87, 94] and estimation [18, 19, 47, 102, 108] primarily rely on a small range of photometric augmentations and flipping.

Nevertheless, photometric augmentations help model the diverse range of illumination conditions and colors of objects that may populate the scene; geometric augmentations can simulate the various camera parameters, i.e., image resizing (zooming) can model changes in focal length, and scene arrangements, i.e., image flipping. However, block artifacts, loss during resampling, and intensity saturation are just some of the many undesirable side-effects of traditional augmentations to the image and sparse depth map for unsupervised learning of geometric tasks. To avoid compromising the supervision signal, we compute the typical reconstruction loss on the original input image and sparse depth map instead of the augmented inputs, which bypasses negative effects of reconstruction artifacts due to photometric and geometric augmentations. However, there exists a mis-alignment between the original input (e.g., image, sparse depth), and the model depth estimate as geometric augmentations induce a change in coordinates. Hence, we *undo* the geometric augmentations by inverting them in the output space to align the model estimate with the training target.

Amongst the many geometric tasks, we focus on *unsupervised depth completion*, the task of inferring a dense depth map from an image and sparse depth map, where augmentations have seen limited use. Here, a training sample includes the input sparse depth map, its associated image and additional images of neighboring views of the same 3D scene. Our method is also applicable to *unsupervised monocular depth estimation*, which omits the sparse depth modality. Augmentations have traditionally been restricted to a limited range of photometric transformations and flipping – due to the need to preserve photometric consistency across a sequence of video frames used during training, and the sparse set of 3D points projected onto the image frame as a 2.5D range map; degradation to either modalities directly impacts the supervision signal as the loss function is conventionally computed on the augmented inputs. By using our

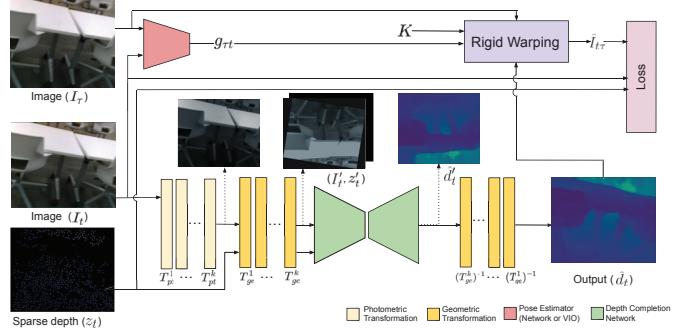


Fig. 1: Overview. We apply photometric augmentations to the input image, and the same set of geometric augmentations to both the input image and sparse depth map. We warp the output depth back to the original reference frame with the inverse geometric transformations. This enables image and sparse depth reconstruction losses to be computed on the original inputs while unlocking previously-infeasible augmentations.

method, loss functions involving sparse depth and image reconstruction from other views can be computed on the original inputs while applying augmentations that were previously not viable for the task. **Our hypothesis:** By “undoing” the augmentations, one can expand the viable set and scale up their use in training, leading to improved model performance and generalizability.

To this end, we introduce AugUndo, an augmentation framework that enables one to apply a wide range of photometric and geometric transformations on the inputs, and to “undo” them during loss computation. This allows computing the unsupervised loss on the original images and sparse depth maps, free of artifacts, through a warping of the output depth map – obtained from augmented input – onto the input frame of reference based on the inverse geometric transformation. In addition to group transformations that allow for output alignment, we combine them with commonly employed photometric augmentations. To the best of our knowledge, we are the first to propose a unified augmentation scheme for photometric and geometric augmentations for unsupervised depth completion and estimation. We demonstrate AugUndo on recent methods on indoor and outdoor settings, where we consistently improve all methods across all datasets.

Our contributions are as follows: (1) We propose AugUndo, a simple-yet-effective framework to scale up photometric and geometric augmentations for unsupervised depth completion and estimation, without compromising the supervision signal; AugUndo can be applied in a plug-and-play manner to existing methods with negligible increase in computational costs during training. (2) We enable previously-infeasible augmentations to be used for training unsupervised methods and comprehensively ablate combinations of eleven types of augmentations to study the performance benefits of each. (3) We show that AugUndo can consistently improve robustness to shifts in sparse point densities for completion, model performance as well as zero-shot generalization for both depth completion and estimation for indoor and outdoor scenarios; thus, validating our hypothesis.

2 Related Work

Data augmentation for depth completion and estimation. While photometric transformations such as color jitter are often applied for unsupervised depth completion [48, 82–84], and estimation [13, 18–20, 37, 38, 53, 59, 86, 100, 101], geometric augmentations other than simple flipping are less commonly adopted. Most supervised depth completion methods [11, 24, 25, 27, 30, 40, 42, 55–57, 90, 94, 105, 106] similarly limit their augmentations to color jitter and horizontal flips due to sparse depth maps being decimated by rotation and scaling, causing points to be interpolated away. Nevertheless, for supervised training, it is straightforward to directly apply the same transformation to the ground truth for training. Indeed, some supervised depth completion methods [5, 6, 12, 50, 51, 66, 76, 99] adopted random scaling, translation, in-plane rotation. Translation augmentations are also commonly applied in supervised depth estimation methods [9, 10, 32, 41, 44, 58, 59, 75, 81, 88, 92, 97]. We posit that the reason that such transformations can be applied to the ground truth in supervised settings is due to artifacts caused by the transformation of a piece-wise smooth depth map being less severe than those of an RGB image and its intensities. However, such assumptions do not hold in unsupervised training. These artifacts would affect the training signal for unsupervised methods, which rely on photometric correspondences and observed sparse points. Our approach aims to resolve this to allow diverse geometric augmentations to be applied in a plug-and-play fashion in unsupervised training.

Unsupervised depth completion methods [48, 63, 82–84, 91, 94] leverage image and sparse depth reconstructions as training signals by minimizing errors between the input image and its reconstruction from other views, and errors between the input sparse depth map and the predicted depth along with a local smoothness regularizer. [48] used Perspective-n-Point [39] and RANSAC [14] to align consecutive video frames. [94] learns a depth prior conditioned on the image. [82] uses synthetic data to learn a prior on the shapes populating a scene, while [46] translates synthetic data to real domain to leverage rendered depth. [83] proposed an adaptive scheme to reduce penalties incurred on occluded regions. [87] maps the image onto the 3D scene through calibrated back-projection. [91] decouples structure and scale. [26] uses line features from visual SLAM and [45] introduced monitored distillation for positive congruent training. Augmentations for these methods are limited to a small range of photometric perturbations and image flipping. Operations such as rotation, resizing, and translation require resampling, which creates artifacts and affects the reconstruction quality. This causes performance degradation since loss is typically computed on the augmented images. Loss in sparse depth maps is further impacted as resampling and interpolation may cause loss of sparse points. Contrary to these limited augmentation schemes, we enable a large range of photometric and geometric augmentations to be introduced during training.

Unsupervised monocular depth estimation, like depth completion, also minimizes photometric reconstruction error. [15] frames depth estimation as a novel view synthesis problem. [18] improves [15] by imposing a consistency loss on the depth predicted from left and right images. [108] uses a pose network to

enable unsupervised training on video sequences. [19] introduced auto-masking and min reprojection loss. To improve the supervision signal, [7, 73, 79] leverage noisy proxy labels and [54] uses a binocular assumption. Additional loss terms based on visual odometry [77], iterative closest point [49], surface normals [96], and semantic segmentation [21, 31] were also introduced; [53, 93] further included predictive uncertainty. To handle rigid and non-rigid motions in the scene, previous works explored multi-task learning to include optical flow and moving object estimation [4, 60, 95, 98, 109], and used semantics to filter out outlier regions [29]. [47] redesigned the skip connection and decoders to extract high-resolution features, [107] combined global and local representations and [102] introduced a lightweight architecture with dilated convolution and attention. In addition to depth completion, our method also shows consistent improvements on monocular depth estimation for [19, 47, 102].

3 Method Formulation

Let $I : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}_+^3$ be an RGB image captured by a calibrated camera, $z : \Omega_z \subset \Omega \rightarrow \mathbb{R}_+$ the corresponding sparse point cloud projected onto the image plane as a sparse depth map, and $K \in \mathbb{R}^{3 \times 3}$ the intrinsic calibration matrix. Given an image and its sparse depth map, depth completion aims to learn a function $f_\theta(I, z)$ that recovers the distance between the camera to points in the 3D scene as a dense depth map. In another mode, if sparse depth maps are not given, then the problem reduces to monocular depth estimation which learns a function $f_\theta(I)$ to map a single image to a depth map $\Omega \rightarrow \mathbb{R}_+$. For the ease of notation, we denote the output depth map for both depth completion and estimation as $\hat{d} \in \mathbb{R}_+^{H \times W}$ where H and W are its height and width.

Unsupervised depth completion relies on photometric and sparse depth reconstruction errors as its primary supervision signal. Without loss of generality, we assume an input of (I_t, z_t) for an RGB image and associated sparse depth map captured at time t and during training, an additional set of temporally adjacent images I_τ for $\tau \in \mathcal{T} \doteq \{t-1, t+1\}$. The reconstruction $\hat{I}_{t\tau}$ of I_t from image I_τ is given by the reprojection based on estimated depth $\hat{d}_t := f_\theta(\cdot)$

$$\hat{I}_{t\tau}(x) = I_\tau(\pi g_{\tau t} K^{-1} \bar{x} \hat{d}_t(x)) \quad (1)$$

where $\bar{x} = [x^\top, 1]^\top$ is the homogeneous coordinates of $x \in \Omega$, $g_{\tau t} \in SE(3)$ the relative pose of the camera from time t to τ , K the intrinsic calibration matrix, and π the canonical perspective projection.

Using Eq. (1), a depth completion or estimation network f_θ minimizes

$$\arg \min_{\theta} \sum_{\tau \in \mathcal{T}} \sum_{x \in \Omega} \alpha \rho(\hat{I}_{t\tau}(x), I_t(x)) + \sum_{x \in \Omega_z} \beta \psi(\hat{d}_t(x), z_t(x)) + \lambda R(\hat{d}_t) \quad (2)$$

where ρ denotes the photometric reconstruction error, typically L_1 difference in pixel values and/or structural similarity (SSIM), ψ the sparse depth reconstruction error, typically L_1 or L_2 , R the regularizer that biases the depth map to

be piece-wise smooth with depth discontinuities aligned with edges in the image (commonly used by previous works [19,48,84,87]), and α , β and λ their respective weighting. Note: monocular depth omits sparse depth term, i.e., $\beta = 0$.

Inverse transformation. Let \mathcal{A}_{pt} be the set of possible photometric transformations, and \mathcal{A}_{ge} be the set of all geometric transformations. Given $T_{ge} \in \mathcal{A}_{ge}$, we wish to obtain an inverse transformation T_{ge}^{-1} such that $T_{ge} \circ T_{ge}^{-1} \approx Id$ the identity function³. At each time step, we can sample a sequence of transformations $\{T_{pt}^1 \dots T_{pt}^k\}$ and $\{T_{ge}^1 \dots T_{ge}^m\}$ respectively from \mathcal{A}_{pt} and \mathcal{A}_{ge} to construct transformations $T_{pt} = T_{pt}^1 \circ \dots \circ T_{pt}^k$ and $T_{ge} = T_{ge}^1 \circ \dots \circ T_{ge}^m$. We denote the composition of a collection of augmentation transformations by $T = T_{ge} \circ T_{pt}$ where T_{pt} denotes photometric transformation, and T_{ge} denotes geometric transformation. Furthermore, we denote the inverse geometric transformations by $T_{ge}^{-1} = (T_{ge}^m)^{-1} \circ (T_{ge}^{m-1})^{-1} \circ \dots \circ (T_{ge}^1)^{-1}$, which operates on the space of depth maps to reverse the geometric transformation so that we can warp the output depth map onto the reference frame of the original image. Through inverse warping, this process is differentiable. In practice, some transformations cause border regions of the image to be warped out of frame, i.e., translated off the image plane or cropped away. Hence, after warping our output depth map back to the original frame of reference using the inverse geometric transformations, border extensions (edge paddings) are introduced to handle out-of-frame regions.

AugUndo. We apply each geometric transformation over image coordinates and resample (bilinear for image and nearest neighbor for sparse depth):

$$[x' \ 1]^\top = T_{ge} [x \ 1]^\top \quad (3)$$

$$I'(x') = T_{pt}(I)(x); \ z'(x') = z(x) \quad (4)$$

where T_{ge} is the geometric transformation, $x \in \Omega$ and $x' \in \Omega$ are coordinates in the image grid, and I' is the image after the transformation; for ease of notation, we hereafter denote $I' = T(I) = T_{ge} \circ T_{pt}(I)$ to include both augmentations through composition. Note that x is in the original image reference frame and x' is in the transformed image reference frame. Naturally, this process can be extended to multiple geometric transformations by composing them, i.e., $T_{ge} = T_{ge}^1 \circ T_{ge}^2 \circ \dots \circ T_{ge}^m$. The reverse process is simply inverting the transformations where $T_{ge}^{-1} = (T_{ge}^m)^{-1} \circ (T_{ge}^{m-1})^{-1} \circ \dots \circ (T_{ge}^1)^{-1}$:

$$[x'' \ 1]^\top = T_{ge}^{-1} [x' \ 1]^\top \quad (5)$$

$$\hat{d}(x'') = \hat{d}'(x') \quad (6)$$

where \hat{d}' is the depth map inferred from augmented inputs (I', z') . Once reverted back to the original reference frame, \hat{d} is used to reconstruct I_t from I_τ for $\tau \in \Upsilon$ as in Eq. (1). More details can be found in Alg. 1 in the Supp. Mat.

