

Self-Supervised Monocular Depth Estimation: Solving the Edge-Fattening Problem

Xingyu Chen¹ Ruonan Zhang¹ Ji Jiang¹ Yan Wang¹ Ge Li¹ Thomas H. Li ^{1,2,3}

¹School of Electronic and Computer Engineering, Peking University ²Advanced Institute of Information Technology, Peking University

³Information Technology R&D Innovation Center of Peking University

{cxy, zhangrn, jiangji, wyan}@stu.pku.edu.cn geli@ece.pku.edu.cn tli@aiit.org.cn

<https://github.com/xingyuuchen/tri-depth>

Abstract

Self-supervised monocular depth estimation (MDE) models universally suffer from the notorious edge-fattening issue. Triplet loss, as a widespread metric learning strategy, has largely succeeded in many computer vision applications. In this paper, we redesign the patch-based triplet loss in MDE to alleviate the ubiquitous edge-fattening issue. We show two drawbacks of the raw triplet loss in MDE and demonstrate our problem-driven redesigns. First, we present a min. operator based strategy applied to all negative samples, to prevent well-performing negatives sheltering the error of edge-fattening negatives. Second, we split the anchor-positive distance and anchor-negative distance from within the original triplet, which directly optimizes the positives without any mutual effect with the negatives. Extensive experiments show the combination of these two small redesigns can achieve unprecedented results: Our powerful and versatile triplet loss not only makes our model outperform all previous SoTA by a large margin, but also provides substantial performance boosts to a large number of existing models, while introducing no extra inference computation at all.

1. Introduction

Estimating how far each pixel in the image is away from the camera is a fundamental problem in 3D computer vision. This technique is desirable in various fields, such as autonomous driving [43, 49], AR [29] and robotics [13]. The majority of these applications adopt the appealing off-the-shelf hardware, e.g. Lidar sensor or RGB-D cameras, to enable agents. On the contrary, monocular videos or stereo pairs are much easier to obtain. Hence, quantities of researches [51, 10, 45] have been conducted to estimate promising dense depth map from only a single RGB image. Most of them and their follow-ups formulated the problem

as image reconstruction [10, 51, 45, 46, 24]. In particular, given a target image, the network infers its pixel-aligned depth map. Next, with a known or estimated camera ego-motion, every pixel in the target image can be reprojected to the reference image(s) which is taken from different view-point(s) of the same scene. The reconstructed image can be generated by sampling from the source image, and the training loss is based on the photometric distance between the reconstructed and target image. In this way, the network is trained under self-supervision.

Nevertheless, these approaches suffered severally from the notorious ‘edge-fattening’ problem, where the objects’ depth predictions are always ‘fatter’ than themselves. We visualize the problem and analyse its cause in Fig. 1. Disappointingly, there is no corresponding one-size-fits-all solution, let alone a lightweight one.

Deep metric learning seeks to learn a feature space where semantically similar samples are mapped to close locations, while semantically dissimilar samples are mapped to distant locations. Pioneeringly, [24] first introduced the patch-based semantics-guided triplet loss into MDE. Its key idea is to encourage pixels within each object instance to have similar depths, while those across semantic boundaries to have depth differences, as shown in Fig. 2.

However, we find that the straightforward application of the triplet loss only produces poor results. In this paper, we dig into the weakness of the patch-based triplet loss, and improve it through a problem-driven manner, reaching unprecedented performances.

First, in some boundary regions the edge-fattening areas could be thin, their contributions are small compared to the non-fattening area, in which case the thin but defective regions’ error could be covered up. Therefore, we change the optimizing strategy to only focus on the fattening area. The problematic case illustrated in Fig. 3 motivates our strategy, where we leave the normal area alone, and concentrate the optimization on the poor-performing negatives.

