

# Robust Semi-Supervised Monocular Depth Estimation with Reprojected Distances

Vitor Guizilini Jie Li Rareş Ambruş Sudeep Pillai Adrien Gaidon  
Toyota Research Institute (TRI)  
firstname.lastname@tri.global

**Abstract:** Dense depth estimation from a single image is a key problem in computer vision, with exciting applications in a multitude of robotic tasks. Initially viewed as a direct regression problem, requiring annotated labels as supervision at training time, in the past few years a substantial amount of work has been done in self-supervised depth training based on strong geometric cues, both from stereo cameras and more recently from monocular video sequences. In this paper we investigate how these two approaches (supervised & self-supervised) can be effectively combined, so that a depth model can learn to encode true scale from sparse supervision while achieving high fidelity local accuracy by leveraging geometric cues. To this end, we propose a novel supervised loss term that complements the widely used photometric loss, and show how it can be used to train robust semi-supervised monocular depth estimation models. Furthermore, we evaluate how much supervision is actually necessary to train accurate scale-aware monocular depth models, showing that with our proposed framework, very sparse LiDAR information, with as few as 4 beams (less than 100 valid depth values per image), is enough to achieve results competitive with the current state-of-the-art.

**Keywords:** Structure from Motion, Semi-Supervised Learning, Deep Learning, Depth Estimation, Computer Vision

## 1 Introduction

Depth perception is an essential component of any autonomous agent, enabling it to interact with objects and properly react to its surrounding environment. While there are sensors capable of providing such information directly, estimating depth from monocular images is particularly appealing, since cameras are inexpensive, compact, with low power consumption and capable of providing dense textured information. Recent breakthroughs in learning-based algorithms have allowed the generation of increasingly accurate monocular depth models, however these come with their own shortcomings. *Supervised* methods [1] require additional sensors with precise cross-calibration to provide depth labels and normally do not generalize to non-supervised areas, while *self-supervised* methods are limited by the accuracy of stereo [2, 3, 4] or structure-from-motion [5, 6, 7, 8, 9] reconstruction, with the latter also suffering from scale ambiguity. Nevertheless, *supervised* methods still produce the highest accuracy models, especially when high quality groundtruth information is available, while *self-supervised* methods are highly scalable, being capable of consuming massive amounts of unlabeled data to produce more generic models. Ideally, both approaches should be able to work together and complement each other in order to achieve the best possible solution, given the data that is available at training time.

Following this intuition, the main contribution of this paper is a novel supervised loss term that minimizes reprojected distances in the image space, and therefore operates under the same conditions as the photometric loss [10], which constitutes the basis for appearance-based self-supervised monocular depth learning methods. We show that this novel loss not only facilitates the injection of depth labels into self-supervised models, to produce *scale-aware* estimates, but it also further improves the quality of these estimates, even in the presence of very sparse labels. The result is a novel semi-supervised training methodology that combines the best of both worlds, being able to consume massive amounts of unlabeled data, in the form of raw video sequences, while also prop-

erly exploiting the information contained in depth labels when they are available. The ability to properly leverage sparse information also greatly facilitates the generation of depth labels, eliminating the need for expensive 64 or 128-beam LiDAR sensors in favor of cheaper and easily available alternatives as the source of supervision.

## 2 Related Work

### 2.1 Supervised Methods

Supervised methods use ground truth depth, usually from LiDAR in outdoor scenes, to train a neural network as a regression model. Eigen et al. [11] was amongst the first to propose convolutional neural networks as a solution to this problem, generating initially a coarse prediction and then refining it using another neural network to produce more accurate results. Since then, substantial work has been done to improve the accuracy of supervised depth estimation from monocular images, including the use of Conditional Random Fields (CRFs) [12], the inverse Huber distance loss function [13], joint optimization of surface normals [14], fusion of multiple depth maps [15] and more recently its formulation as an ordinal classification problem [1]. LiDAR data is also sparse relative to the camera field of view, and consequently supervised methods are unable to produce meaningful depth estimates in non-overlapping areas of the image.

### 2.2 Self-Supervised Methods

As supervised techniques for depth estimation advanced rapidly, the availability of target depth labels became challenging, especially for outdoor applications. To this end, [5, 16] provided an alternative strategy involving training a monocular depth network with stereo cameras, without requiring ground-truth depth labels. By leveraging Spatial Transformer Networks [17], Godard et al [16] use stereo imagery to geometrically transform the right image plus a predicted depth of the left image into a synthesized left image. The loss between the resulting synthesized and original left images is then defined in a fully-differentiable manner, using a Structural Similarity [10] term and additional depth regularization terms, thus allowing the depth network to be self-supervised in an end-to-end fashion. Following [16] and [18], Zhou et al. [6] generalize this to self-supervised training in the *purely* monocular setting, where a depth and pose network are simultaneously learned from unlabeled monocular videos. Several methods [19, 7, 20, 21, 22, 23, 24] have since then advanced this line of work by incorporating additional loss terms and constraints.