By modeling T_{ge}^{-1} , Eq. (3)-Eq. (6) allow us to apply a wide range of augmentations, while still establishing the correspondence between I_t and I_τ . Specifically,

³ Note that in practice, not all transformations are bijective. For instance, image rescaling with a fixed size image plane; hence strict equality is not always possible.

for minimizing Eq. (2), one may simply augment the image and sparse depth by T to yield (I'_t, z'_t) as input, and reconstruct the original image and sparse depth (I_t, z_t) from other views I_τ and the aligned output depth \hat{d} (Eq. (6)). We note that the inverse transformation is critical for enabling the sparse depth reconstruction term to be computed properly in Eq. (2); if computed in the transformed reference frame, i.e., on z'_t , many of the sparse points would be decimated by interpolation during rotation and resizing (downsampling) augmentations – leaving a lack of supervision from the sparse depth term of the loss function. We note that the original image sequence is fed to the pose estimator, rather than the augmented images. This ensures the estimated relative pose used during reprojection (Eq. (1)) also corresponds to the original images.

Modeling geometric augmentations. Depth completion differs from typical image-based (e.g., monocular) depth estimation, where operations like resizing are susceptible to scale ambiguity and can either correspond to changes in focal length, or distance from the camera (i.e., due to changes in camera pose). In depth completion, the sparse depth map grounds image pixels to specific depth values. Hence, there is no ambiguity in distance of each sparse point from the camera. Our goal is to synthesize new training data via geometric transformations that are consistent with the original 3D scene. Given that there is no scale ambiguity, this can be achieved by either (1) changes in camera pose (extrinsics), or (2) changes in camera parameters (intrinsic).

Augmentation via (1) would require warping the inputs for view synthesis. Since we operate under the unsupervised setting, accurate warping (preserving 3D scene) without access to dense ground truth depth values is difficult to achieve. An exception to this is rotation, which we address below. On the other hand, (2) can be achieved via standard 2D image transformations while preserving the 3D structure of the scene. This motivates our modeling of geometric transformations, i.e., resizing, translation, as changes in camera parameters.

For example, augmenting focal length (i.e. zooming in/out) is akin to “resizing”. Similarly, translation can be used to model shifting of the optic center, i.e. capturing the same 3D scene using cameras with different principal point offsets. Note: image rotation can be seen as rotating a camera about its optical axis, which does not change the distance of a point from the camera. Thus, the camera position and the scene are kept constant across augmentations, eliminating the need for adjusting depth values via warping during the process. Instead, we only have to realign the output depth back to the original frame of reference (Eqs. (5) and (6)), which makes our method computationally efficient.

Augmentations. Our choices are based on common nuisance variabilities, i.e., changes in illumination, occlusions, and object color, scene layout (flip), and camera parameters (resize, translation) and orientation (rotation).

Photometric. We include brightness, contrast, saturation, and hue, where all follow standard augmentation pipeline in existing works [19, 84, 87, 102]. Applying the inverse of photometric augmentation can be viewed as recovering the original image; hence, we directly use original image instead of applying transformations to the intensities (which are not recoverable if saturated at the max value).

Occlusion. We consider patch removal and sparse point removal. For the former, we randomly select a percentage of pixels $x \in \Omega$ in the image and remove (with zero-fill) an arbitrary-sized patch around it. For the latter, we randomly sample a percentage of points $x \in \Omega_z$ in the sparse depth map. The inverse transformation of this is simply the original image and sparse depth map, both used in loss computation before augmentation.

Flip. We consider horizontal and vertical flips. When applied, the same flip operation is used for both input image and sparse depth maps. We record the flip type during data augmentation. During loss computation, we reverse the flip direction to align output depth with the original image and sparse depth map.

Resize. We define a new image plane of the input resolution and generate a random scaling factor to be applied along both height and width directions. The image is warped to the new image plane, where any point warped out of the plane is excluded; borders of the warped image are extended to the bounds of the image plane by edge replication. Sparse points occupying multiple pixels are eroded to a single point for resizing with a factor greater than 1. During loss computation, we warp the output depth map onto a new image plane of the same dimensions by the inverse scaling matrix of the recorded scaling factors; borders of the warped depth map are extended by edge replication if necessary. We view resizing factors greater and less than 1 (zooming in and out) as two separate forms of augmentations to distinguish their contributions.

Rotation. Naive rotation leads to the loss of large areas of the image, i.e., cropped away to retain the same-size image, discarding large portions of possible co-visible regions and also supervisory signals. To preserve the entire image, we first warp the image by a randomly generated angle to a new (larger) image plane, so that the rotated image fits tightly within the new image. As image sizes within a batch can vary depending on the rotation angle, we center-pad (uniformly on each side) each image in the batch to the maximum width and height of the augmented batch. To reverse the rotation on the output depth, we warp the output depth map back by the inverse rotation matrix, then perform a center crop on the depth map to align with the original input.

Translation. We define a new image plane of the same dimensions as the input and generate random height and width translation factors. The coordinates of the input are translated and its pixels are inverse warped onto the new image plane. Any pixel warped out of the image plane is excluded. Borders of the warped image are extended to the bounds of the image plane by edge replication. During loss computation, we warp the output depth map onto a new image plane of the same dimensions by the inverse translation matrix. Borders of the warped depth map are extended to the bounds of the image plane by edge replication.

4 Experiments

We demonstrate AugUndo on recent unsupervised depth completion (VOICED [84], FusionNet [82], KBNet [87]) and estimation (Monodepth2 [19], HRDepth [47], LiteMono [102]) methods on two datasets (KITTI [17, 74], VOID [84]) here.

Due to space limitations, we defer the results of MonDi [45] and DesNet [91] to Sec. K in the Supp. Mat. To evaluate zero-shot generalization, we test on three additional datasets for each task: NYUv2 [65], ScanNet [8], Waymo [69], and Make3d [62] – where we transfer models on VOID to NYUv2 (Tab. 3) and ScanNet, and from KITTI to Waymo (Tab. 4 in Supp. Mat.) for completion and Make3d [62] (Tab. 5 in Supp. Mat.) for estimation. We also test model sensitivity to different sparse depth input densities in Tab. 1 (right) – see Sec. H of the Supp. Mat for an extensive study. Additionally, we present an comprehensive ablation study in Sec. G in the Supp. Mat. to test the effect of each augmentation. We also provide results on modeling AugUndo as changes in camera motion or parameters in Tab. 7 of the Supp. Mat.; interpreting geometric augmentations as different camera parameters yields better performance. Tabs. 8, 9 in Supp. Mat. show that unsupervised models cannot be trained with naive geometric augmentations. See Sec. E of the Supp. Mat. for descriptions of datasets.

Experiment setup. We followed the original settings of the open-sourced code for each work and modified the data handling and loss function to incorporate AugUndo. We perform 4 independent trials for each experiment and report their means and standard deviations. To ensure a fair comparison, we train all models from scratch. Below, we report the best combination of augmentations found empirically through extensive experiments on each dataset. See Secs. A and B in Supp. Mat. for implementation details, and evaluation metrics, respectively.

Augmentations. Through a search over augmentation types and values, we found a consistent set that tends to yield improvements across all methods with small changes to degree of augmentation catered to each method. See Sec. C in Supp. Mat. for details of augmentation parameters. Rows in Tabs. 1 to 6 marked with "+ AugUndo" denote models trained with our method. We note that performance gain can be obtained by typical set and ranges of augmentations and does not require a meticulous selection of hyper-parameters (see Sec. J of Supp. Mat. for a sensitivity study). Results reported here aims to study how far one can push performance and generalization, purely from augmentations.

Results on VOID. We begin by presenting quantitative results for AugUndo on unsupervised depth completion. Tab. 1 (left, VOID1500) shows our main results on the VOID benchmark. By training the models with AugUndo, we observe an average overall improvement of 18.3% across all methods and metrics on VOID1500. Specifically, we improve VOICED by 26.4%, FusionNet by 16.3%, and KBNet by 12.2%. These experiments validate our hypothesis that by applying a wider range of augmentations, we are able to improve the baseline model's performance. They also illustrate the lack in use of augmentations in existing works: incorporating standard augmentations (albeit requiring modification to augmentation and loss computation pipelines) can yield a large performance boost. Fig. 2 shows a head-to-head comparison between KBNet trained using standard procedure in [87] and KBNet trained using AugUndo. We observe qualitative improvements from AugUndo where we improve KBNet in homogeneous regions and image borders, i.e., pillar (left), cabinet (middle), wall (right). Applying geometric augmentations yields models with fewer border

Table 1: *Depth completion on VOID*. AugUndo improves performance by an average of 18.3% across all methods and metrics on VOID1500. When models trained on VOID1500 are tested VOID500, AugUndo improves by 23.1%, as translation, resizing, and occlusion augmentations vary sparsity by removing sparse points from the input.

Method	VOID1500				VOID500			
	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓
VOICED	74.78±2.69	139.75±4.57	39.20±1.46	71.98±2.54	137.01±4.23	235.80±7.82	71.36±1.86	130.63±5.66
+ AugUndo	52.73±0.41	111.09±0.92	26.93±0.54	54.46±0.38	92.99±1.11	176.94±1.38	46.43±0.85	91.10±1.64
FusionNet	52.11±0.44	113.30±1.18	28.53±0.52	58.79±2.01	97.73±0.73	194.32±1.36	58.65±1.31	122.95±3.04
+ AugUndo	41.16±0.18	99.21±0.39	22.23±0.35	53.07±1.30	74.97±0.69	162.71±2.39	40.44±1.39	92.11±5.79
KBNet	38.11±0.77	95.22±1.72	19.51±0.14	46.70±0.48	78.44±1.39	178.17±3.27	37.56±0.61	83.43±1.89
+ AugUndo	33.32±0.18	85.67±0.39	16.61±0.29	41.24±0.60	66.97±0.81	151.55±2.03	31.63±0.53	71.90±0.82

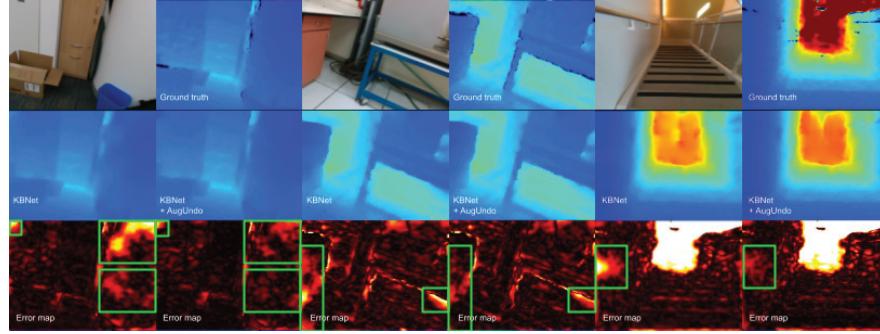


Fig. 2: *Depth completion on VOID*. We compare KBNet trained with standard augmentations and with AugUndo. AugUndo consistently produces lower errors with reduced border artifacts and improves on homogeneous regions, i.e., pillar (left), cabinet (middle), wall of staircase (right). Error maps are highlighted for comparison.

artifacts as we apply random translation, which warps part of the image out of frame. This allows us to simulate border occlusion, where there lacks correspondence in an adjacent frame. While training with standard protocol results in failures to recover structures near the image border, training with AugUndo can render models robust to them by computing the loss on the original frame of reference (where we do have correspondence) through our inverse transformations of the depth map. Additionally, resizing allows the model to learn multiple resolutions of the input, akin to zooming in and out, which can also impose smoothness in homogeneous regions through the scale-space transition; whereas, rotation can simulate diverse camera orientations. Fig. 2 shows that translation, rotation, and resizing can model these effect in the input space to yield models robust to these nuisance variability.

AugUndo also is applicable for monocular depth estimation. Tab. 2 shows a comparison using the standard augmentation procedure of Monodepth2, HR-Depth, and Lite-Mono and using AugUndo. Tab. 2 shows that AugUndo consistently improves all models across all error and accuracy metrics. Thus, validating

Table 2: *Monocular depth estimation on VOID.* AugUndo is applicable for monocular depth estimation and consistently improves three monocular depth models.

Method	MAE ↓	RMSE ↓	AbsRel ↓	SqRel ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
Monodepth2	283.861±3.732	395.947±5.728	0.183±0.002	0.100±0.003	0.717±0.005	0.922±0.004	0.975±0.002
+ AugUndo	277.696±4.861	388.088±5.768	0.178±0.003	0.095±0.004	0.724±0.007	0.925±0.004	0.978±0.002
HR-Depth	286.282±7.059	399.112±9.184	0.185±0.004	0.100±0.004	0.714±0.012	0.919±0.006	0.975±0.002
+ AugUndo	283.086±6.787	394.261±9.133	0.181±0.005	0.097±0.005	0.718±0.013	0.922±0.004	0.977±0.002
Lite-Mono	319.910±15.00	446.005±22.97	0.209±0.013	0.129±0.019	0.669±0.016	0.892±0.011	0.963±0.006
+ AugUndo	308.010±0.859	426.626±0.484	0.200±0.003	0.114±0.001	0.674±0.005	0.901±0.002	0.969±0.002

the hypothesis that AugUndo can be applied generically to improve monocular depth estimation. Specifically, we observe a boost in the most difficult accuracy metric ($\delta < 1.25$), where Monodepth2 improves from 0.717 to 0.724, HR-Depth from 0.714 to 0.718 and Lite-Mono from 0.669 to 0.674. At the same time, for AbsRel metric, Monodepth2 improves from 0.183 to 0.178, HR-depth improves from 0.185 to 0.181, and Lite-Mono improves from 0.209 to 0.200.