Second, we point out that the training objective of the original triplet loss [37] is to distinguish / discriminate, the network only has to make sure the correct answer’s score (the anchor-positive distance \mathcal{D}^+) wins the other choices’ scores (the anchor-negative distances \mathcal{D}^-) by a predefined margin, *i.e.*, $\mathcal{D}^- - \mathcal{D}^+ > m$, while the absolute value of \mathcal{D}^+ is not that important, see an example in Fig. 4. However, depth estimation is a regression problem since every pixel has its unique depth solution. Here, we have no idea of the exact depth differences between the intersecting objects, thus, it is also unknown how much \mathcal{D}^- should exceed \mathcal{D}^+ . But one thing for sure is that, in depth estimation, the smaller the \mathcal{D}^+ , the better, since depths within the same object are generally the same. We therefore split \mathcal{D}^- and \mathcal{D}^+ from within the original triplet and optimize them in isolation, where either error of the positives or negatives will be penalized individually and more directly. The problematic case illustrated in Fig. 5 motivates this strategy, where the negatives are so good enough to *cover up* the badness of the positives. In other words, even though \mathcal{D}^+ is large and needs to be optimized, \mathcal{D}^- already exceeds \mathcal{D}^+ more than m , which hinders the optimization for \mathcal{D}^+ .

To sum up, this paper’s contributions are:

- We show two weaknesses of the raw patch-based triplet optimizing strategy in MDE, that it (i) could miss some thin but poor fattening areas; (ii) suffered from mutual effects between positives and negatives.
- To overcome these two limitations, (i) we present a *min.* operator based strategy applied on all negative samples, to prevent the good negatives sheltering the error of poor (edge-fattening) negatives; (ii) We split the anchor-positive distance and anchor-negative distance from within the original triplet, to prevent the good negatives sheltering the error of poor positives.
- Our redesigned triplet loss is powerful, generalizable and lightweight: Experiments show that it not only makes our model outperform all previous methods by a large margin, but also provides substantial boosts to a large number of existing models, while introducing no extra inference computation at all.

2. Related Work

2.1. Self-Supervised Monocular Depth Estimation

Garg *et al.* [7] firstly introduced the novel concept of estimating depth without depth labels. Then, SfMLearner presented by Zhou *et al.* [51] required only monocular videos to predict depths, because they employed an additional pose network to learn the camera ego-motion. Godard *et al.* [10] presented *Monodepth2*, with surprisingly simple methods handling occlusions and dynamic objects in a non-learning

manner, both of which add no network parameters. Multiple works leveraged additional supervisions, *e.g.* estimating depth with traditional stereo matching method [45, 38] and semantics [22]. HR-Depth [30] proved that higher-resolution input images can reduce photometric loss with the same prediction. Manydepth [46] proposed to make use of multiple frames available at test time and leverage the geometric constraint by building a cost volume, achieving superior performance. [35] integrated wavelet decomposition into the depth decoder, reducing its computational complexity. Some other recent works estimated depths in more severe environments, *e.g.* indoor scenes [21] or in night-time [40, 27]. Innovative loss functions were also developed, *e.g.* constraining 3D point clouds consistency [20]. To deal with the notorious edge-fattening issue as illustrated in Fig. 1, most existing methods utilized an occlusion mask [11, 52] in the way of removing the incorrect supervision under photometric loss. We argue that although this exclusion strategy works to some extent, the masking technique could prevent these occluded regions from learning, because no supervision exists for them any longer. In contrast, our triplet loss closes this gap by providing additional supervision signals directly to these occluded areas.

2.2. Deep Metric Learning

The idea of comparing training samples in the high-level feature space [3, 1] is a powerful concept, since there could be more task-specific semantic information in the feature space than in the low-level image space. The contrastive loss function (*a.k.a.* discriminative loss) [16] is formulated by whether a *pair* of input samples belong to the same class. It learns an embedding space where samples within the same class are close in distance, whereas unassociated ones are farther away from each other. The triplet loss [37] is an extension of contrastive loss, with three samples as input each time, *i.e.* the anchor, the positive(s) and the negative(s). The triplet loss encourages the anchor-positive distance to be smaller than the anchor-negative loss by a margin m . Generally, the triplet loss function was introduced to face recognition and re-identification [37, 54], image ranking [41], to name a few. Jung *et al.* [24] first plugged triplet loss into self-supervised MDE, guided by another pretrained semantic segmentation network. In our experiments, we show that without other contributions in [24], the raw semantic-guided triplet loss only yielded a very limited improvement. We tackle various difficulties when plugging triplet loss into MDE, allowing our redesigned triplet loss outperforms existing ones by a large margin, with unparalleled accuracy.