### 2.3 Semi-Supervised Methods

Unlike both of the above mentioned categories, there has not been much work on semi-supervised depth estimation. Most notably, Guo et al. [25] and Luo et al. [26] introduced multiple sequential self-supervised and supervised training stages, where each stage is conducted independently. Kuznetsov et al. [27] proposed adding the supervised and self-supervised loss terms together, allowing both sources of information to be used simultaneously. The same concept was applied in [28], with the introduction of left-right consistency to avoid post-processing at inference time. Our work follows a similar direction, focusing on how to properly incorporate depth labels into appearance-based self-supervised learning algorithms at training time, so we can both produce scale-aware models and further improve on the quality of depth estimates, even in the presence of very sparse labels.

## 3 Methodology

Our proposed semi-supervised learning methodology is composed of two training stages, as depicted in Figure 1. The first consists of a self-supervised monocular structure-from-motion setting, where we aim to learn: (i) a monocular depth model  $f_D : I \rightarrow D$ , that predicts the scale-ambiguous depth  $\hat{D} = f_D(I(p_t))$  for every pixel  $p_t$  in the target image  $I_t$ ; and (ii) a monocular ego-motion model  $f_x : (I_t, I_S) \rightarrow \mathbf{x}_{t \rightarrow S}$ , that predicts the set of 6-DoF rigid transformations for all  $s \in S$  given by  $\mathbf{x}_{t \rightarrow s} = \begin{pmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0} & 1 \end{pmatrix} \in \text{SE}(3)$ , between the target image  $I_t$  and the set of source images  $I_s \in S$  considered as part of the temporal context. Afterwards, if supervision is available, this information can be used to (i) collapse the scale ambiguity inherent to a single camera configuration into a metrically accurate model, and (ii) further improve the depth and ego-motion models by leveraging

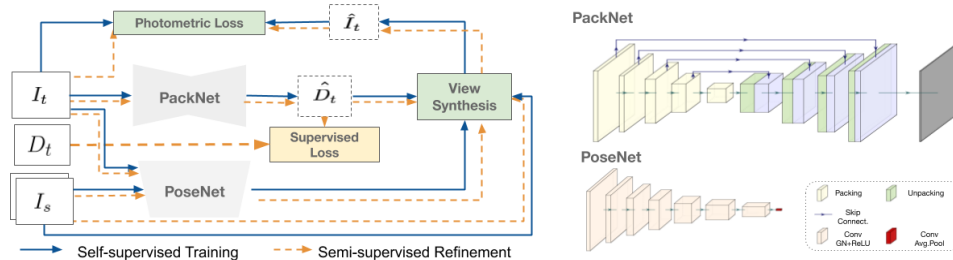


Figure 1: **Diagram of the proposed semi-supervised monocular depth estimation framework.**

cues that are not appearance-based. More details on network architectures and objective function terms are given in the following sections.

### 3.1 Depth Network

Our depth network is based on the architecture introduced by Guizilini et al. [9], that proposes novel packing and unpacking blocks to respectively downsample and upsample feature maps during the encoding and decoding stages (Fig. 1). Skip connections [29] are used to facilitate the flow of information and gradients throughout the network. The decoder produces intermediate inverse depth maps, that are upsampled before being concatenated with their corresponding skip connection and unpacked feature maps. They also serve as output to produce a 4-scale inverse depth pyramid, from which the loss is calculated. However, instead of incrementally super-resolving each inverse depth map, as described in [9], here they are all upsampled directly to the highest resolution using bilinear interpolation. As noted in [30], this reduces copy-based artifacts and photometric ambiguity, leading to better results as shown in experiments.

### 3.2 Pose Network

For camera ego-motion estimation, we use the neural network architecture proposed by [6] *without* the explainability mask, which we found not to improve results. Following [6], the pose network consists of 7 convolutional layers followed by a final  $1 \times 1$  convolutional layer and a 6-channel average pooling operation (Fig. 1). The input to the network consists of a target  $I_t$  and context  $I_s$  images, concatenated together, and the output is the set of 6 DoF transformations between  $I_t$  and  $I_s$  that constitute  $T_{s \rightarrow t}$ . If more than one contextual image is considered, this process is repeated for each  $I_s \in \mathcal{S}$  to produce independent transformations. Translation is parametrized in Euclidean coordinates  $\{x, y, z\}$  and rotation uses Euler angles  $\{\alpha, \beta, \gamma\}$ .