Sensitivity study. One limitation of existing depth completion training pipelines is that there are little to no augmentations applied to sparse depth modality. However, in real-world applications, sparse depth has, in fact, high variability, i.e., features tracked in SLAM/VIO systems will vary depending on the scene (points are dropped or added to the state), and point cloud densities returned by a sensor will differ based on specifications. To further examine the effect of AugUndo on sparse depth, we study the sensitivity to changes in sparse depth density by testing models on VOID1500 (\approx 1500 points) on VOID500 (\approx 500). For the $3\times$ reduction in sparse points, AugUndo improves robustness by an average of 23.1% across all methods. These improvements result from the geometric and occlusion augmentations made possible via AugUndo, which greatly increases sparse depth variations, i.e., decimating them through resizing, re-orienting their configuration through rotation, translating them out of frame, and randomly occluding them, to avoid overfitting particular sparse depth configurations. We further push their limits in Sec. H of the Supp. Mat., where we test them on VOID150 with a $10\times$ reduction in sparse points from the training set (VOID1500) and observe the same trend of improvements. This shows that AugUndo significantly improves robustness of models to changes in sparse depth.

Zero-shot generalization. We test depth completion models trained on VOID on NYUv2 and ScanNet. Tab. 3 shows an average of 23.2% improvement on NYUv2 and 27.6% on ScanNet. Applying AugUndo to VOICED greatly improve its generalizability to both NYUv2 and ScanNet. This is likely due to the scaffolding densification employed by VOICED, where the network can overfit to scaffoldings of particular sparse depth configurations and therefore, does not generalize well when presented with sparse depth maps with different configurations. Like our sensitivity study (Fig. 2, VOID500), occlusion and geometric augmentations introduce variation into the sparse depth configurations and densities, which alleviates overfitting to specific point clouds or 3D scenes; hence,

Table 3: Zero-shot transfer from VOID to NYUv2 and ScanNet for depth completion. AugUndo improves generalization of models trained on VOID to novel datasets by an average of 25.4% for all evaluation metrics (in millimeters) across both datasets.

Method	NYUv2				ScanNet			
	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓
VOICED	2240±143.90	2427±143.49	211±9.60	238±10.89	1562±136.79	1764±146.39	270±17.25	311±17.36
+ AugUndo	990±82.48	1181±65.55	110±6.11	132±6.01	638±44.74	791±79.48	131±7.68	170±6.32
FusionNet	132.24±2.12	236.16±4.59	28.68±0.42	61.87±1.20	109.47±3.01	206.33±6.11	55.45±1.56	122.52±2.04
+ AugUndo	124.93±3.05	227.23±4.96	25.70±0.41	54.09±0.87	100.64±2.31	195.85±5.85	45.98±0.78	99.90±5.00
KBNet	138.31±5.74	257.99±10.36	25.48±0.63	51.77±0.99	103.05±4.99	217.12±13.35	36.23±1.12	76.55±2.90
+ AugUndo	118.60±3.44	231.13±8.85	22.06±0.31	47.07±0.70	82.53±4.33	175.30±11.13	29.87±1.06	63.78±1.30

Table 4: Zero-shot transfer from VOID to NYUv2 and ScanNet for depth estimation. AugUndo consistently improves generalization (in meters) for monocular depth models.

Dataset	Method	MAE ↓	RMSE ↓	AbsRel ↓	SqRel ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
NYUv2	Monodepth2	0.432±0.003	0.556±0.005	0.205±0.001	0.159±0.001	0.683±0.005	0.907±0.001	0.975±0.001
	+ AugUndo	0.415±0.005	0.537±0.006	0.196±0.002	0.148±0.003	0.700±0.008	0.915±0.002	0.977±0.001
	HR-Depth	0.424±0.003	0.549±0.003	0.201±0.002	0.154±0.002	0.692±0.003	0.910±0.002	0.976±0.001
	+ AugUndo	0.421±0.009	0.542±0.011	0.199±0.005	0.152±0.006	0.696±0.009	0.913±0.004	0.976±0.001
	Lite-Mono	0.480±0.012	0.616±0.018	0.231±0.008	0.199±0.015	0.637±0.009	0.879±0.009	0.964±0.004
	+ AugUndo	0.468±0.003	0.595±0.003	0.225±0.004	0.187±0.005	0.646±0.002	0.889±0.003	0.968±0.001
	Monodepth2	0.284±0.004	0.368±0.005	0.177±0.002	0.097±0.002	0.741±0.006	0.931±0.003	0.980±0.001
	+ AugUndo	0.270±0.003	0.351±0.004	0.169±0.001	0.088±0.002	0.759±0.004	0.937±0.001	0.982±0.001
	HR-Depth	0.282±0.004	0.366±0.005	0.175±0.002	0.095±0.002	0.743±0.005	0.929±0.003	0.979±0.002
	+ AugUndo	0.274±0.007	0.357±0.009	0.172±0.005	0.092±0.005	0.754±0.009	0.935±0.003	0.981±0.001
	Lite-Mono	0.296±0.002	0.388±0.007	0.185±0.003	0.109±0.006	0.731±0.002	0.921±0.002	0.976±0.001
	+ AugUndo	0.296±0.001	0.382±0.001	0.185±0.001	0.105±0.002	0.728±0.002	0.924±0.001	0.977±0.001

improving VOICED from 49.9% and 52.6% on NYUv2 and ScanNet, respectively. For FusionNet and KBNet, we still see nontrivial improvement: FusionNet improves by 8% and 12.1% on NYUv2 and ScanNet, and KBNet by 11.7% and 18.3%. This further validates our hypothesis that by applying a more diverse set of transformations, we are able to improve generalization to new datasets.

For depth estimation, we tested Monodepth2, HR-Depth, Lite-Mono (trained on VOID) for zero-shot generalizability to NYUv2 and ScanNet. The results are shown in Tab. 4, where training with AugUndo consistently improves generalization errors across all evaluation metrics for both NYUv2 and ScanNet. Notably, for RMSE metric, Monodepth2 improves from 0.556 to 0.537, HR-Depth from 0.549 to 0.542, Lite-Mono from 0.616 to 0.595. For ScanNet, similar improvement on RMSE can also be observed, where Monodepth2 improves from 0.368 to 0.351, HR-Depth from 0.366 to 0.357, and Lite-Mono from 0.388 to 0.382. The improvement in RMSE metric, which is sensitive to outliers, highlights AugUndo’s ability to model different input data distributions, collected by a different camera of different orientation with different object colors and layouts.

Results on KITTI. We begin with quantitative results for depth completion. While AugUndo consistently improves all methods (Tab. 5), we note that the improvement is less, but respectable, in this case: ≈5.2% overall with the largest gain in FusionNet of 6.41%. This is largely due to the small scene varia-

Table 5: *Depth completion on KITTI.* AugUndo consistently improves all methods for all metrics (in millimeters) by approximately 5.2% on average.

Method	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓
VOICED	318.59±7.74	1,213.60±17.49	1.30±0.05	3.72±0.04
+ AugUndo	295.41±0.30	1,159.27±5.44	1.20±0.02	3.49±0.03
FusionNet	285.55±2.16	1,174.47±10.67	1.20±0.03	3.45±0.08
+ AugUndo	267.69±1.85	1,157.07±4.61	1.08±0.02	3.19±0.03
KBNet	263.90±3.63	1,130.66±6.22	1.05±0.01	3.24±0.04
+ AugUndo	256.37±1.00	1,114.53±3.79	1.01±0.01	3.13±0.03

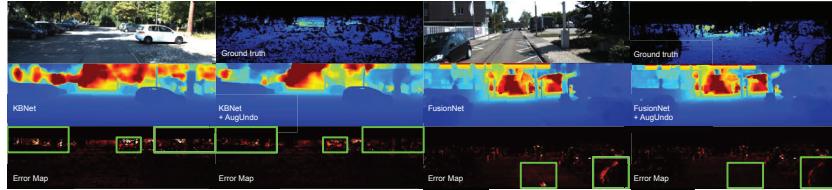


Fig. 3: *Depth completion on KITTI.* We compare KBNet and FusionNet trained with standard augmentations and with AugUndo. AugUndo consistently produces lower errors in highlighted regions where structures may have arbitrary orientation (vegetation) and regions near image border that typically lacks correspondence during training.

tions in the outdoor driving scenarios, i.e., ground plane with vehicles, buildings on the sides, horizontal lidar scans, and largely planar motion. The dataset bias is strong enough to render vertical flip detrimental to performance. This is also evident in existing works as none utilizes vertical flip as augmentation. Nonetheless, AugUndo still improves performance, where point removal models different lidar returns patterns, and resizing simulates large variations in scales of objects observed in road scenes. Similar to indoors, translations help model occlusions by shifting the projection with different principal points. Fig. 3 compares FusionNet and KBNet trained using conventional augmentations and AugUndo. AugUndo improves over objects that may have diverse orientations (i.e., tree branches and vegetation), thanks to rotation augmentations. Improvements are also observed in “small” objects at far regions where resizing can simulate zooming in/out to emulate objects of different sizes. Through translation and occlusion, AugUndo also improves on occlusion boundaries during training (highlighted), making estimates near image borders and object boundaries more robust.

We additionally show results for monocular depth estimation on KITTI in Tab. 6, where we compare the standard augmentation pipelines used by [19, 47, 102] and AugUndo. We observe similar trends in performance gain as we did in depth completion trained on KITTI: Applying our set of augmentations improves most metrics for all methods. We observe notable improvements in $\delta < 1.25$, the most difficult accuracy metric, from 0.869 to 0.879 in Monodepth2, from 0.879 to 0.883 in HR-Depth, and 0.862 to 0.863 in Lite-Mono. For Monodepth2, we also observe a large reduction in AbsRel error, improving it from 0.117 to 0.113.

Table 6: Monocular depth estimation on KITTI. AugUndo consistently improve on evaluation metrics (in meters) across different models.

Method	MAE ↓	RMSE ↓	AbsRel ↓	SqRel ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
Monodepth2	2.315±0.005	4.794±0.035	0.117±0.001	0.845 ±0.030	0.869±0.004	0.959±0.001	0.982 ±0.001
+ AugUndo	2.237 ±0.014	4.739 ±0.032	0.113 ±0.000	0.862±0.030	0.879 ±0.002	0.960 ±0.001	0.982 ±0.001
HR-Depth	2.226±0.004	4.626±0.032	0.113±0.001	0.797±0.022	0.879±0.002	0.961±0.001	0.982±0.000
+ AugUndo	2.185 ±0.013	4.610 ±0.029	0.111 ±0.001	0.794 ±0.021	0.883 ±0.001	0.962 ±0.001	0.983 ±0.001
Lite-Mono	2.338±0.005	4.821±0.027	0.121±0.001	0.875±0.007	0.862±0.002	0.955 ±0.001	0.980±0.000
+ AugUndo	2.314 ±0.030	4.780 ±0.055	0.120 ±0.001	0.849 ±0.019	0.863 ±0.004	0.955 ±0.001	0.981 ±0.001

Similarly, for Lite-Mono, we also boosted performance across all metrics, with a particularly high reduction in SqRel from 0.875 to 0.849.

Remarks. As AugUndo only augments the inputs and modifies the loss computations, which add negligible time to training, we note that the performance improvements from it are obtained nearly for free. Yet, the percentage gain, however, is similar to improvements obtained by each successive state-of-the-art: For example, the improvement of Lite-Mono over HR-Depth, and that of HR-Depth over PackNet for monocular depth estimation; similarly, the improvement of FusionNet over VOICED and KBNet over FusionNet for depth completion.

5 Discussion

Conventionally, data augmentation aims to seek visual invariance and create a collection of equivalent classes, i.e., identifying an image and its augmented variant as the same. For geometric tasks, the underlying equivalence is in the 3D scene structures under various illumination conditions, camera, viewpoints, occlusion, etc. Assuming a rigid scene, the shapes populating it should persist regardless of the nuisance variables. This motivates the use of geometric augmentations, as it simulates the nuisances within the image. However, adoption of geometric augmentations in unsupervised geometric tasks are obstructed by artifacts introduced during transformations (see Sec. D of Supp. Mat.). AugUndo lifts this obstacle by “undo-ing” the augmentations before computing the loss.

While AugUndo enables scaling up augmentations for unsupervised training, i.e., depth completion, it may also be applicable for supervised methods; though, we posit that the gain to be less as the artifacts induced from photometric and geometric augmentations of an image tend to be larger than those of a piecewise smooth depth map. AugUndo is also limited to 2D augmentations; whereas, nuisances modeled by it are projections of the 3D scene (see limitations in Sec. M in Supp. Mat.). We leave extensions to 3D for future work. We also only consider a single image and sparse depth map as input. Likewise, extensions can be made towards multi-frame tasks such as stereo, optical flow, pose estimation, etc., but one must account for their specifics and problem setups, i.e., stereo assumes frontoparallel views. This is outside of our scope, so we leave them as future directions. This work paves way for the empirical success in unsupervised geometric tasks that we have observed in other visual recognition tasks.

Acknowledgements. This work was supported by was supported by NSF 2112562 Athena AI Institute.