3. Analysis of the Edge-fattening Problem

Before delving into our powerful triplet loss, it is necessary to make the problem clear. We first show the

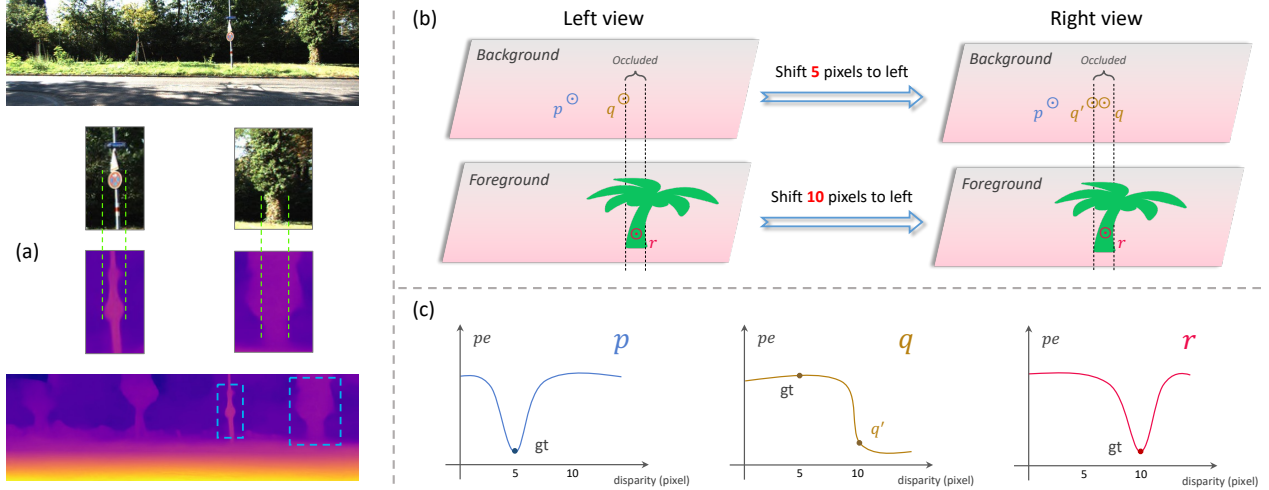


Figure 1. **Analysis of the edge-fattening issue.** (a) Example of the edge-fattening issue. The depth predictions of foreground objects (e.g. the tree-trunk and poles) are ‘fatter’ than the objects themselves. (b) In the left view, pixel p and q are located in the background with a disparity of 5 pixels, and q will be occluded if any further to the right. r is on the tree with a disparity of 10 pixels. (c) p and r are OK - their gt disparity is the global optimum of the photometric error (loss). q suffers from edge-fattening issue. Since q is occluded by the tree in the right view, the photometric error of its gt disparity 5 is large. The photometric loss function therefore struggles to find another location that has a small loss, i.e., shifting another 5 pixels to reach the nearest background pixel q' . However, q' is not the true correspondence of q . As a result, disparity of the background q equals to that of the foreground r , leading to the edge-fattening issue. Details in Sec. 3.

behaviour of the so-called edge-fattening issue in self-supervised MDE, then analyse its cause in Fig. 1. This motivates us to introduce our resigned triplet loss.

The ubiquitous edge-fattening problem, limiting performances of vast majorities of self-supervised MDE models [10, 46, 45, 24, 51, 30], manifests itself as yielding inaccurate object depths that partially leak into the background at object boundaries, as shown in Fig. 1a.

We first lay out the final conclusion: *The networks misjudge the background near the foreground as the foreground, so that the foreground looks fatter than itself.*

The cause could be traced back to occlusions of background pixels as illustrated in Fig. 1b&c. The background pixels visible in the target (left) image but invisible in the source (right) image suffer from incorrect supervision under the photometric loss, since no exact correspondences in the source image exist for them at all.