### 3.3 Objective Function

The objective function used in our work has two components: a *self-supervised* term, that operates on appearance matching  $\mathcal{L}_{photo}$  between the target  $I_t$  and synthesized images  $I_{s \rightarrow t}$  from the context set  $\mathcal{S} = \{I_s\}_{s=1}^S$ , with masking  $M_{photo}$  and depth smoothness  $\mathcal{L}_{smooth}$ ; and a *supervised* term, that operates on the reprojected  $\mathcal{L}_{rep}$  distances between predicted and ground-truth depth values. The coefficients  $\lambda$  are responsible for weighting the various terms relative to the photometric loss. The full objective function is as follows:

$$\mathcal{L}(I_t, \mathcal{S}) = \mathcal{L}_{photo} \odot M_{photo} + \lambda_{smooth} \cdot \mathcal{L}_{smooth} + \lambda_{rep} \cdot \mathcal{L}_{rep} \quad (1)$$

**Appearance Matching Loss.** Following [6] the pixel-level similarity between the target image  $I_t$  and the synthesized target image  $I_{t \rightarrow s}$  is estimated using the Structural Similarity (SSIM) [10] term combined with an L1 pixel-wise loss term, inducing an overall photometric loss given by:

$$\mathcal{L}_{photo}(I_t, I_{s \rightarrow t}) = \alpha \frac{1 - \text{SSIM}(I_t, I_{s \rightarrow t})}{2} + (1 - \alpha) \|I_t - I_{s \rightarrow t}\| \quad (2)$$

While multi-view projective geometry provides strong cues for self-supervision, errors due to parallax and out-of-bounds objects have an undesirable effect incurred on the photometric loss, that adds noise to the training stage and should not be learned by the depth model. Following Godard et al.

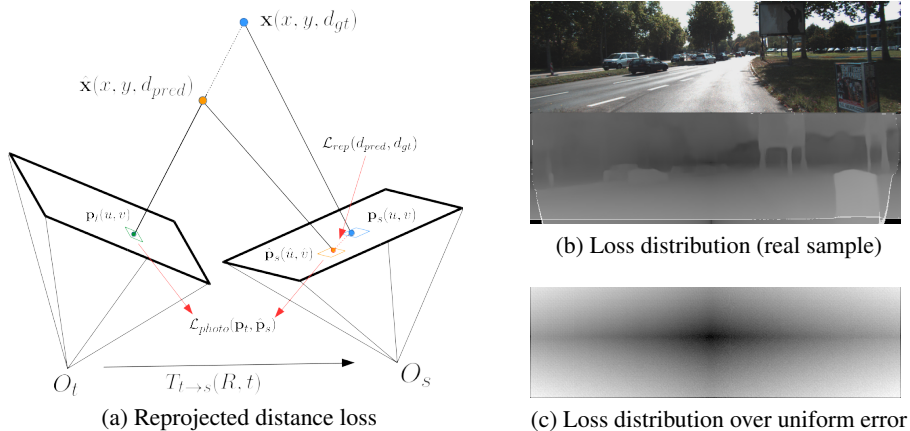


Figure 2: **Diagram of the proposed reprojected distance loss** and its behavior relative to  $\mathbf{p}_t$  pixel coordinates. In (b), the bottom row depicts the per-pixel reprojected loss values when there is a constant error of 1m relative to the predicted depth values from the top row image. Similarly, (c) depicts the per-pixel loss values when the same 1m error is applied to a constant depth of 50m throughout the entire image.

[30], we mitigate these undesirable effects by calculating the minimum photometric loss per pixel for each source image in the context  $\mathcal{S}$ , so that:

$$\mathcal{L}_{photo}(I_t, \mathcal{S}) = \min_{s \in \mathcal{S}} \mathcal{L}_{photo}(I_t, I_{s \rightarrow t}) \quad (3)$$

The intuition is that the same pixel will not be occluded or out-of-bounds in all context images, and that the association with minimal photometric loss should be the correct one. Furthermore, also following [30] we mask out static pixels by removing those which have a *warped* photometric loss  $\mathcal{L}_{photo}(I_t, I_{t \rightarrow s})$  higher than their corresponding *unwarped* photometric loss  $\mathcal{L}_{photo}(I_t, I_s)$ , calculated using the original source image without view synthesis. This mask removes pixels whose appearance does not change between frames, which includes static scenes and dynamic objects moving at a similar speed as the camera, since these will have a smaller photometric loss when we assume no ego-motion.

$$M_{photo} = \min_{s \in \mathcal{S}} \mathcal{L}_{photo}(I_t, I_s) > \min_{s \in \mathcal{S}} \mathcal{L}_{photo}(I_t, I_{s \rightarrow t}) \quad (4)$$