References

1. Berger, Z., Agrawal, P., Liu, T.Y., Soatto, S., Wong, A.: Stereoscopic universal perturbations across different architectures and datasets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15180–15190 (2022)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
3. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* **40**(4), 834–848 (2017)
4. Chen, Y., Schmid, C., Sminchisescu, C.: Self-supervised learning with geometric constraints in monocular video: Connecting flow, depth, and camera. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7063–7072 (2019)
5. Cheng, X., Wang, P., Guan, C., Yang, R.: Cspn++: Learning context and resource aware convolutional spatial propagation networks for depth completion. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 10615–10622 (2020)
6. Cheng, X., Wang, P., Yang, R.: Depth estimation via affinity learned with convolutional spatial propagation network. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 103–119 (2018)
7. Choi, H., Lee, H., Kim, S., Kim, S., Kim, S., Sohn, K., Min, D.: Adaptive confidence thresholding for monocular depth estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12808–12818 (2021)
8. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5828–5839 (2017)
9. Eigen, D., Fergus, R.: Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In: Proceedings of the IEEE international conference on computer vision. pp. 2650–2658 (2015)
10. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems* **27** (2014)
11. Eldekokey, A., Felsberg, M., Holmquist, K., Persson, M.: Uncertainty-aware cnns for depth completion: Uncertainty from beginning to end. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12014–12023 (2020)
12. Ezhov, V., Park, H., Zhang, Z., Upadhyay, R., Zhang, H., Chandrappa, C.C., Kadambi, A., Ba, Y., Dorsey, J., Wong, A.: All-day depth completion. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE (2024)

13. Fei, X., Wong, A., Soatto, S.: Geo-supervised visual depth prediction. *IEEE Robotics and Automation Letters* **4**(2), 1661–1668 (2019)
14. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24**(6), 381–395 (1981)
15. Garg, R., Bg, V.K., Carneiro, G., Reid, I.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII* 14. pp. 740–756. Springer (2016)
16. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* **32**(11), 1231–1237 (2013)
17. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: *2012 IEEE conference on computer vision and pattern recognition*. pp. 3354–3361. IEEE (2012)
18. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 270–279 (2017)
19. Godard, C., Mac Aodha, O., Firman, M., Brostow, G.J.: Digging into self-supervised monocular depth estimation. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 3828–3838 (2019)
20. Guizilini, V., Ambrus, R., Pillai, S., Raventos, A., Gaidon, A.: 3d packing for self-supervised monocular depth estimation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2485–2494 (2020)
21. Guizilini, V., Hou, R., Li, J., Ambrus, R., Gaidon, A.: Semantically-guided representation learning for self-supervised monocular depth. In: *International Conference on Learning Representations* (2019)
22. Han, M., Bushong, E.A., Segawa, M., Tiard, A., Wong, A., Brady, M.R., Momcilovic, M., Wolf, D.M., Zhang, R., Petcherski, A., et al.: Spatial mapping of mitochondrial networks and bioenergetics in lung cancer. *Nature* **615**(7953), 712–719 (2023)
23. Harris, C.G., Stephens, M., et al.: A combined corner and edge detector. In: *Alvey vision conference*. vol. 15, pp. 10–5244. Citeseer (1988)
24. Hu, M., Wang, S., Li, B., Ning, S., Fan, L., Gong, X.: Penet: Towards precise and efficient image guided depth completion. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 13656–13662. IEEE (2021)
25. Jaritz, M., De Charette, R., Wirbel, E., Perrotton, X., Nashashibi, F.: Sparse and dense data with cnns: Depth completion and semantic segmentation. In: *2018 International Conference on 3D Vision (3DV)*. pp. 52–60. IEEE (2018)
26. Jeon, J., Lim, H., Seo, D.U., Myung, H.: Struct-mdc: Mesh-refined unsupervised depth completion leveraging structural regularities from visual slam. *IEEE Robotics and Automation Letters* **7**(3), 6391–6398 (2022)
27. Kam, J., Kim, J., Kim, S., Park, J., Lee, S.: Costdcnet: Cost volume based depth completion for a single rgb-d image. In: *European Conference on Computer Vision*. pp. 257–274. Springer (2022)
28. Kingma, D.P., Ba, J.L.: Adam: A method for stochastic gradient descent. In: *ICLR: International Conference on Learning Representations*. pp. 1–15. ICLR US. (2015)
29. Klingner, M., Termöhlen, J.A., Mikolajczyk, J., Fingscheidt, T.: Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX* 16. pp. 582–600. Springer (2020)

30. Krishna, S., Vandrotti, B.S.: Deepsmooth: Efficient and smooth depth completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3357–3366 (2023)
31. Kumar, V.R., Klingner, M., Yogamani, S., Milz, S., Fingscheidt, T., Mader, P.: Syndistnet: Self-supervised monocular fisheye camera distance estimation synergized with semantic segmentation for autonomous driving. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 61–71 (2021)
32. Laina, I., Rupprecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks. In: 2016 Fourth international conference on 3D vision (3DV). pp. 239–248. IEEE (2016)
33. Lao, D., Sundaramoorthi, G.: Minimum delay moving object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4250–4259 (2017)
34. Lao, D., Sundaramoorthi, G.: Extending layered models to 3d motion. In: Proceedings of the European conference on computer vision (ECCV). pp. 435–451 (2018)
35. Lao, D., Sundaramoorthi, G.: Minimum delay object detection from video. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5097–5106 (2019)
36. Lao, D., Wang, C., Wong, A., Soatto, S.: Diffeomorphic template registration for atmospheric turbulence mitigation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 25107–25116 (2024)
37. Lao, D., Wu, Y., Liu, T.Y., Wong, A., Soatto, S.: Sub-token vit embedding via stochastic resonance transformers. In: International Conference on Machine Learning. PMLR (2024)
38. Lao, D., Yang, F., Wang, D., Park, H., Lu, S., Wong, A., Soatto, S.: On the viability of monocular depth pre-training for semantic segmentation. In: European Conference on Computer Vision. Springer (2024)
39. Lepetit, V., Moreno-Noguer, F., Fua, P.: Epnp: An accurate $O(n)$ solution to the pnp problem. International journal of computer vision **81**(2), 155 (2009)
40. Li, A., Yuan, Z., Ling, Y., Chi, W., Zhang, C., 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. pp. 32–40 (2020)
41. Li, B., Shen, C., Dai, Y., Van Den Hengel, A., He, M.: Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1119–1127 (2015)
42. Lin, Y., Cheng, T., Zhong, Q., Zhou, W., Yang, H.: Dynamic spatial propagation network for depth completion. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 1638–1646 (2022)
43. Liu, C., Chen, L.C., Schroff, F., Adam, H., Hua, W., Yuille, A.L., Fei-Fei, L.: Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 82–92 (2019)
44. Liu, F., Shen, C., Lin, G., Reid, I.: Learning depth from single monocular images using deep convolutional neural fields. IEEE transactions on pattern analysis and machine intelligence **38**(10), 2024–2039 (2015)
45. Liu, T.Y., Agrawal, P., Chen, A., Hong, B.W., Wong, A.: 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. pp. 35–53. Springer (2022)

46. Lopez-Rodriguez, A., Busam, B., Mikolajczyk, K.: Project to adapt: Domain adaptation for depth completion from noisy and sparse sensor data. In: Proceedings of the Asian Conference on Computer Vision (2020)
47. Lyu, X., Liu, L., Wang, M., Kong, X., Liu, L., Liu, Y., Chen, X., Yuan, Y.: Hr-depth: High resolution self-supervised monocular depth estimation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 2294–2301 (2021)
48. Ma, F., Cavalheiro, G.V., Karaman, S.: Self-supervised sparse-to-dense: Self-supervised depth completion from lidar and monocular camera. In: International Conference on Robotics and Automation (ICRA). pp. 3288–3295. IEEE (2019)
49. Mahjourian, R., Wicke, M., Angelova, A.: Unsupervised learning of depth and ego-motion from monocular video using 3d geometric constraints. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5667–5675 (2018)
50. Park, H., Gupta, A., Wong, A.: Test-time adaptation for depth completion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20519–20529 (2024)
51. Park, J., Joo, K., Hu, Z., Liu, C.K., Kweon, I.S.: Non-local spatial propagation network for depth completion. In: European Conference on Computer Vision, ECCV 2020. European Conference on Computer Vision (2020)
52. Perez, L., Wang, J.: The effectiveness of data augmentation in image classification using deep learning. arXiv preprint arXiv:1712.04621 (2017)
53. Poggi, M., Aleotti, F., Tosi, F., Mattoccia, S.: On the uncertainty of self-supervised monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3227–3237 (2020)
54. Poggi, M., Tosi, F., Mattoccia, S.: Learning monocular depth estimation with unsupervised trinocular assumptions. In: 2018 International conference on 3d vision (3DV). pp. 324–333. IEEE (2018)
55. Qiu, J., Cui, Z., Zhang, Y., Zhang, X., Liu, S., Zeng, B., Pollefeys, M.: 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. pp. 3313–3322 (2019)
56. Qu, C., Liu, W., Taylor, C.J.: Bayesian deep basis fitting for depth completion with uncertainty. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 16147–16157 (2021)
57. Qu, C., Nguyen, T., Taylor, C.: Depth completion via deep basis fitting. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 71–80 (2020)
58. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12179–12188 (2021)
59. Ranftl, R., Lasinger, K., Hafner, D., Schindler, K., Koltun, V.: Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. IEEE transactions on pattern analysis and machine intelligence **44**(3), 1623–1637 (2020)
60. Ranjan, A., Jampani, V., Balles, L., Kim, K., Sun, D., Wulff, J., Black, M.J.: Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12240–12249 (2019)

61. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 779–788 (2016)
62. Saxena, A., Sun, M., Ng, A.Y.: Make3d: Learning 3d scene structure from a single still image. *IEEE transactions on pattern analysis and machine intelligence* **31**(5), 824–840 (2008)
63. Shrivakumar, S.S., Nguyen, T., Miller, I.D., Chen, S.W., Kumar, V., Taylor, C.J.: Dfusenet: Deep fusion of rgb and sparse depth information for image guided dense depth completion. In: 2019 IEEE Intelligent Transportation Systems Conference (ITSC). pp. 13–20. IEEE (2019)
64. Shorten, C., Khoshgoftaar, T.M.: A survey on image data augmentation for deep learning. *Journal of big data* **6**(1), 1–48 (2019)
65. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Indoor segmentation and support inference from rgbd images. In: European conference on computer vision. pp. 746–760. Springer (2012)
66. Singh, A.D., Ba, Y., Sarker, A., Zhang, H., Kadambi, A., Soatto, S., Srivastava, M., Wong, A.: Depth estimation from camera image and mmwave radar point cloud. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9275–9285 (2023)
67. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7262–7272 (2021)
68. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8934–8943 (2018)
69. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2446–2454 (2020)
70. Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10781–10790 (2020)
71. Taylor, L., Nitschke, G.: Improving deep learning with generic data augmentation. In: 2018 IEEE symposium series on computational intelligence (SSCI). pp. 1542–1547. IEEE (2018)
72. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: European conference on computer vision. pp. 402–419. Springer (2020)
73. Tosi, F., Aleotti, F., Poggi, M., Mattoccia, S.: Learning monocular depth estimation infusing traditional stereo knowledge. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9799–9809 (2019)
74. Uhrig, J., Schneider, N., Schneider, L., Franke, U., Brox, T., Geiger, A.: Sparsity invariant cnns. In: 2017 international conference on 3D Vision (3DV). pp. 11–20. IEEE (2017)
75. Upadhyay, R., Zhang, H., Ba, Y., Yang, E., Gella, B., Jiang, S., Wong, A., Kadambi, A.: Enhancing diffusion models with 3d perspective geometry constraints. *ACM Transactions on Graphics (TOG)* **42**(6), 1–15 (2023)
76. Van Gansbeke, W., Neven, D., De Brabandere, B., Van Gool, L.: Sparse and noisy lidar completion with rgb guidance and uncertainty. In: 2019 16th International Conference on Machine Vision Applications (MVA). pp. 1–6. IEEE (2019)

77. Wang, C., Buenaposada, J.M., Zhu, R., Lucey, S.: Learning depth from monocular videos using direct methods. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2022–2030 (2018)
78. Wang, F., Galliani, S., Vogel, C., Speciale, P., Pollefeys, M.: Patchmatchnet: Learned multi-view patchmatch stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14194–14203 (2021)
79. Watson, J., Firman, M., Brostow, G.J., Turmukhambetov, D.: Self-supervised monocular depth hints. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2162–2171 (2019)
80. Wong, A., Chen, A., Wu, Y., Cicek, S., Tiard, A., Hong, B.W., Soatto, S.: Small lesion segmentation in brain mrIs with subpixel embedding. In: International MICCAI Brainlesion Workshop. pp. 75–87. Springer (2021)
81. Wong, A., Cicek, S., Soatto, S.: Targeted adversarial perturbations for monocular depth prediction. *Advances in neural information processing systems* **33**, 8486–8497 (2020)
82. Wong, A., Cicek, S., Soatto, S.: Learning topology from synthetic data for unsupervised depth completion. *IEEE Robotics and Automation Letters* **6**(2), 1495–1502 (2021)
83. Wong, A., Fei, X., Hong, B.W., Soatto, S.: An adaptive framework for learning unsupervised depth completion. *IEEE Robotics and Automation Letters* **6**(2), 3120–3127 (2021)
84. Wong, A., Fei, X., Tsuei, S., Soatto, S.: Unsupervised depth completion from visual inertial odometry. *IEEE Robotics and Automation Letters* **5**(2), 1899–1906 (2020)
85. Wong, A., Mundhra, M., Soatto, S.: Stereopagnosia: Fooling stereo networks with adversarial perturbations. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 2879–2888 (2021)
86. Wong, A., Soatto, S.: Bilateral cyclic constraint and adaptive regularization for unsupervised monocular depth prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5644–5653 (2019)
87. Wong, A., Soatto, S.: Unsupervised depth completion with calibrated backprojection layers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12747–12756 (2021)
88. Xu, D., Ricci, E., Ouyang, W., Wang, X., Sebe, N.: Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5354–5362 (2017)
89. Xu, H., Zhang, J.: Aanet: Adaptive aggregation network for efficient stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1959–1968 (2020)
90. Xu, Y., Zhu, X., Shi, J., Zhang, G., Bao, H., Li, H.: Depth completion from sparse lidar data with depth-normal constraints. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2811–2820 (2019)
91. Yan, Z., Wang, K., Li, X., Zhang, Z., Li, J., Yang, J.: Desnet: Decomposed scale-consistent network for unsupervised depth completion. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 3109–3117 (2023)
92. Yang, F., Feng, C., Chen, Z., Park, H., Wang, D., Dou, Y., Zeng, Z., Chen, X., Gangopadhyay, R., Owens, A., Wong, A.: Binding touch to everything: Learning unified multimodal tactile representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 26340–26353 (2024)