The crux of the matter is that, for the occluded pixel (e.g. pixel q in Fig. 1b), the photometric loss still needs to seek a pixel with a similar appearance (in the right view) to be its fake correspondence. Generally, for a background pixel, only another background pixel could show a small photometric loss, so the problem turns into finding the nearest background pixel for the occluded pixel (q) in the source (right) view. Since the foreground has a smaller depth Z than the background, it has a larger disparity d owing to the geometry projection constrain: $Z = \frac{f \cdot b}{d}$, where b is the fixed camera baseline. Consequently, the photometric loss has to shift more to the left to find the nearest background

pixel (the fake solution, e.g. pixel q' in Fig. 1b&c). In this way, the occluded background pixels share the same disparities as the foreground, forming the edge-fattening issue.

4. Methodology

4.1. Self-supervised Depth Estimation

Following [51, 10], given a monocular and/or stereo video, we first train a depth network ψ_{depth} consuming a single target image I_t as input, and outputs its pixel-aligned depth map $D_t = \psi_{depth}(I_t)$. Then, we train a pose network ψ_{pose} taking temporally adjacent frames as input, and outputs the relative camera pose $T_{t \rightarrow t+n} = \psi_{pose}(I_t, I_{t+n})$.

Suppose we have access to the camera intrinsics K , along with D_t and $T_{t \rightarrow t+n}$, we warp I_{t+n} into I_t to generate the reconstructed image \tilde{I}_{t+n} :

$$\tilde{I}_{t+n} = I_{t+n} \langle proj(D_t, T_{t \rightarrow t+n}, K) \rangle, \quad (1)$$

where $\langle \cdot \rangle$ is the differentiable bilinear sampling operator according to [10].

To make use of multiple input frames, we build a cost volume [46] using discrete depth values from a predefined range $[d_{min}, d_{max}]$. Moreover, d_{min} and d_{max} are dynamically adjusted during training to find the best scale [46].

In order to evaluate the reconstructed images \tilde{I}_{t+n} , we adopt the edge-aware smoothness loss [18] and the photometric reprojection loss measured by $\mathcal{L}_1 + \mathcal{L}_{ssim}$ [51, 10]:

$$\mathcal{L}_{pe} = \frac{\alpha}{2} (1 - SSIM(I_t, \tilde{I}_{t+n})) + (1 - \alpha) \| I_t - \tilde{I}_{t+n} \|_1, \quad (2)$$

where $\alpha = 0.85$ by default and *SSIM* [44] computes pixel similarity over a 3×3 window. See [46] for more details of our network architecture.

4.2. Redesigned Triplet Loss

We first introduce our baseline depth-patch-based triplet loss, and then give detailed analyses of our two redesigns.

4.2.1 Patch-based Triplet Loss

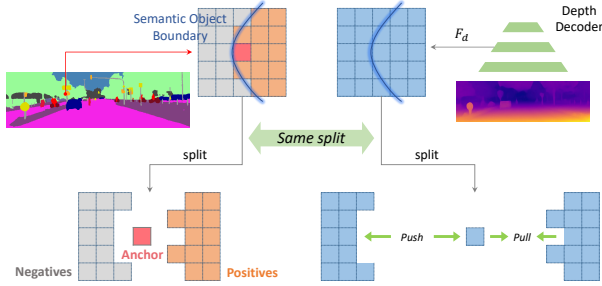


Figure 2. **The baseline semantic-aware triplet loss.** For each pixel in the semantic boundary region, we group the local patch of its corresponding depth feature (F_d) into a triplet according to the semantic patch. Next, the triplet loss minimizes anchor-positive distance (\mathcal{D}^+) and maximizes the anchor-negative distance (\mathcal{D}^-).