**Depth Smoothness Loss.** In order to regularize the depth in texture-less low-image gradient regions, we incorporate an edge-aware term, similar to [16]. The loss is weighted for each of the pyramid levels, starting from 1 and decaying by a factor of 2 on each scale:

$$\mathcal{L}_{smooth}(\hat{D}_t) = |\delta_x \hat{D}_t| e^{-|\delta_x I_t|} + |\delta_y \hat{D}_t| e^{-|\delta_y I_t|} \quad (5)$$

**Reprojected Distance Loss.** Most supervised depth training methods operate on a direct regression basis, mapping input RGB images into output depth values without exploiting motion information or camera geometry. However, if such information is available — as is the case in self-supervised learning —, we show here that it can be leveraged to produce more accurate and consistent depth models. This is achieved using a novel supervised loss term (Eq. 6) that operates under the same conditions as the photometric loss, by reprojecting depth errors back onto the image space as observed by a context camera, as shown in Fig. 2a. For each pixel  $\mathbf{p}_t^i$  in the target image, this reprojected depth error  $\mathcal{L}_{rep}$  corresponds to the distance between its true  $\mathbf{p}_s^i$  and predicted  $\hat{\mathbf{p}}_s^i$  associations in the source image, from which the photometric loss  $\mathcal{L}_{photo}$  is calculated:

$$\begin{aligned} \mathcal{L}_{rep}(\hat{D}_t, D_t) &= \frac{1}{V} \sum_{i \in \mathcal{V}_t} \|\hat{\mathbf{p}}_s^i - \mathbf{p}_s^i\| = \frac{1}{V} \sum_{i \in \mathcal{V}_t} \|\pi_s(\hat{\mathbf{x}}_t^i) - \pi_s(\mathbf{x}_t^i)\| \\ &= \frac{1}{V} \sum_{i \in \mathcal{V}_t} \|\pi_s(\hat{d}_t^i K^{-1} \mathbf{u}_t^i) - \pi_s(d_t^i K^{-1} \mathbf{u}_t^i)\| \end{aligned} \quad (6)$$

where  $\mathbf{u}_t^i = (u, v, 1)_s^{i,T}$  denotes the homogeneous coordinates of pixel  $i$  in target image  $I_t$ , and  $\hat{\mathbf{x}}_t^i$  and  $\mathbf{x}_t^i$  are the homogeneous coordinates of its reconstructed 3D points given respectively the predicted  $\hat{d}^i$  and ground truth  $d^i$  depths values.  $\hat{\mathbf{p}}_s^i = (\hat{u}, \hat{v})_s^{i,T}$  and  $\mathbf{p}_s^i = (u, v)_s^{i,T}$  denote respectively the 2D projected pixel coordinates of points  $\hat{\mathbf{x}}_t^i$  and  $\mathbf{x}_t^i$  onto source frame  $I_s$ , produced by the project function  $\pi_t = \pi(\mathbf{x}; K, T_{t \rightarrow s})$ . The effects of operating in this reprojected space are shown in Figs. 2b and 2c, where we respectively see how depth ranges are weighted differently, and so are pixel coordinates closer to the vanishing point. The concept of depth weighting has been explored before, such as with the inverse Huber loss [13] and spacing-increasing discretization [1], however here this weighting is not artificially introduced, but rather comes as a direct by-product of camera geometry, with errors being proportional to their respective projections onto the image space. Similarly, weighting based on distance to the vanishing point directly correlates to the baseline used in structure-from-motion calculation. If there is no relative baseline (i.e. it coincides with the error vector  $\mathbf{e}_t^i = \mathbf{x}_t^i - \hat{\mathbf{x}}_t^i$ ), this is an ill-defined problem, and therefore no depth information can be produced by that particular configuration.

Note that this proposed loss term is not appearance-based, and therefore in theory any transformation matrix  $T$  could be used to produce the reprojections from which the distance is minimized. However, by enforcing  $T = T_{s \rightarrow t}$  we can (a) back-propagate through the pose network, so it can also be directly updated with label information and remain consistent with the depth network; and (b) operate on the same reprojected distances that are used by the photometric loss, only on a *scale-aware* capacity, so its inherent ambiguity is forced to collapse into a metrically accurate model.

## 4 Experimental Results

### 4.1 Datasets

We use the popular KITTI [31] dataset for all experiments, to facilitate comparisons with other related methods. The evaluation method introduced by Eigen [11] for KITTI uses raw reprojected LIDAR points, and does not handle occlusions, dynamic objects or ego-motion, which leads to wrong image reprojections. A new set of high quality depth maps for KITTI was introduced in [32], making use of 5 consecutive frames to increase the number of available information and handling moving objects using the stereo pair. Throughout our experiments, we refer to this new set as *Annotated*, while the original set is referred to as *Original*. For training, we use the pre-processing methodology described in [6] to remove static frames, resulting in 39810 training images from both left and right cameras (note that no stereo information is used in this work) and 652 evaluation images. Context information for monocular depth and pose learning is generated using the immediate previous  $t - 1$  and posterior  $t + 1$  frames.