93. Yang, N., Stumberg, L.v., Wang, R., Cremers, D.: D3vo: Deep depth, deep pose and deep uncertainty for monocular visual odometry. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1281–1292 (2020)
94. Yang, Y., Wong, A., Soatto, S.: Dense depth posterior (ddp) from single image and sparse range. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3353–3362 (2019)
95. Yang, Z., Wang, P., Wang, Y., Xu, W., Nevatia, R.: Every pixel counts: Unsupervised geometry learning with holistic 3d motion understanding. In: Proceedings of the European conference on computer vision (ECCV) workshops. pp. 0–0 (2018)
96. Yang, Z., Wang, P., Wang, Y., Xu, W., Nevatia, R.: Lego: Learning edge with geometry all at once by watching videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 225–234 (2018)
97. Yin, W., Liu, Y., Shen, C., Yan, Y.: Enforcing geometric constraints of virtual normal for depth prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5684–5693 (2019)
98. Yin, Z., Shi, J.: Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1983–1992 (2018)
99. Yu, Z., Sheng, Z., Zhou, Z., Luo, L., Cao, S.Y., Gu, H., Zhang, H., Shen, H.L.: Aggregating feature point cloud for depth completion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8732–8743 (2023)
100. Zeng, Z., Wang, D., Yang, F., Park, H., Soatto, S., Lao, D., Wong, A.: Wordepth: Variational language prior for monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9708–9719 (2024)
101. Zhan, H., Garg, R., Weerasekera, C.S., Li, K., Agarwal, H., Reid, I.: Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 340–349 (2018)
102. Zhang, N., Nex, F., Vosselman, G., Kerle, N.: Lite-mono: A lightweight cnn and transformer architecture for self-supervised monocular depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18537–18546 (2023)
103. Zhang, X., Pak, D.H., Ahn, S.S., Li, X., You, C., Staib, L., Sinusas, A.J., Wong, A., Duncan, J.S.: Heteroscedastic uncertainty estimation for probabilistic unsupervised registration of noisy medical images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer (2024)
104. Zhang, X., Stendahl, J.C., Staib, L., Sinusas, A.J., Wong, A., Duncan, J.S.: An adaptive correspondence scoring framework for unsupervised image registration of medical images. In: European Conference on Computer Vision. Springer (2024)
105. Zhang, Y., Funkhouser, T.: Deep depth completion of a single rgb-d image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 175–185 (2018)
106. Zhang, Y., Guo, X., Poggi, M., Zhu, Z., Huang, G., Mattoccia, S.: Completion-former: Depth completion with convolutions and vision transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18527–18536 (2023)
107. Zhao, C., Zhang, Y., Poggi, M., Tosi, F., Guo, X., Zhu, Z., Huang, G., Tang, Y., Mattoccia, S.: Monovit: Self-supervised monocular depth estimation with a vision

transformer. In: 2022 International Conference on 3D Vision (3DV). pp. 668–678. IEEE (2022)

108. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1851–1858 (2017)
109. Zou, Y., Luo, Z., Huang, J.B.: Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In: Proceedings of the European conference on computer vision (ECCV). pp. 36–53 (2018)

AugUndo: Scaling Up Augmentations for Monocular Depth Completion and Estimation

Supplementary Material

Summary. In Sec. A, we provide implementation details and hyperparameters, such as learning rate schedule and crop size, used to reproduce our results. In B, we provide augmentation parameters used to train each method on each dataset. Sec. C lists the evaluation metrics used in our quantitative results. To illustrate the motivation behind AugUndo, we provide examples of image and sparse depth degradation through typical augmentation routines and add an extended discussion in Sec. D. We provide details on datasets used in our experiments in Sec. E. In Sec. F, we provide a detail walk-through of the AugUndo algorithm. To study the contribution of each of the proposed augmentation, we present a comprehensive ablation study in Sec. G. We further extend the sensitivity study to sparsity from the main paper conducted on VOID to VOID500 and VOID150 in Sec. H. We present our results on zero-shot generalization from KITTI to Waymo (depth completion) and Make3d (monocular depth estimation) in Sec. I. We also show a sensitivity study on the effect of different augmentations in Sec. J. In Sec. K, we include additional results for depth completion, including extensive quantitative results such as evaluations on MonDi and DesNet, comparing modeling AugUndo as change in camera pose and parameters, and the effects of naively applying geometric augmentations during unsupervised training. In Sec. L, we show additional results on monocular depth estimation. Finally, we discuss limitations in Sec. M.

A Implementation details

For unsupervised depth completion, we implemented our method in PyTorch and incorporated our augmentation pipeline into the codebases of VOICED [84], FusionNet [82], KBNet [87], MonDi [45], and DesNet [91]. For monocular depth estimation, we implement our method in PyTorch according to [19]. Specifically, we implemented our augmentation pipeline into the codebases of Monodepth2 [19], HR-Depth [47], and Lite-Mono [102]. Details of each task are described below.

Unsupervised depth completion. The models are optimized using Adam [28] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. For VOID: We used an input batch size of 12 with random crop size of 416×512 for KBNet and DesNet, a batch size of 8 without cropping for FusionNet and VOICED, and a batch size of 8 with random crop size of 448×576 for MonDi. We trained each KBNet and DesNet for 40 epochs with base learning rate of 1×10^{-4} for 20 epochs and decreased to 5×10^{-5} for the last 20 epochs. We trained FusionNet and VOICED for 20 epochs with a base learning rate of 1×10^{-4} for 10 epochs and decreased to 5×10^{-5} for the last 10 epochs. For KITTI: We used a batch size of 8 and random crop size of

320 by 768 for all models. We trained KBNet for 60 epochs, with 5×10^{-5} for 2 epochs, 1×10^{-4} until 8th epoch, 2×10^{-4} until 30th epoch, 1×10^{-4} until the 45th epoch and 5×10^{-5} until the 60th epoch. We trained FusionNet for 30 epochs 2×10^{-4} for 16 epochs, 1×10^{-4} until 24th epoch and 5×10^{-5} until the 30th epoch. We trained VOICED for 30 epochs 2×10^{-4} for 16 epochs, 6×10^{-5} until 24th epoch and 3×10^{-5} until the 30th epoch.

We ensure that all baseline methods can reproduce or exceed the numbers originally reported by the authors. Since the authors of DesNet did not publish their code implementation, we re-implemented their method to the best of our ability. All reported results of AugUndo are based on those same settings with the exception of the augmentation scheme.

Unsupervised monocular depth estimation. The models are optimized using Adam [28] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. For VOID, we used a input batch size of 12 and random crop size of 256×448 for all models. For Monodepth2, we trained with learning rate 1×10^{-8} for 2 epochs, 1×10^{-5} for the next 23 epochs, and 1×10^{-6} for the next 25 epochs. For HR-Depth, we train with learning rate 1^{-4} for 15 epochs, and 1×10^{-5} for the final 5 epochs. For Lite-Mono, we train with learning rate 5×10^{-4} for 35 epochs. For KITTI dataset, We used a input batch size of 12 and a random crop size of 192×640 . We trained the models for 20 epochs with an initial learning rate of 1×10^{-4} and drop the learning rate to 1×10^{-5} at the 15th epoch. The smoothness loss weight for KITTI is set to 0.001 as per [19] and 0.01 for VOID as VOID contains more indoor scenes with homogeneous surfaces.

We ensure that all baseline methods can reproduce or exceed the numbers originally reported by the authors. All reported results of AugUndo are based on those same settings with the exception of the augmentation scheme. Note: For Monodepth2 and HR-Depth, we initialize the ResNet encoder weight with the Imagenet-pretrained weight downloaded from PyTorch website, as specified in their Github repository. However, we cannot locate the pretrained weight used for Lite-Mono throughout their repository, which left us no choice but to train their model from scratch. Nonetheless, this does not affect the validity of the study as the comparison is made between models initialized from scratch where the difference is only in the augmentations schemes: standard convention used in the original papers and repositories, and AugUndo.

B Augmentations

Through a search over augmentation types and values, we found a consistent set of augmentations that tends to yield improvements across all methods with small changes to degree of augmentation catered to each method. All augmentations are applied with a 50% probability. We note that performance gain can be obtained by typical set and ranges of augmentations and does not require a meticulous selection of hyper-parameters (see Fig. 5).

VOID. For depth completion, we applied photometric transformations of random brightness from 0.5 to 1.5, contrast from 0.5 to 1.5, saturation from 0.5

to 1.5, and hue from -0.1 to 0.1. We applied image patch removal by selecting between 0.1% to 0.5% of the pixels and removing 5×5 patches centered on them – approximately removing between 2.5% to 12.5% of the image. We applied random sparse depth point removal at a rate between 60% and 70% of all sparse points. We applied geometric transformations of random horizontal and vertical flips, up to 10% of translation and between -25 to 25 degrees of rotation. For KBNet and FusionNet, we applied random resize factor between 0.6 to 1.1, while for VOICED, we used 0.7 to 1.1 because the scaffolding step of VOICED requires at least 3 points, but as discussed above, resampling causes loss and factors smaller than 0.7 often cause all points to be dropped. For monocular depth estimation, we applied photometric transformations of random brightness from 0.5 to 1.5, contrast from 0.5 to 1.5, saturation from 0.5 to 1.5, hue from -0.1 to 0.1. We applied geometric transformation of random rotation between -10 to 10 degrees and random horizontal flipping. We further applied random resize factor between 0.8 to 1.

KITTI. For depth completion, we applied random brightness, contrast, saturation from 0.5 to 1.5 and random hue from -0.1 to 0.1. We applied image patch removal by selecting between 0.1% to 0.5% of the pixels and removing 5×5 patches centered on them. We applied random sparse depth point removal at a rate between 60% and 70% of all sparse points. We further applied random horizontal flips, up to 10% of translation, resizing factors between 0.8 to 1.2, and between -20 to 20 degrees of rotation. We found that vertical flips are detrimental to performance. For monocular depth estimation, we applied random brightness, contrast, saturation from 0.5 to 1.5 and random hue from -0.1 to 0.1. We applied a random rotation between -30 to 30 degrees and a random translation of up to 30% of the image. We further apply horizontal flip to the image.

Table 7: Error metrics for depth completion and monocular depth estimation. d_{gt} denotes ground truth depth and evaluated where values are available for a given image.

Metric	Definition
MAE	$\frac{1}{ \Omega } \sum_{x \in \Omega} \hat{d}(x) - d_{gt}(x) $
RMSE	$\left(\frac{1}{ \Omega } \sum_{x \in \Omega} \hat{d}(x) - d_{gt}(x) ^2 \right)^{1/2}$
iMAE	$\frac{1}{ \Omega } \sum_{x \in \Omega} 1/\hat{d}(x) - 1/d_{gt}(x) $
iRMSE	$\left(\frac{1}{ \Omega } \sum_{x \in \Omega} 1/\hat{d}(x) - 1/d_{gt}(x) ^2 \right)^{1/2}$
AbsRel	$\frac{1}{ \Omega } \sum_{x \in \Omega} \frac{ \hat{d}(x) - d_{gt}(x) }{d_{gt}(x)}$
SqRel	$\frac{1}{ \Omega } \sum_{x \in \Omega} \frac{ \hat{d}(x) - d_{gt}(x) ^2}{d_{gt}(x)}$
Accuracy	% of $z(x)$ s.t. $\delta \doteq \max\left(\frac{z(x)}{z_{gt}(x)}, \frac{z_{gt}(x)}{z(x)}\right) < \text{threshold}$

C Evaluation metrics

The evaluation metrics used for depth completion and monocular depth estimation are shown in Tab. 7. Depth completion models are evaluated with MAE,

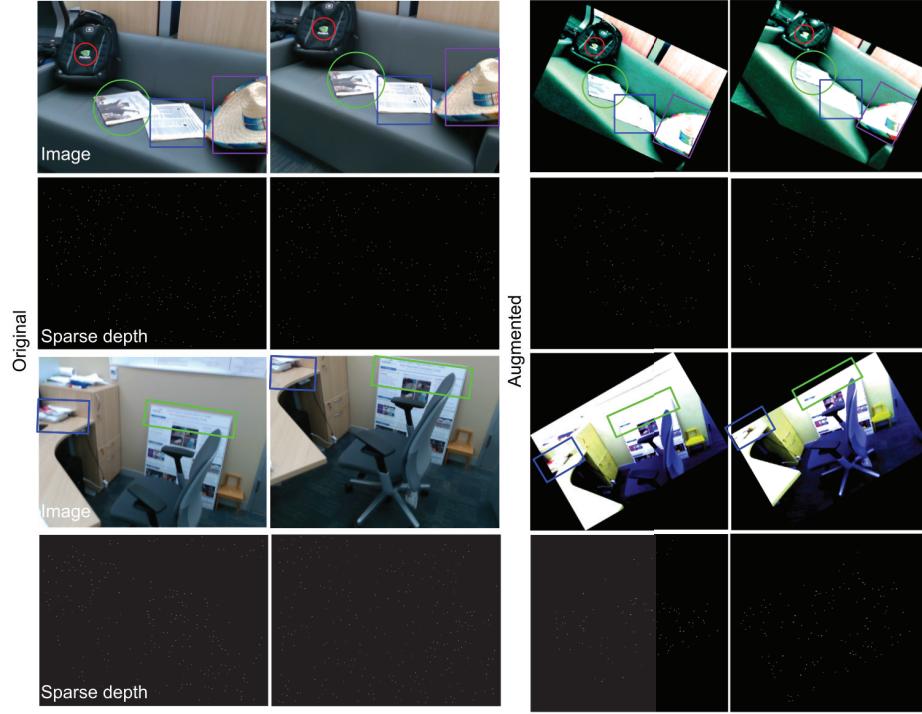


Fig. 4: Motivation. After augmenting the inputs with photometric and geometric transformations, the previously exciting image textures (for establishing correspondence across images) and sparse points in the original inputs, that would have served as supervision, are now largely saturated in intensity and homogeneous, and lost due to resampling, respectively (best viewed in 2 \times).