There lies a fact that pixels within a certain object have similar depths, while pixels across object boundaries may have large depth differences. However, in many cases, the depth boundary does not align to the semantic boundary as shown in Fig. 1. Therefore, similar to [24], we overcome this problem through deep metric learning [25, 48]. Specifically, we group the pixels in a local patch into triplets: the central pixel as the *anchor*; those share the same semantic class with the anchor as *positives*; while those have a different semantic class from the anchor as *negatives*. Then, we refer to the sets of positive and negative pixels of the anchor \mathcal{P}_i as \mathcal{P}_i^+ and \mathcal{P}_i^- , respectively. For example, $|\mathcal{P}_i^-| = 0$ implies \mathcal{P}_i is located inside an object.

The anchor-positive distance \mathcal{D}^+ and anchor-negative distance \mathcal{D}^- are defined as the mean of Euclidean distance of L_2 normalized depth features [24]:

$$\mathcal{D}^+(i) = \frac{1}{|\mathcal{P}_i^+|} \sum_{j \in \mathcal{P}_i^+} \left\| \hat{F}_d(i) - \hat{F}_d(j) \right\|_2^2, \quad (3)$$

$$\mathcal{D}^-(i) = \frac{1}{|\mathcal{P}_i^-|} \sum_{j \in \mathcal{P}_i^-} \left\| \hat{F}_d(i) - \hat{F}_d(j) \right\|_2^2, \quad (4)$$

where $\hat{F}^d = F_d / \|F_d\|$.

Intuitively, \mathcal{D}^+ should be minimized, whereas \mathcal{D}^- should be maximized. However, naively maximizing \mathcal{D}^- as large as possible is unfavourable, since the depth differences between two spatially adjacent objects are not always

large. Hence, a margin m is employed to prevent \mathcal{D}^- from exceeding \mathcal{D}^+ immoderately [37]. In this way, we want:

$$\mathcal{D}^+ + m < \mathcal{D}^-, \quad (5)$$

where m is the threshold controlling the least distance between \mathcal{D}^+ and \mathcal{D}^- .

In addition, since we also adopt an off-the-shelf pre-trained semantic segmentation model [53] whose predictions could also be inaccurate, we only optimize the boundary anchors with $|\mathcal{P}_i^+|$ and $|\mathcal{P}_i^-|$ both larger than k [24]:

$$\mathcal{D}^+(i) + m < \mathcal{D}^-(i), \quad \forall \mathcal{P}_i \in \Gamma, \quad (6)$$

$$\Gamma = \left\{ \mathcal{P}_i \mid (|\mathcal{P}_i^+| > k) \wedge (|\mathcal{P}_i^-| > k) \right\}, \quad (7)$$

where Γ is the set holding all semantic boundary pixels that fulfill the aforementioned constraint.

Consequently, the baseline triplet loss is defined as:

$$\mathcal{L}_{tri} = \frac{1}{|\Gamma|} \sum_{\mathcal{P}_i \in \Gamma} [\mathcal{D}^+(i) - \mathcal{D}^-(i) + m]_+, \quad (8)$$

where $[\cdot]_+$ is the hinge function.

In this way, we do not require any ground truth annotations, neither in depth nor segmentation, which enables a fully self-supervised setting and can therefore better generalize to unseen scenarios.

4.2.2 Motivation 1: The Poor Negatives' Error may be Sheltered

Nevertheless, naively encouraging the *mean* \mathcal{D}^- in the local patch to increase leads to poor results. We trace its cause to the average operator in computing the anchor-negative distance. An example is demonstrated in Fig. 3, where the majorities of the negatives perform well, while only a few pixels inside the gap between the depth and semantic boundary perform poorly, triplet loss with the strategy of averaging \mathcal{D}^- is likely to mask the error of these poor negatives, as their contributions are small compared to large numbers of other well-performed negatives. As a result, the average anchor-negative distance is still not small enough to be optimized by the triplet loss with margin m , which further hinders the optimization for these thin but poor negatives.