### 4.2 Implementation Details

We use PyTorch [33] for all our experiments<sup>1</sup>, training across 8 distributed Titan V100 GPUs. We use the Adam optimizer [34], with  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . When training a model from scratch, we follow Guizilini et al. [9] and optimize the depth and pose networks for 100 epochs, with a batch size of 4 and initial depth and pose learning rates of  $2 \cdot 10^{-4}$ , that are halved every 40 epochs. We set the SSIM weight to  $\alpha = 0.85$  (Eq. 2), the depth smoothness weight to  $\lambda_{smooth} = 10^{-3}$  (Eq. 1) and, when applicable, the reprojected distance weight to  $\mathcal{L}_{rep} = 10^4$  (Eq. 1). When refining a model only 15 epochs are considered, with depth and pose learning rates of  $10^{-4}$  that are halved every 6 epochs, and the same weight values for the loss function are used. Input images are downsampled to a  $640 \times 192$  resolution, and horizontal flipping and jittering are used as data augmentation.

### 4.3 Depth Estimation

In this section we present and discuss our results in monocular depth estimation, and how they compare with similar works found in the literature. An ablative study of the different components of our proposed semi-supervised training methodology can be found in Table 1, starting with quantitative evidence that our proposed modifications to the losses used by the original PackNet-SfM framework [9] are able to improve self-supervised depth training, resulting in significantly more accurate

<sup>1</sup>Inference code and pretrained weights for our models are available upon request.



Method	Supervision	MS	RF	Loss	Abs.Rel	Sq.Rel	RMSE	RMSE <sub>log</sub>	$\delta < 1.25$
PackNet-SfM [9]	Self	✓			0.086	0.460	3.712	0.132	0.918
<b>Ours (baseline)</b>	Self	✓			0.078	0.417	3.487	0.121	0.931
<b>Ours</b>	Semi	✓	✓	Rep	<b>0.069</b>	<b>0.313</b>	<b>3.097</b>	<b>0.110</b>	<b>0.941</b>
<b>Ours</b>	Sup		✓	L1	0.078	0.378	3.330	0.121	0.927
<b>Ours</b>	Semi			L1	0.084	0.437	3.68	0.133	0.913
<b>Ours</b>	Semi		✓	L1	0.078	0.383	3.340	0.121	0.927
<b>Ours</b>	Semi			Rep	0.077	0.362	3.274	0.119	0.929
<b>Ours</b>	Semi		✓	Rep	<b>0.072</b>	<b>0.340</b>	<b>3.265</b>	<b>0.116</b>	<b>0.934</b>

Table 1: **Ablation study of our proposed framework**, on the KITTI dataset with  $640 \times 192$  input resolution and *Annotated* depth maps when applicable. *MS* indicates the use of ground-truth median-scaling at inference time. *Self*, *Sup* and *Semi* respectively indicate self-supervised, supervised and semi-supervised training. *RF* indicates the use of **Ours (baseline)** as a pretrained model for refinement. *L1* and *Rep* respectively indicate supervision using the L1 loss (absolute depth distance) and the proposed reprojected distance loss (Sec. 3.3).

models. These results are compared against other techniques in Table 2 (*Original* depth maps), and constitute a new state-of-the-art in self-supervised depth learning, surpassing even stereo methods. Given this high-quality unscaled model, we proceed to show how depth labels can be efficiently incorporated at training time to further improve results. Initially, we train two supervised models with the traditionally used L1 loss (absolute depth distance), both starting from scratch and from our pretrained self-supervised checkpoint. We then introduce the semi-supervision methodology described in this paper, by using both the photometric and L1 losses during the refinement process from the same pretrained self-supervised checkpoint. Finally, we remove the L1 loss and introduce our proposed reprojected distance loss, training again in a semi-supervised manner two models, one from scratch and another from the pretrained self-supervised checkpoint.

Interestingly, we can see from these results that the introduction of depth labels for semi-supervision, when using the L1 loss, mostly enabled the model to learn scale, however it was unable to further improve the overall quality of depth estimates (i.e. going from 0.078 with median-scaling to 0.078 without median-scaling). In contrast, semi-supervision with our proposed reprojected distance loss effectively improved the accuracy of the refined depth model, going from 0.078 with median-scaling to 0.072 without median-scaling. Furthermore, we show in Fig. 3 that our proposed loss term also enables the learning of better estimates in areas where there is no LiDAR supervision, such as the upper portions of the image. We attribute this behavior to the fact that the reprojected distance loss operates in the same image space as the photometric loss, and thus is better suited to address its




Figure 3: **Example of depth maps** produced by an unscaled self-supervised model, and after its refinement with different supervised losses (the colormaps used for plotting are produced to be scale agnostic). As expected, our proposed Reprojected Distance loss is able to better reconstruct areas not observed by the LiDAR sensor used as source of supervision. For comparison, we also show results using the current state-of-the-art supervised method from [1], which also fails to properly reconstruct areas not observed by the LiDAR sensor at training time.