RMSE, iMAE, iRMSE. Monocular depth estimation models are evaluated with MAE, RMSE, AbsRel, SqRel, and accuracy ($\delta < 1.25$, $\delta < 1.25^2$, $\delta < 1.25^3$). We note that as monocular depth is inferred, at most, up to an unknown scale, we perform scale matching during evaluation by using median scaling with respect to the ground truth before computing each error or accuracy metric.

D Extended Discussion of Motivation

In the main paper, we discussed the motivation behind the approach. Here, we provide an extended discussion: Training unsupervised depth completion methods [82–84, 87, 94] relies on a photometric reconstruction term, sparse depth reconstruction term, and a (generic) regularizer, such as local smoothness; a related problem, unsupervised monocular depth estimation [13, 19, 38, 81, 86, 100], omits the sparse depth term. The photometric reconstruction term constraints the solution in regions where there are sufficiently exciting textures and co-visible between images such as those highlighted in Fig. 4. The sparse depth

reconstruction term constraints the solution anywhere with a sparse point. Everywhere else in the image are inherently ambiguous, so we must rely on the regularizer. Fig. 4 shows that after one has applied a number of photometric and geometric transformations, the previously exciting textures that would have served as our supervision are now largely saturated in intensity and homogenous. The number of points in the sparse depth maps have also been greatly reduced.

In conventional augmentation scheme for *other* (semantic) vision tasks, i.e. classification, detection, and segmentation, it is typical for one to introduce saturation in image intensities, block artifacts from interpolation and loss during resampling. These are some of the side-effects that are *desirable* in the aforementioned tasks: introducing these nuisance variabilities enables the model to learn to be invariant to them (i.e., to learn them away), which yields more generalizable and robust representations. This is also true for supervised depth completion [12, 51, 66, 76] and estimation methods [9, 58, 59, 75, 97]. However, this is not the case for learning unsupervised depth. Because the supervision signal comes from reconstructing the input image and sparse depth map, the more we augment the data (causing a loss of photometric correspondences across image frames and a loss of points in the sparse depth map), the more we degrade the supervision signal (see Tabs. 14 and 15 in Sec. K for empirical evidence of this phenomenon). So, it is not too surprising that conventional augmentation procedures, both photometric and geometric, have seen limited use beyond small changes in image intensities, and flipping.

Nevertheless, photometric augmentations help model the diverse range of illumination conditions and colors of object that may populate the scene. Geometric augmentations can simulate the various camera parameters and motion, for example, image translation can simulate principle point offsets or it can approximate small baseline movements, image resizing with cropping can simulate different focal lengths (zooming) or it can model camera movements, while in-plane rotations can model camera orientation. These augmentations are often viewed as essential to training pipelines for other vision tasks, but are detrimental for unsupervised depth completion; yet, without them, one may encounter robustness and generalization issues. As the root of the problem lies in the supervision signal, we investigate an approach to “undo” the augmentations during (or right before) the loss computation step. Doing so enables one to feed forward augmented inputs, but compute the loss on the original inputs with no loss in the training signal. Extensive experiments show that our approach improves performance and zero-shot generalization across a number of methods for both indoor and outdoor scenarios.

E Datasets

We conduct experiments on six datasets (two for training and four for generalization) in total. Details of each dataset are provided below.

KITTI [16] contains 61 driving scenes with research in autonomous driving and computer vision. It contains calibrated RGB images with synchronized point

clouds from Velodyne lidar, inertial, GPS information, etc. For depth completion [74], there are $\approx 80,000$ raw image frames and associated sparse depth maps, each with a density of $\approx 5\%$. Ground-truth depth is obtained by accumulating 11 neighbouring raw lidar scans. Semi-dense depth is available for the lower 30% of the image space. We test on the official validation set of 1,000 samples because the online test server has submission restrictions to accommodate multiple trials. For depth estimation, we used Eigen split [9], following [108] to preprocess and remove static frames. The remaining training set contains 39,810 monocular triplets and the validation set contains 4,424 triplets. The testing set contains 697 monocular images. We follow the evaluation protocol of [10, 74], where output depth is evaluated where ground truth exists between 0 to 80.0 meters.

VOID [84] comprises indoor (laboratories, classrooms) and outdoor (gardens) scenes with synchronized 640×480 RGB images and sparse depth maps. XIVO [13], a VIO system, is used to obtain the sparse depth maps that contain approximately 1500 sparse depth points with a density of about 0.5%. Active stereo is used to acquire the dense ground-truth depth maps. In contrast to the typically planar motion in KITTI, VOID has 56 sequences with challenging 6 DoF motion captured on rolling shutter. 48 sequences (about 45,000 frames) are assigned for training and 8 for testing (800 frames). We follow the evaluation protocol of [84] where output depth is evaluated where ground truth exists between 0.2 and 5.0 meters.

NYUv2 [65] consists of 372K synchronized 640×480 RGB images and depth maps for 464 indoors scenes (household, offices, commercial), captured with a Microsoft Kinect. The official split consisting in 249 training and 215 test scenes. We use the official test set of 654 images. Because there are no sparse depth maps provided, we sampled ≈ 1500 points from the depth map via Harris corner detector [23] to mimic the sparse depth produced by SLAM/VIO. We test models trained on VOID to evaluate their generalization to NYUv2. We follow the evaluation protocol of [84] where output depth is evaluated where ground truth exists between 0.2 and 5.0 meters.

ScanNet [8] consists of RGB-D data for 1,513 indoor scenes with 2.5 million images and corresponding dense depth map. Because there are no sparse depth maps provided, we sampled ≈ 1500 points from the depth map via Harris corner detector [23] to mimic the sparse depth produced by SLAM/VIO. We followed [8] and used 100 scenes (scene707-scene806), for zero-shot generalization for models trained on VOID. The output depth is evaluated where ground truth exists between 0.2 and 5.0 meters.

Waymo Open Dataset [69] contains 1920×1280 RGB images and lidar scans from autonomous vehicles. The training set contains ≈ 158 K images from 798 scenes and the validation set ≈ 40 K images from 202 scenes, collected at 10Hz. Objects are annotated across the full 360° field. Sparse depth maps are obtained by reprojecting the point cloud scan from the top lidar to the camera frame. Ground truth is obtained by reprojecting both front facing lidars as well as those collected 10 time steps forward and backwards (approximately 1 second of capture) to a given camera frame at a specific time step to densify the

Algorithm 1 AUGUNDO

Require: Depth completion network f_θ , Images I_t, I_τ , Sparse depth z_t ,
 Relative pose g_{rt} , Intrinsic K

- 1: Sample $\{T_{pt,I}^1 \dots T_{pt,I}^k\}$ from $T_{pt,I}^i \in \mathcal{A}_{pt,I}$, and compose

$$T_{pt,I} = T_{pt,I}^1 \circ T_{pt,I}^2 \circ \dots \circ T_{pt,I}^k$$
- 2: Sample $\{T_{pt,z}^1 \dots T_{pt,z}^j\}$ from $T_{pt,z}^i \in \mathcal{A}_{pt,z}$, and compose

$$T_{pt,z} = T_{pt,z}^1 \circ T_{pt,z}^2 \circ \dots \circ T_{pt,z}^j$$
- 3: Sample $\{T_{ge}^1 \dots T_{ge}^m\}$ from $T_{ge}^i \in \mathcal{A}_{ge}$, and compose

$$T_{ge} = T_{ge}^1 \circ T_{ge}^2 \circ \dots \circ T_{ge}^m$$
- 4: Compose the inverse geometric transform

$$T_{ge}^{-1} = (T_{ge}^m)^{-1} \circ (T_{ge}^{m-1})^{-1} \circ \dots \circ (T_{ge}^1)^{-1}$$
- 5: Compute the coordinates after geometric transform

$$[x' \ 1]^\top = T_{ge} [x \ 1]^\top$$
 (Eqn. 3 from main paper)
- 6: Augment I_t with photometric and geometric transformations

$$I'_t(x') = T_{pt,I}(I_t)(x)$$
 (Eqn. 4 from main paper)
- 7: Augment z_t with occlusion and geometric transformations

$$z'_t(x') = T_{pt,z}(z_t)(x)$$
 (Eqn. 4 from main paper)
- 8: Obtain depth prediction $\hat{d}_t' = f_\theta(I'_t, z'_t)$
- 9: Compute coordinates of the inverse geometric transformation

$$[x'' \ 1]^\top = T_{ge}^{-1} [x' \ 1]^\top$$
 (Eqn. 5 from main paper)
- 10: Apply inverse geometric transformation on output depth map:

$$\hat{d}_t(x'') = \hat{d}_t(x')$$
 (Eqn. 6 from main paper)
- 11: Reconstruct I_t from I_τ using Eqn. 1 from main paper, i.e., $\hat{I}_{t\tau} = I_\tau(\pi g_{\tau t} K^{-1} \bar{x} \hat{d}_t)$
- 12: Minimize reconstruction losses between $\hat{I}_{t\tau}$ and I_t , and \hat{d}_t and z_t , and the regularizer (Eqn. 2 from main paper)

sparse depth. We used the object annotations to remove all moving objects to ensure that reprojected points respects the static scene assumption. We also performed outlier removal to filter out erroneous (noisy) points. The output depth is evaluated where ground truth exists between a 1.5 and 80.0 meters range.

Make3d [62] contains 134 test images with 2272×1707 resolution. Ground-truth depth maps are given at 305×55 resolution and must be rescaled and interpolated. We use the central cropping proposed by [18] to get a 852×1707 center crop of the image. We use standard Make3d evaluation protocol and metrics. We use Make3d to test the generalization of monocular depth estimation models trained on KITTI.

F The AugUndo Algorithm

We assume that we are given (i) $I : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}_+^3$ an RGB image I_t , (ii), its associated sparse point cloud projected onto as a depth map $z : \Omega_z \subset \Omega \rightarrow \mathbb{R}_+$, z_t , (iii) camera intrinsic calibration matrix $K \in \mathbb{R}^{3 \times 3}$, (iv) a sequence of images I_τ for $\tau \in \{t-1, t+1\}$ during training, and (v) the relative pose $g_{\tau t}$ between the image frame I_t and some temporally adjacent image I_τ . Given a image and its associated sparse depth map, a depth completion network learns a map

Table 8: *Ablation study for KBNet on VOID dataset.* BRI stands for brightness, CON for contrast, SAT for saturation, FLP for horizontal and vertical flip, TRN for translation, ROT for rotation, RZD for resize down, RZU for resize up, RMP for point removal, RMI for image patch removal. **Bold** denotes AugUndo, *italicized* the standard augmentation protocol. RMP, FLP, RZD, RZU have the highest influence; only when BRI, CON, SAT, HUE (color jitter) are disabled do they have non-negligible effect. Best results are achieved by using AugUndo.

Augmentation settings										Evaluation metrics				
BRI	CON	SAT	HUE	FLP	TRN	ROT	RZD	RZU	RMP	RMI	MAE	RMSE	iMAE	iRMSE
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	33.32 ± 0.18	85.67 ± 0.39	16.61 ± 0.29	41.24 ± 0.60
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	<i>38.11</i> ± 0.77	<i>95.22</i> ± 1.72	<i>19.51</i> ± 0.14	<i>46.70</i> ± 0.48
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	33.46 ± 0.27	86.02 ± 0.69	16.84 ± 0.22	41.78 ± 0.63
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	33.79 ± 0.09	86.38 ± 0.36	17.12 ± 0.19	42.12 ± 0.54
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	33.73 ± 0.30	85.84 ± 0.22	17.05 ± 0.16	41.73 ± 0.47
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	33.60 ± 0.16	86.47 ± 0.69	16.95 ± 0.27	41.65 ± 0.44
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	34.14 ± 0.37	87.26 ± 0.74	17.09 ± 0.17	42.57 ± 1.33
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	44.38 ± 0.89	106.87 ± 0.89	22.91 ± 0.68	52.55 ± 0.57
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	33.92 ± 0.19	87.19 ± 0.55	17.09 ± 0.03	42.23 ± 0.23
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	33.68 ± 0.19	85.86 ± 0.51	16.92 ± 0.03	41.47 ± 0.18
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	35.77 ± 0.43	89.88 ± 0.84	17.81 ± 0.23	42.75 ± 0.38
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	35.69 ± 0.25	90.40 ± 0.71	17.75 ± 0.12	42.82 ± 0.20
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	37.73 ± 0.27	92.62 ± 0.20	19.44 ± 0.18	45.47 ± 0.37
✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	33.64 ± 0.22	86.26 ± 0.23	16.80 ± 0.13	41.60 ± 0.64

the inputs to the output depth map $\hat{d}_t := f_\theta(I_t, z_t) \in \mathbb{R}_+^{H \times W}$. In the main paper, we denoted photometric and occlusion augmentations as \mathcal{A}_{pt} for ease of notation. Here, for specificity, we define $\mathcal{A}_{pt,I}$ as the set of possible photometric (including occlusion) transformations for the image, $\mathcal{A}_{pt,z}$ as the possible set of occlusion augmentations for sparse depth maps, and \mathcal{A}_{ge} as the possible set of all geometric transformations. We additionally assume we have the set of geometric transformations T_{ge} used during augmentation and their inverse transformations T_{ge}^{-1} . Alg. 1 is the procedural algorithm of AugUndo and details the step by step augmentation and loss computation pipelines, given our inputs.