4.2.3 Solution 1: Focus only on the Fattening Area

We point out that it is not a wise idea to simply increase the margin m to let cases like Fig. 3 participate in training, because it is not only these cases matter, we should find out the right m suitable for as many cases as possible. Instead, inspired by *PointNetVLAD* [39], a point cloud retrieval network, we introduce selecting the hardest negative to participate in training, while leaving other negatives alone. In this




Figure 3. **How the error of the low-proportion poor negatives is sheltered by the large-proportion well-performing negatives.** (a) An instance of the edge-fattening issue (detailed analysis in Fig. 1), where the depth prediction is ‘fatter’ than the object itself. In other words, negatives in the ‘fatter’ area perform poorly - their depth has to be the same as the background but not the foreground object. (b) The Euclidean distance of depth feature between all pixels in the local patch and the anchor. The yellow plane indicates \mathcal{D}^+ (mean anc-pos distances); the grey plane indicates \mathcal{D}^- (mean anc-neg distances); the green plane specifies the decision boundary of whether this triplet participates in training (see the hinge function in Eq. 8). \mathcal{D}^- is large such that the green plane lies beneath the grey. Thus, no learning happens. Disappointingly, the poor negative pixels in the ‘fatter’ area (e.g. the *poor negatives* marked in (a)) get no optimization. This is because of the *average* operator - the low-proportion poor negatives’ contributions are weakened by the large-proportion good negatives.

way, the anchor-negative distance becomes:

$$\mathcal{D}^{-'}(i) = \min_{j \in \mathcal{P}_i^-} \left\| \hat{F}^d(i) - \hat{F}^d(j) \right\|_2^2. \quad (9)$$

This comes from the fact that if the worst negative is good enough, other negatives from the local patch would no longer need optimizations. Compared to other computer vision tasks like image classification [50], point cloud retrieval [39], person re-identification [19], where hard-negative mining is expensive because the whole set of training data has to be revisited to compute the one with the largest \mathcal{D}^- , our strategy is much more efficient since our mining process happens only in the local patch, and it is quite easy to proceed all triplets in batches.

4.2.4 Motivation 2: The Positives’ Error may be Sheltered




Figure 4. In a classification task, both network 1&2 perform well (both *s.t.* Eq. 5), they undoubtedly judge the anchor as a cat because there is a clear distance gap ($> margin$) between the two choices.

In a classification task, it is OK (tolerable) that the anchor-positive distance (\mathcal{D}^+) is a bit large (e.g. Network 2 in Fig. 4), because it only has to be smaller enough than any of the samples belonging to a different class (e.g. the dog). In other words, for classification, it is the *comparison relationship* that $\mathcal{D}^- \gg \mathcal{D}^+$ counts, and neither the individual \mathcal{D}^- nor \mathcal{D}^+ counts. A classification network only has to make sure that the score of the correct answer wins all the others, while the absolute score of the correct answer is not that important. In fact, it makes no difference for \mathcal{D}^+ to be any value $\in [0, \mathcal{D}^- - m)$ in Fig. 4 when inferring.




Figure 5. **Example of how the error of poor positives is sheltered by the good negatives.** (a) From top to bottom: An RGB image; its predicted depth; semantic segmentation. (b) Image patch in (a), where the depth prediction of the car window (positives) is wrong - it has to be the same as the car body, while depth predictions of the background (negatives) are extremely accurate. (c) The average of anchor-negative distances (\mathcal{D}^-) is so large that the grey plane lies above the green plane. That is, the triplet loss stops working as Eq. 5 is satisfied. However, all poor positive car window pixels (the large \mathcal{D}^+) get no optimization.

However, when it comes to depth estimation, the situ-

ation changes. We emphasize that depth estimation is a regression problem, which aims to predict accurate depth values for every pixel. In this case, \mathcal{D}^- and \mathcal{D}^+ matter individually. For example, as depicted in Fig. 5b, the depth prediction of the car window (positives) should have been close to that of the car body (anchor), but in fact it doesn't. When using the original triplet loss, the depth prediction of the background negatives is so good (\mathcal{D}^- is so large enough) that they can shelter / cover up the error of the positives. In consequence, Eq. 5 is satisfied and no learning happens - all poor positive pixels (car window) are not going to be optimized.