	Method	Supervision	Resolution	Dataset	Abs.Rel	Sq.Rel	RMSE	RMSE <sub>log</sub>	$\delta < 1.25$
Original	Monodepth2. [30] <sup>‡</sup>	Self (M)	640 x 192	K	0.129	1.112	5.180	0.205	0.851
	PackNet-SfM [9]	Self (M)	640 x 192	K	0.120	1.018	5.136	0.198	0.865
	Monodepth2 [30] <sup>‡</sup>	Self (S)	640 x 192	K	0.115	1.010	5.164	0.212	0.858
	Monodepth2 [30] <sup>‡</sup>	Self (M)	640 x 192	K	0.115	0.903	4.863	0.193	0.877
	SuperDepth [2]	Self (S)	1024 x 384	K	0.112	0.875	4.958	0.207	0.852
	<b>Ours</b>	Self (M)	640 x 192	K	<b>0.111</b>	<b>0.785</b>	<b>4.601</b>	<b>0.189</b>	<b>0.878</b>
Annotated	Kuznietsov et al. [27] <sup>‡</sup>	Semi (S)	640 x 192	K	0.089	0.478	3.610	0.138	0.906
	SemiDepth [28]	Semi (S)	640 x 192	C+K	0.078	0.417	3.464	0.126	0.923
	SVSM FT [26] <sup>‡</sup>	Semi (S)	640 x 192	F+K	0.077	0.392	3.569	0.127	0.919
	DORN (VGG) [1] <sup>‡</sup>	Sup	640 x 192	K	0.081	0.376	3.056	0.132	0.915
	DORN (ResNet) [1] <sup>‡</sup>	Sup	640 x 192	K	<b>0.072</b>	<b>0.307</b>	<b>2.727</b>	0.120	0.930
	<b>Ours</b>	Semi (M)	640 x 192	K	<b>0.072</b>	0.340	3.265	<b>0.116</b>	<b>0.934</b>
	<b>Ours (64 beams)</b>	Semi (M)	640 x 192	K	0.074	0.355	3.349	0.118	0.930
	<b>Ours (32 beams)</b>	Semi (M)	640 x 192	K	0.076	0.363	3.361	0.121	0.929
	<b>Ours (16 beams)</b>	Semi (M)	640 x 192	K	0.078	0.371	3.388	0.122	0.928
	<b>Ours (8 beams)</b>	Semi (M)	640 x 192	K	0.078	0.395	3.587	0.126	0.922
	<b>Ours (4 beams)</b>	Semi (M)	640 x 192	K	0.082	0.424	3.732	0.131	0.917

Table 2: **Quantitative comparison** between different learning-based monocular depth estimation techniques. *Self*, *Sup* and *Semi* respectively indicate self-supervised, supervised and semi-supervised training, with (M) and (S) respectively indicating monocular and stereo self-supervision. *F*, *C* and *K* indicate respectively the FlyingThings3D, Cityscapes and KITTI datasets. † indicates ImageNet [35] pretraining on the depth and/or pose networks. All *Mono* methods use median-scaling at inference time, to produce metrically accurate estimates.

inherent scale ambiguity, consistently collapsing the entire image to a metrically accurate model. We also compare these results with other similar techniques found in the literature (Table 2, *Annotated* depth maps), and show that they are competitive with the current state-of-the-art, surpassing all other semi-supervised methods and achieving similar performance as the Deep Ordinal Regression Networks proposed in [1]. However, because our proposed approach also leverages unlabeled data via the photometric loss, we are able to process the entire image at training time, thus producing more visually consistent depth maps, as shown in Fig. 3d. Note that our pose network also learns to produce metrically-accurate estimates ( $t_{rel}$  and  $r_{rel}$  of 2.4296 and 0.5747 respectively on the KITTI odometry benchmark for training sequences 00/03/04/05/07, and 6.8017 and 2.7142 for testing sequences 01/02/06/08/10), even though there was no direct pose supervision.