G Ablation Study for Depth Completion

In the main paper, we demonstrated AugUndo for unsupervised depth completion methods on VOID1500 and VOID500 benchmarks using the most performant set of augmentations we found. Here, we provide an ablation study for each of the augmentations used to study their individual contributions.

Tab. 8 shows a comprehensive ablation study following augmentation settings above. From our best reported settings on KBNet, we removed individual augmentations to show their empirical contribution. Note: random brightness, contrast, saturation, and hue (BRI, SAT, CON, HUE) together is the standard color jitter employed by all current methods. We also test removing all color jitter augmentations to further quantify its effect.

We found that random flips (FLP) has the largest impact of all of the augmentations – increasing the error by an average of 30%; this is because it simulates

Table 9: *Sensitivity study of depth completion on VOID*. Reported scores are mean and standard deviation over four independent trials. We compare the sensitivity of models trained on VOID1500, with standard augmentations and AugUndo, by testing them on VOID500 and VOID150. Sparse depth maps within VOID1500 contains approximately 1500 points, and those within VOID500 and VOID150 contain approximately 3 \times and 10 \times less, respectively. AugUndo improves performance by an average of 19.66% and 19.62% across all unsupervised methods and evaluation metrics on VOID500 and VOID150, respectively. For distillation method (marked by *), MonDi, AugUndo improves by an average of 9.02% and 10.25% on VOID500 and VOID150, respectively. Despite MonDi is supervised by pseudo ground truth in addition to typical unsupervised losses, AugUndo can still boost performance.

Method	VOID500				VOID150			
	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓
VOICED	137.01 \pm 4.23	235.80 \pm 7.82	71.36 \pm 1.86	130.63 \pm 5.66	209.59 \pm 5.18	329.71 \pm 10.01	130.45 \pm 5.63	229.79 \pm 14.09
+ AugUndo	92.99 \pm 1.11	176.94 \pm 1.38	46.43 \pm 0.85	91.10 \pm 1.64	151.77 \pm 1.99	262.61 \pm 2.19	88.65 \pm 0.76	169.29 \pm 2.71
FusionNet	97.73 \pm 0.73	194.32 \pm 1.36	58.65 \pm 1.31	122.95 \pm 3.04	158.03 \pm 1.97	284.23 \pm 3.05	113.67 \pm 2.03	223.41 \pm 0.93
+ AugUndo	74.97 \pm 1.15	162.71 \pm 1.86	40.44 \pm 1.09	92.11 \pm 1.01	126.16 \pm 1.44	246.16 \pm 2.46	86.13 \pm 4.46	181.08 \pm 8.00
KBNet	78.44 \pm 1.39	178.17 \pm 3.27	37.56 \pm 0.61	83.43 \pm 1.89	149.13 \pm 3.29	306.30 \pm 8.74	70.74 \pm 2.26	136.75 \pm 5.55
+ AugUndo	66.97 \pm 0.81	151.55 \pm 2.03	31.63 \pm 0.53	71.90 \pm 0.82	117.16 \pm 4.51	239.60 \pm 10.96	57.65 \pm 1.47	112.81 \pm 1.75
DesNet	74.89 \pm 0.67	170.32 \pm 2.03	35.62 \pm 0.72	78.30 \pm 1.31	139.54 \pm 4.46	287.95 \pm 12.69	65.00 \pm 1.24	123.81 \pm 2.17
+ AugUndo	67.78 \pm 1.10	153.46 \pm 2.64	32.09 \pm 0.48	71.96 \pm 1.05	117.93 \pm 3.77	239.49 \pm 8.44	58.13 \pm 1.14	112.78 \pm 2.33
MonDi*	73.90 \pm 0.61	191.69 \pm 2.64	31.78 \pm 1.52	72.67 \pm 2.17	146.33 \pm 5.57	343.25 \pm 13.70	60.43 \pm 3.54	114.18 \pm 2.21
+ AugUndo	69.90 \pm 3.05	157.02 \pm 7.46	29.61 \pm 0.12	68.48 \pm 0.57	127.08 \pm 4.51	299.42 \pm 10.96	54.47 \pm 1.47	108.23 \pm 1.75

different scene layouts. As there are no viable intensity augmentations for sparse depth, RMP is the only one to explicitly increase variability in the data modality hence it also has large influence. We note that other geometric augmentations contribute as well, i.e., rotation, resizing, and translation, to discard points. By computing the loss on the original inputs, we reconstruct the removed points, which in fact serves as an additional training signal to map RGB to depth.

Photometric augmentations, individually, have small effect on performance. To get a non-negligible effect, one must disable all color jitter (Tab. 8, row 7), which still yielded small increases in errors. This shows the limitations of existing augmentations, which relies heavily on color jittering. We note that we use a larger value range in color jitter than existing works, so the impact is expected to be even smaller for existing works. Admittedly, the most important one (FLP) is currently being used by all methods, but resize, image patch and point removal, play a large role. The best results are obtained when using all of the proposed augmentations, demonstrating the importance of scaling up both photometric and geometric augmentations.

H Sensitivity Study for Depth Completion

Given that conventional augmentation pipelines are not applied towards sparse depth modality, it is possible that a model will overfit to the sparse point cloud, which describes the coarse 3D scene structure. Overfitting to scenes, hence, can

limit generalization and increase sensitivity to the configuration of sparse points. In the main paper, we presented results for depth completion on VOID1500 and VOID500. Here, we test the effect of AugUndo on various sparse depth input densities. We presented results on VOID500 (repeated for side-by-side comparison) and VOID150 in Tab. 9 (left and right, respectively). VOID500 contains approximately 500 sparse points per point cloud and VOID150 contains approximately 150 points. All models tested are trained on VOID1500, which contains 1500 sparse points.

For VOID500 (Tab. 9, left), which is a $3\times$ reduction in density, AugUndo improves the sensitivity to changes in the sparse points by an average of 19.66% across all methods, and 30.57%, 23.92%, 14.79%, 9.35% for VOICED, FusionNet, KBNet, and DesNet, respectively. For an even more challenging scenario, also the closest setup to the density of sparse points tracked by a Simultaneous Localization and Mapping (SLAM) and Visual Inertial Odometry (VIO) system, we consider VOID150 (Tab. 9, right). Here, AugUndo improves by an average of 19.62% across all methods, and 26.57%, 19.18%, 19.80%, and 12.95% for VOICED, FusionNet, KBNet, and DesNet, respectively. We attribute this to geometric and occlusion augmentations: translation, resize, rotation and sparse points removal. All of these augmentations not only affect photometry, but also the sparse point cloud where points are dropped due to resampling or explicitly removed, and additionally point cloud orientation is also altered.

We note that AugUndo is also helpful for distillation methods like MonDi [45]. As mentioned in the main text, we conjecture that artifacts caused from transformation of a piece-wise smooth depth map (in the case of supervised or distillation methods) are less severe than those of image and sparse point clouds. Hence, we were expecting the gains for supervised or distillation methods to be small compared to unsupervised methods. However, as shown in the last row of Tab. 9, we observe a surprisingly nontrivial gain when applying AugUndo to MonDi (marked by *). For VOID500, we improve MonDi by 9.02% on average across on metrics; for VOID150, we improve by 10.25%.

I Zero-shot Generalization from KITTI

In the main paper, we demonstrated AugUndo for three unsupervised depth completion (VOICED [84], FusionNet [82], and KBNet [87]) and three unsupervised monocular depth estimation (Monodepth2 [19], HR-Depth [47]), and Lite-Mono [102]) methods on the KITTI dataset. Due to space constraints, here, we provide additional results for zero-shot generalization from KITTI to Waymo Open Dataset [69] for depth completion and to Make3d [62] for monocular depth estimation. Similar to our generalization experiments on indoors (VOID to NYUv2 and ScanNet), we will train on KITTI using the conventional augmentation schemes employed by each respective method and compare the resulting models with those trained using AugUndo.

We begin by presenting results on depth completion in Tab. 10. Here, we evaluate VOICED, FusionNet, and KBNet models (trained on KITTI using standard

Table 10: *Zero-shot transfer from KITTI to Waymo for depth completion.* AugUndo improves generalization of models trained on KITTI to Waymo by an average of 13.3% for all evaluation metrics. We note that the sparse depth maps provided in Waymo is considerably denser than KITTI as they are merged from two front separate lidars; hence, FusionNet, which employs a learned (frozen) densification network performs similarly, i.e., the bias introduced by the same frozen densification network (ScaffNet) is strong enough that FusionNet yields similar results for both conventional augmentation scheme and AugUndo. Nonetheless, we still observe considerable improvements.

Dataset	Method	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓
Waymo	VOICED	6781.27±317.21	7734.57±339.77	24.58±1.45	27.87±1.63
	+ AugUndo	5965.07 ±367.47	7029.66 ±509.43	18.54 ±1.83	22.10 ±3.96
	FusionNet	530.55±39.23	1734.23±114.97	1.27±0.12	2.82±0.29
	+ AugUndo	512.29 ±8.43	1707.34 ±48.25	1.21 ±0.09	2.75 ±0.18
	KBNet	625.00±12.47	2167.74±92.17	1.76±0.13	5.46±0.93
	+ AugUndo	541.29 ±15.16	2014.14 ±76.52	1.34 ±0.11	3.43 ±0.30

Table 11: *Zero-shot transfer from KITTI to Make3d for monocular depth estimation.* All models are trained on KITTI. Note: for Monodepth2, we use the numbers reported by [19] and the best trial on KITTI.

Dataset	Method	MAE ↓	RMSE ↓	Abs Rel ↓	Sq Rel ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
Make3d	Monodepth2	-	7.417	0.322	3.589	-	-	-
	+ AugUndo	4.109	6.803	0.272	2.769	0.606	0.848	0.936
	HR-Depth	4.136	6.505	0.281	2.484	0.562	0.839	0.938
	+ AugUndo	4.023	6.428	0.272	2.393	0.584	0.848	0.94
	Lite-Mono	5.116	8.061	0.358	4.676	0.511	0.793	0.91
	+ AugUndo	4.728	7.518	0.321	3.887	0.529	0.817	0.926

augmentation pipelines and AugUndo) on the Waymo Open Dataset. Overall, training with AugUndo improves existing methods by an average 13.3% across all evaluation metrics. We note that AugUndo improves VOICED and KBNet by larger amounts than FusionNet. For FusionNet, whether trained with conventional augmentation scheme or AugUndo, both seem to perform similarly. This is because FusionNet employs a learning-based densification network (ScaffNet), which is pretrained on synthetic datasets and frozen, that is used in both models. As the sparse depth maps in Waymo are much denser than KITTI, ScaffNet is able to approximate the dense depth map with small amounts of errors. This serves as an inductive bias for the downstream FusionNet, which learns the residual over the approximated depth map. As the reconstruction from ScaffNet exhibits high fidelity, the bias induced by ScaffNet on FusionNet causes FusionNet to perform only minor modifications to approximated depth map, leading to similar outputs whether trained with conventional augmentations or AugUndo. Nonetheless, we still observe consistent (albeit smaller) improvements when FusionNet is trained with AugUndo.

For monocular depth estimation, we similarly evaluate Monodepth2, HR-Depth, and Lite-Mono models (trained on KITTI using standard augmentation pipelines and AugUndo) on Make3d. Tab. 11 shows that models trained on KITTI with AugUndo generalizes well to Make3d, gaining an average of around

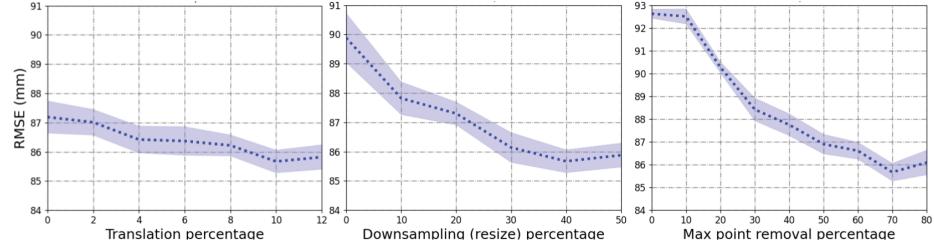


Fig. 5: Effect of different degree of augmentation on VOID1500.

8% improvement over all metrics and all models. Note: training Monodepth2 from their code repository reproduces their results on KITTI, but produces worse generalization results on Make3d than the weights they released; hence, we take the original weights released by the authors for evaluation. Nonetheless, we still improve over their best result.

J Sensitivity to Choice of Hyperparameters

In the main paper and Sec. A above, we described in details of the set of augmentations and degrees of each augmentation used in our experiments to achieve the reported numbers. We note that the performance gains can be obtained by typical set and ranges of augmentations and does not require a meticulous selection of hyper-parameters. In the paper, we tried to push the limits and tested increasing degrees of augmentation until performance saturated. Fig. 5 shows the trends of improvement for resizing (down), translation, and points removal. Choosing even small degrees of augmentations yields some performance gain. We note that there are diminishing returns as we scale up augmentations to large amounts, i.e. 50% or more in downsampling, which reduces the size of the observed image and predictions will have little to no detail. When computing the loss on these predictions on the original data, there is a considerable amount of noise and lends little to learning. Hence, at extreme degrees, we observe that performance eventually saturates. We also note that the choice of augmentation would depends on the task and dataset (i.e. vertical flip not applicable for KITTI since the camera is typical right-side-up for driving scenarios).