4.2.5 Solution 2: Optimize Pos and Neg in Isolation

To this end, we no longer compare the value of $(\mathcal{D}^- - \mathcal{D}^+)$ with m . Instead, we easily split \mathcal{D}^+ and \mathcal{D}^- from within the original triplet loss (Eq. 8). We compare \mathcal{D}^- with a new margin m' , and optimize \mathcal{D}^+ directly. Concretely, our redesigned the triplet loss becomes:

$$\mathcal{L}'_{tri} = \frac{1}{|\Gamma|} \sum_{P_i \in \Gamma} \left(\mathcal{D}^+(i) + [m' - \mathcal{D}^-(i)]_+ \right), \quad (10)$$

where either error of the positives or negatives appears individually without any mutual effects. *E.g.*, in Fig. 5, the poor positives (wrong car window depth predictions) will be penalized directly, since \mathcal{D}^+ is no longer restricted to the comparison with \mathcal{D}^- . This is a bit like contrastive loss [16], but we point out that we are still a triplet loss, since we still

follow the idea of searching and optimizing both positives and negatives of the given anchor simultaneously. Another reason why we do not optimize the value of $(\mathcal{D}^- - \mathcal{D}^+)$ is that, we have no prior knowledge of how much the depth difference between the two-side objects should be. In fact, it varies with different objects. But one thing is certain: the smaller the \mathcal{D}^+ , the better, since we always want the depths of different parts of the same object to be the same.

5. Experiments

5.1. Implementation Details

We implement our methods in PyTorch [31]. We set the total training epoch as 20, the learning rate as 10^{-4} . Our triplet loss is applied throughout every layer of the depth decoder. We set the triplets' patch size to be 5×5 , and $k = 4$ [24] to filter out anchors located in potentially inaccurate semantic predictions. Because we use the minimum \mathcal{D}^- of all negative samples, we increase the original margin $m = 0.3$ to $m' = 0.65$. The rest of our model's setting is the same as [46]. Note that our triplet loss is not restricted to a fixed baseline, when plugging it into a new model, just leave all the original settings exactly unchanged.

5.2. Comparison to State-of-the-Arts

KITTI dataset [8] is selected as our testing benchmark where calibrated stereo images with depth information from the Lidar point clouds are provided. In detail, we use the KITTI Eigen split [6] to evaluate the depth predictions,

Method	PP	W × H	Data	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Ranjan <i>et al.</i> [36]	✗	832 × 256	S	0.148	1.149	5.464	0.226	0.815	0.935	0.973
EPC++ [28]	✗	832 × 256	S	0.141	1.029	5.350	0.216	0.816	0.941	0.976
Structure2depth [2]	✗	416 × 128	M	0.141	1.026	5.291	0.215	0.816	0.945	0.979
Videos in the wild [12]	✗	416 × 128	M	0.128	0.959	5.230	0.212	0.845	0.947	0.976
Guizilini <i>et al.</i> [15]	✗	640 × 192	M	0.102	<u>0.698</u>	<u>4.381</u>	0.178	0.896	0.964	0.984
Johnston <i>et al.</i> [23]	✗	640 × 192	M	0.106	0.861	4.699	0.185	0.889	0.962	0.982
Packnet-SFM [14]	✗	640 × 192	M	0.111	0.785	4.601	0.189	0.878	0.960	0.982
Li <i>et al.</i> [26]	✗	416 × 128	M	0.130	0.950	5.138	0.209	0.843	0.948	0.978
Patil <i>et al.</i> [32]	✗	640 × 192	M	0.111	0.821	4.650	0.187	0.883	0.961	0.982
Wang <i>et al.</i> [42]	✗	640 × 192	M	0.106	0.799	4.662	0.187	0.889	0.961	0.982
Monodepth2 MS [10]	✗	640 × 192	MS	0.106	0.818	4.750	0.196	0.874	0.957	0.979
Zhou <i>et al.</i> [51]	✗	640 × 192	M	0.183	1.595	6.709	0.270	0.734	0.902	0.959
WaveletMonodepth [35]	✗	640 × 192	S	0.109	0.845	4.800	0.196	0.870	0.956	0.980
HR-Depth [30]	✗	640 × 192	MS	0.107	0.785	4.612	0.185	0.887	0.962	0.982
FSRE-Depth [24]	✗	640 × 192	M	0.105	0.722	4.547	0.182	0.886	0.964	0.984
Depth-Hints [45]	✗	640 × 192	S	0.109	0.845	4.800	0.196	0.870	0.956	0.980
CADepth [47]	✗	640 × 192	S	0.106	0.849	4.885	0.204	0.869	0.951	0.976
SuperDepth [33]	✗	640 × 192	S	0.112	0.875	4.958	0.207	0.852	0.947	0.977
ManyDepth [46]	✗	640 × 192	M	<u>0.098</u>	0.770	4.459	<u>0.176</u>	<u>0.900</u>	<u>0.965</u>	<u>0.983</u>
Refine&Distill [34]	✗	-	S	<u>0.098</u>	0.831	4.656	0.202	0.882	0.948	0.973
TriDepth (Ours)	✗	640 × 192	M	0.093	0.665	4.272	0.172	0.907	0.967	0.984