#### 4.4 Sparse Depth Labels

The *Annotated* depth maps, used in this work and in most related works on supervised depth learning, have been carefully curated and are composed of accumulated LiDAR beams, resulting in information that is much denser than what is produced by traditional range sensors. In fact, sparse

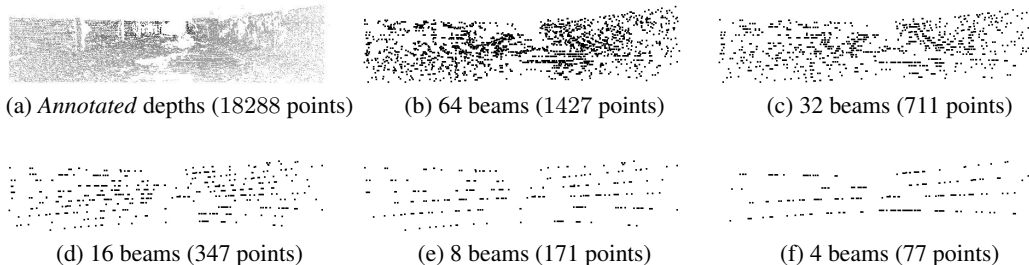


Figure 4: **Effects of LiDAR beam decimation** on the sparsity of depth labels. Relative to a full annotated depth map, a 4-beam representation contains only 0.42% of the number of valid depth values, or 0.06% of the total number of pixels in the image.

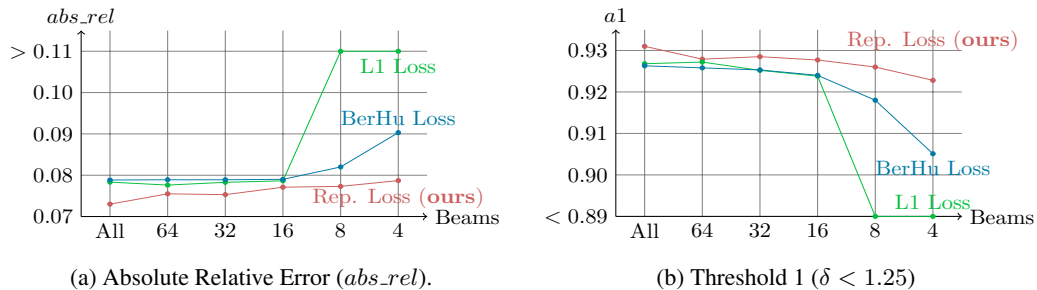


Figure 5: **Depth degradation with a decreasing number** of LiDAR beams, for semi-supervised training with different losses. Loss weight for each type of loss are optimized through grid search.

LiDAR sensors are not uncommon in most real-world robotics platforms, and devising ways to properly leverage this information for learning-based methods would be highly valuable as a way to decrease costs and increase data collection rate. In this section we investigate the effects that sparse information has in supervised depth training, and how our proposed semi-supervised training methodology is able to mitigate degradation in depth estimates when labels are substantially sparser than commonly reported in related works. This sparsity effect can be achieved by selectively masking out pixels from each of the 64 LiDAR beams, keeping only those belonging to equally spaced beams at increasing intervals. To maintain consistency with previous experiments, these masks are produced individually for each frame and the *original* depth values at each valid pixel is substituted by the corresponding *annotated* values. The effects of decimating depth maps in such a way is shown in Fig. 4, where we can see a substantial decrease in the amount of information available for supervised training as more beams are removed, reaching less than 100 valid pixels per image when only four beams are considered.

Even so, in Fig. 5 we can see that our proposed approach is capable of producing accurate depth estimates even when using as few as 4 beams for depth supervision at training time. In contrast, semi-supervision with other traditional losses starts to degrade after reaching a certain level of sparsity, empirically determined to be at 16 beams, or around 2% of the original number of valid pixels. This behavior is more visible on the L1 loss (green line), while the BerHu loss (blue line) degrades at a slower pace, however it still starts to decay exponentially, while our proposed Reprojected Distance loss (red line) maintains a roughly linear decay relative to the number of considered beams. This decay is also numerically represented in Table 2, for all considered metrics. In fact, our results when using only 8 beams are comparable to Amiri et al. [28] and Luo et al. [26], considered the current state-of-the-art for semi-supervised monocular depth estimation, and our results when using only 4 beams are better than Kuznietsov et al. [27].

## 5 Conclusion

This paper introduces a novel semi-supervised training methodology for monocular depth estimation, that both improves on the current state of the art and is also more robust to the presence of sparse labels. To accomplish this, we propose a new supervised loss term that operates in the image space, and thus is compatible with the widely used photometric loss in the semi-supervised setting. We show, using the popular KITTI benchmark, that our proposed methodology can efficiently incorporate information from depth labels into pretrained self-supervised models, allowing them to produce metrically accurate estimates while further improving the overall quality of resulting depth maps. Further analysis also indicates that our model presents a strong robustness to the degradation of available supervised labels, reaching results competitive with the current state-of-the-art even when as few as 4 beams are considered at training time. Future work will focus on the scalability of the proposed semi-supervised training methodology, investigating its application to different datasets, sources of labeled information and robustness to domain transfer.