K Additional Results for Depth Completion

In the main paper, we demonstrated AugUndo for three unsupervised methods [82, 84, 87] on the VOID dataset. Due to space constraints, here, we show additional results on VOID1500 and VOID500 for a recent unsupervised distillation method, MonDi [45] as well as a recent unsupervised depth completion method, DesNet [91], in Tab. 12. Additionally, we show MonDi and DesNet for

VOID150 in Tab. 9. We applied photometric transformations of random brightness from 0.5 to 1.5, contrast from 0.5 to 1.5, saturation from 0.5 to 1.5, and hue from -0.1 to 0.1. We applied image patch removal by selecting between 0.1% to 0.5% of the pixels and removing 5×5 patches centered on them – approximately removing between 2.5% to 12.5% of the image. We applied random sparse depth point removal at a rate between 60% and 95% of all sparse points. We further applied geometric transformations of random horizontal and vertical flips, up to 10% of translation, between -25 to 25 degrees of rotation, and random resize factor between 0.6 to 1.1.

Table 12: Quantitative results of *MonDi*, a distillation-based unsupervised method, on VOID. Reported scores are mean and standard deviation over four independent trials. We evaluate *MonDi* trained on VOID1500, with standard augmentations and AugUndo, by testing them on VOID1500 and VOID500. Sparse depth maps within VOID1500 contains approximately 1500 points, and those within VOID500 contain approximately 3 \times less, respectively. AugUndo improves performance by an average of 5.33% and 9.02% across all methods and evaluation metrics on VOID1500 and VOID500, respectively.

Method	VOID1500				VOID500			
	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓
MonDi*	30.56 \pm 0.23	86.67 \pm 0.55	15.08 \pm 0.16	37.58 \pm 0.44	73.90 \pm 0.61	191.69 \pm 2.64	31.78 \pm 1.52	72.67 \pm 2.17
+ AugUndo	29.02\pm0.56	79.34\pm1.17	14.56\pm0.44	35.93\pm1.12	69.90\pm3.05	157.02\pm7.46	29.61\pm0.12	68.48\pm0.57
DesNet	37.41 \pm 0.28	93.31 \pm 0.76	19.17 \pm 0.24	45.57 \pm 0.62	74.89 \pm 0.67	170.32 \pm 2.03	35.62 \pm 0.72	78.30 \pm 1.31
+ AugUndo	33.86\pm0.60	86.05\pm0.43	16.92\pm0.29	41.25\pm0.31	67.78\pm1.10	153.46\pm2.64	32.09\pm0.48	71.96\pm1.05

Results on MonDi and DesNet. Overall, we improve *MonDi* by an average of 8.2% across all evaluation metrics across all sparsity levels (VOID1500, VOID500, and VOID150). This is surprising as amongst all the tested methods, *MonDi* is the most light-weight, with only 5.3M parameters. As we reduce model capacity, one would expect that the network to saturate in the data variations that can be modeled. However, despite the size of network is much smaller (23.2% less than KBNet), there is still a considerable gain when using AugUndo instead of their augmentation pipeline. This demonstrates efficacy and applicability of AugUndo; it can be used to improve methods with a range of capacities from tens of millions to several million. Also, we note that *MonDi* is an unsupervised distillation method (where it distills from unsupervised methods), so it is more closely related to supervised methods in supervision than unsupervised method. Even so, we observe a non-trivial improvement, which validates our discussion in Sec. 5 of the main paper regarding the applicability of AugUndo to *supervised* methods. Meanwhile, we improves *DesNet* by an average of 10.64% accross all evaluation metrics, highlighting again AugUndo’s applicability to future unsupervised depth completion methods.

Modeling AugUndo as change in camera pose versus camera parameters. In the main text, we discussed two ways of modeling our augmentation scheme: treating the augmented data as a result of changes in camera pose (i.e. motion) or

Table 13: Quantitative results of KBNet with AugUndo as changes in camera pose or intrinsics (i.e. with and without depth adjustment, respectively) on VOID1500. With depth adjustments (change in camera pose), the performance of the model is slightly worse ($\approx 3.3\%$ in terms of percent gain) than without depth adjustments (change in camera intrinsics), yet still better than the baseline by a large margin.

Method	MAE \downarrow	RMSE \downarrow	iMAE \downarrow	iRMSE \downarrow	Gain (%)
KBNet	38.11 \pm 0.77	95.22 \pm 1.72	19.51 \pm 0.14	46.70 \pm 0.48	-
+ AugUndo (intrinsics)	33.32 \pm 0.18	85.67 \pm 0.39	16.61 \pm 0.29	41.24 \pm 0.60	+12.29
+ AugUndo (pose)	34.24 \pm 0.25	87.11 \pm 0.56	17.59 \pm 0.23	43.22 \pm 0.39	+8.99

in camera parameters (i.e. intrinsics). (1) Modeling the augmentation as camera motion requires adjusting sparse depth maps; for example, resizing can be treated as forward motion, so distance from camera to world surfaces need to be adjusted in the sparse depth map. Naturally, if choosing (1), then one would need to reproject the sparse depth points according to the camera motion. On the other hand, (2) modeling the augmentation as changes in the camera parameters does not require adjusting the sparse depth maps as the camera and 3D scene are both static; for example, resizing can now be treated as “zooming in” or an increase in focal length.

Tab. 13 compares the two approaches. We test the resizing operation by adjusting sparse depth, and likewise, dense output depth. Specifically, we assume a pin-hole camera model. As perspective projection is a linear, we adjust the depth value by scaling them using the random scale factor recorded during the random resizing operation. We observe the following: Firstly, both methods of modeling improves KBNet on VOID: modeling as (1) improves KBNet by 8.99% and (2) by 12.29%, respectively. This verifies the efficacy of AugUndo under both modeling choice. Secondly, modeling AugUndo as changes in (2) camera parameters improves over (1) camera motion. Particularly, (2) yields an additional 3.3% gain in the baseline and 3.6% relative improvement over KBNet trained using (1). This justifies our choice of modeling AugUndo as changes in camera parameters.

Naive geometric augmentations. In the main paper, we discussed that naively applying geometric augmentations can be detrimental to model performance and even prevent one from training them. As unsupervised depth completion and unsupervised monocular depth estimation assume rigid motion within the image triplet comprising of a training example, geometric augmentations that introduce some form of border padding (e.g., translation, rotation) will yield constant or edge extended borders across images (i.e., no motion in those regions). The lack of motion will result in PoseNet predicting near identity pose. Hence, naively incorporating geometric augmentations will prevent the model from properly learning depth and pose. This is demonstrated in Tab. 14 and 15 for rows marked with “+ Naive Geo Aug”. Nonetheless, one can make modifications (no translation, center cropping on rotation, and resizing to the same shape for a batch) to ensure no borders are introduced during augmentation.

Table 14: *Naive geometric augmentations for unsupervised depth completion on VOID.*

It is infeasible to conduct unsupervised training with naive geometric augmentations. Modifying geometric augmentations to ensure no borders are introduced allows the model to train, but performance degrades. “+ Naive Geo Aug” denotes models trained with naive geometric augmentations and “+ Modified Geo Aug” denotes models trained with modified geometric augmentations.

Method	MAE ↓	RMSE ↓	iMAE ↓	iRMSE ↓
VOICED	74.78±2.69	139.75±4.57	39.20±1.46	71.98±2.54
+ Naive Geo Aug	554.65±91.10	642.13±69.40	639.42±113.11	980.81±78.01
+ Modified Geo Aug	109.40±13.01	205.14±18.22	106.21±13.27	286.71±31.79
+ AugUndo	52.73±0.41	111.09±0.92	26.93±0.54	54.46±0.38
FusionNet	52.11±0.44	113.30±1.18	28.53±0.52	58.79±2.01
+ Naive Geo Aug	82.73±6.96	190.71±9.90	73.61±33.32	224.62±156.04
+ Modified Geo Aug	57.72±6.67	138.52±15.87	31.91±5.19	74.41±14.64
+ AugUndo	41.16±0.18	99.21±0.39	22.23±0.35	53.07±1.30
KBNet	38.11±0.77	95.22±1.72	19.51±0.14	46.70±0.48
+ Naive Geo Aug	1,065.4 ± 81.7	1,216.4 ± 54.4	5,693.5 ± 401.3	7,032.7 ± 283.0
+ Modified Geo Aug	51.77 ± 4.13	118.95 ± 7.67	27.66 ± 3.04	60.92 ± 5.62
+ AugUndo	33.32±0.18	85.67±0.39	16.61±0.29	41.24±0.60

Table 15: *Naive geometric augmentations for unsupervised monocular depth estimation on KITTI.* It is infeasible to conduct unsupervised training when naively applying geometric augmentations. Modifying geometric augmentations to ensure no borders are introduced allows the model to train, but performance degrades. “+ Naive Geo Aug” denotes models trained with naive geometric augmentations and “+ Modified Geo Aug” denotes models trained with modified geometric augmentations.

Method	RMSE ↓	Abs Rel ↓	Sq Rel ↓	$\delta < 1.25 \uparrow$	$\delta < 1.25^2 \uparrow$	$\delta < 1.25^3 \uparrow$
Monodepth2	4.794 ± 0.035	0.117±0.001	0.845± 0.030	0.869± 0.004	0.959± 0.001	0.982±0.001
+ Naive Geo Aug	6.612±0.18	0.174±0.009	2.057±0.32	0.794±0.02	0.932±0.004	0.971±0.002
+ Modified Geo Aug	4.911±0.046	0.119±0.002	0.926±0.023	0.870±0.004	0.958±0.001	0.981±0.001
+ AugUndo	4.739±0.032	0.113±0.000	0.862±0.030	0.879±0.002	0.960±0.001	0.982±0.001
HR-Depth	4.626 ± 0.032	0.113±0.001	0.797± 0.022	0.879±0.002	0.961± 0.001	0.982±0.000
+ Naive Geo Aug	6.06±0.022	0.159±0.003	1.40±0.016	0.795±0.010	0.939±0.003	0.979±0.002
+ Modified Geo Aug	4.718±0.018	0.116±0.001	0.844±0.005	0.875±0.001	0.960±0.001	0.982±0.001
+ AugUndo	4.610±0.029	0.111±0.001	0.794±0.021	0.883±0.001	0.962±0.001	0.983±0.001

Yet, we do not perform the proposed “undo-ing” process. While this will allow the model to train, however, the effects of artifacts, loss during resampling, color distortion, and intensity saturation will lead to performance degradations worse than the baseline. This is shown in Tab. 14 and 15 for rows marked with “+ Modified Geo Aug”.

L Additional Results on Monocular Depth Estimation

Applying our methods yields qualitative improvements in the depth prediction. In Fig. 6, we can see that our method better captures the building in the last

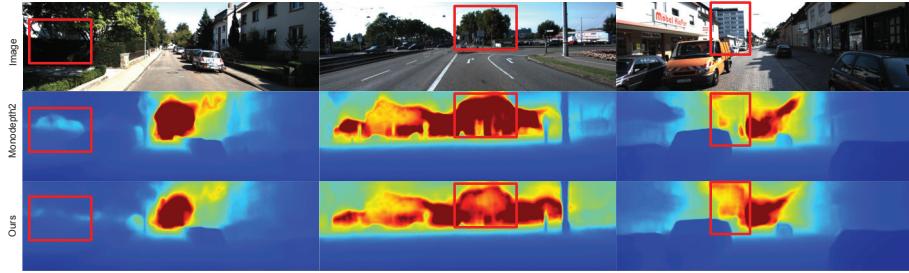


Fig. 6: Qualitative result of MonoDepth2 on KITTI. Red bounding boxes highlight areas where training with our augmentation scheme improves MonoDepth2, i.e., wall and vegetation on left, trees in middle and building on right.

column. Similarly, we observe similar improvements in the middle column for the trees that are located at a distance. This is in part thanks to augmentations such as resizing so that we can simulate objects at far and close distances. Additionally, training with AugUndo also improves over the baseline method on ambiguous regions such as the wall and vegetation in the first column, despite using the same hyper-parameter during training with the exception of data augmentation.

M Limitations

While we have proposed an algorithm to scale up augmentations for depth completion, a multimodal 3D reconstruction problem, admittedly our augmentations are limited to 2D. In the case of this inverse problem, the nuisance variability simulated by AugUndo are, in fact, projections of those in the 3D scene, which we do not model directly. Additionally, augmentations are used in other vision tasks including multiview stereo, binocular stereo, optical flow, etc. While we have shown that AugUndo can be applied to unsupervised and distillation (with dense supervision from pseudo ground truth), we have limited the scope of our method for depth completion, which considers a single image and sparse depth map as input. We foresee that AugUndo can also incorporate other modalities, such as tactile [92], be extended to other settings [50], and tasks, including deep feature visualization [37], semantic segmentation [3, 22, 38, 43, 67, 80], and object detection [2, 33, 35, 61, 70]. Additionally, we hypothesize that AugUndo can also be extended towards multi-frame geometric tasks such as stereo [1, 78, 85, 89], optical flow [34, 36, 68, 72, 103, 104], etc., but one must account for their specifics and problem setups, i.e. stereo assumes frontoparallel views. Lastly, like all scaling problems, AugUndo is eventually limited by diminishing returns. As certain augmentations are pushed to extremes, i.e., maximum brightness such that the image is “white”, large spatial reduction such that the image is small, there is little to no information in the input to infer depth; even if the supervision signal exists, it would be mapping nonsensical inputs to depth maps, which does

not improve performance (as illustrated by Fig. 5), and lending to saturation in performance gain.