## References

- [1] H. Fu, M. Gong, C. Wang, K. Batmanghelich, and D. Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2002–2011, 2018.
- [2] S. Pillai, R. Ambrus, and A. Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. In *arXiv preprint arXiv:1810.01849*, 2018.
- [3] T. Zhou, R. Tucker, J. Flynn, G. Fyffe, and N. Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*, 2018.
- [4] J. Flynn, I. Neulander, J. Philbin, and N. Snavely. Deepstereo: Learning to predict new views from the world’s imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5515–5524, 2016.
- [5] R. Garg, V. K. BG, G. Carneiro, and I. Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European Conference on Computer Vision*, pages 740–756. Springer, 2016.
- [6] T. Zhou, M. Brown, N. Snavely, and D. G. Lowe. Unsupervised learning of depth and ego-motion from video. In *CVPR*, volume 2, page 7, 2017.
- [7] R. Mahjourian, M. Wicke, and A. Angelova. 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*, pages 5667–5675, 2018.
- [8] C. Godard, O. Mac Aodha, and G. J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, 2017.
- [9] V. Guizilini, S. Pillai, R. Ambrus, and A. Gaidon. Packnet-sfm: 3d packing for self-supervised monocular depth estimation. *arXiv preprint arXiv:1905.02693*, 2019.
- [10] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [11] D. Eigen, C. Puhrsch, and R. Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014.
- [12] B. Li, C. Shen, Y. Dai, A. Van, and M. He. Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 201, pages 1119–1127, 2015.
- [13] I. Laina, C. Rupprecht, V. Belagiannis, F. Tombari, and N. Navab. Deeper depth prediction with fully convolutional residual networks. In *International Conference on 3D Vision (3DV)*, pages 239–248, 2016.
- [14] X. Qi, R. Liao, Z. Liu, R. Urtasun, and J. Jia. Geonet: Geometric neural network for joint depth and surface normal estimation. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 283–291, 2018.
- [15] J.-H. Lee, M. Heo, K.-R. Kim, and C.-S. Kim. Single-image depth estimation based on fourier domain analysis. In *International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 330–339, 2018.
- [16] C. Godard, O. Mac Aodha, and G. J. Brostow. Unsupervised monocular depth estimation with left-right consistency. In *CVPR*, volume 2, page 7, 2017.
- [17] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.

- [18] B. Ummenhofer, H. Zhou, J. Uhrig, N. Mayer, E. Ilg, A. Dosovitskiy, and T. Brox. Demon: Depth and motion network for learning monocular stereo. In *IEEE Conference on computer vision and pattern recognition (CVPR)*, volume 5, page 6, 2017.
- [19] Z. Yin and J. Shi. Geonet: Unsupervised learning of dense depth, optical flow and camera pose. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, 2018.
- [20] V. Casser, S. Pirk, R. Mahjourian, and A. Angelova. Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. *AAAI*, 2019.
- [21] Y. Zou, Z. Luo, and J.-B. Huang. Df-net: Unsupervised joint learning of depth and flow using cross-task consistency. In *European Conference on Computer Vision*, 2018.
- [22] M. Klodt and A. Vedaldi. Supervising the new with the old: Learning sfm from sfm. In *European Conference on Computer Vision*, pages 713–728. Springer, 2018.
- [23] C. Wang, J. M. Buenaposada, R. Zhu, and S. Lucey. Learning depth from monocular videos using direct methods. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2022–2030, 2018.
- [24] N. Yang, R. Wang, J. Stückler, and D. Cremers. Deep virtual stereo odometry: Leveraging deep depth prediction for monocular direct sparse odometry. *arXiv preprint arXiv:1807.02570*, 2018.
- [25] X. Guo, H. Li, S. Yi, J. Ren, and X. Wang. Learning monocular depth by distilling cross-domain stereo networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 484–500, 2018.
- [26] Y. Luo, J. Ren, M. Lin, J. Pang, W. Sun, H. Li, and L. Lin. Single view stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 155–163, 2018.
- [27] Y. Kuznetsov, J. Stückler, and B. Leibe. Semi-supervised deep learning for monocular depth map prediction. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6647–6655, 2017.
- [28] A. J. Amiri, S. Y. Loo, and H. Zhang. Semi-supervised monocular depth estimation with left-right consistency using deep neural network. *arXiv preprint arXiv:1905.07542v1*, 2019.
- [29] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4040–4048, 2016.
- [30] C. Godard, O. Mac Aodha, and G. Brostow. Digging into self-supervised monocular depth estimation. *arXiv preprint arXiv:1806.01260v3*, 2018.
- [31] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [32] J. Uhrig, N. Schneider, L. Schneider, U. Franke, T. Brox, and A. Geiger. Sparsity invariant cnns. *3DV*, 2017.
- [33] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [34] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [35] J. Deng, W. Dong, R. Socher, L. jia Li, K. Li, and L. Fei-fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2009.