Table 1. Comparison to previous SoTA on KITTI Eigen split [6]. Metrics are error metrics ↓ and accuracy metrics ↑. The Data column specifies the training data type: S - stereo images, M - monocular video and MS - stereo video. All models are trained with 192×640 images and Resnet18 [17] as backbone. All results are not Post-Processed [9]. Best results are in bold; second best are underlined. We outperform all previous methods by a large margin on exactly all metrics.

Method	Pub.	PP	W × H	Data	Extra time	Abs Rel	Sq Rel	RMSE	RMSE log	δ_1	δ_2	δ_3
Monodepth2 M [10]	ICCV 2019	✗	640 × 192	M	–	0.115	0.903	4.863	0.193	0.877	0.959	0.981
+ Ours	–	✗	640 × 192	M	+ 0ms	0.108	0.744	4.537	0.184	0.883	0.9630	0.983
Zhou <i>et al.</i> [51]	CVPR 2017	✗	640 × 192	M	–	0.183	1.595	6.709	0.270	0.734	0.902	0.959
+ Ours	–	✗	640 × 192	M	+ 0ms	0.148	1.098	5.150	0.212	0.819	0.949	0.980
Monodepth2 S [10]	ICCV 2019	✗	640 × 192	S	–	0.109	0.873	4.960	0.209	0.864	0.948	0.975
+ Ours	–	✗	640 × 192	S	+ 0ms	0.107	0.826	4.822	0.201	0.866	0.953	0.978
FSRE-Depth only SGT [24]	ICCV 2021	✗	640 × 192	M	–	0.113	0.836	4.711	0.187	0.878	0.960	0.982
+ Ours	–	✗	640 × 192	M	+ 0ms	0.108	0.746	4.507	0.182	0.884	0.964	0.983
Monodepth2 MS [10]	ICCV 2019	✗	640 × 192	MS	–	0.106	0.818	4.750	0.196	0.874	0.957	0.979
+ Ours	–	✗	640 × 192	MS	+ 0ms	0.105	0.753	4.563	0.182	0.887	0.963	0.983
Depth-Hints [45]	ICCV 2019	✗	640 × 192	S	–	0.109	0.845	4.800	0.196	0.870	0.956	0.980
+ Ours	–	✗	640 × 192	S	+ 0ms	0.106	0.843	4.774	0.194	0.875	0.957	0.980
HR-Depth [30]	AAAI 2020	✗	640 × 192	M	–	0.109	0.792	4.632	0.185	0.884	0.962	0.983
+ Ours	–	✗	640 × 192	M	+ 0ms	0.107	0.760	4.522	0.182	0.886	0.964	0.984
HR-Depth MS [30]	AAAI 2020	✗	640 × 192	MS	–	0.107	0.785	4.612	0.185	0.887	0.962	0.982
+ Ours	–	✗	640 × 192	MS	+ 0ms	0.105	0.751	4.512	0.181	0.890	0.963	0.983
ManyDepth [46]	CVPR 2021	✗	640 × 192	M	–	0.098	0.770	4.459	0.176	0.900	0.965	0.983
+ Ours	–	✗	640 × 192	M	+ 0ms	0.093	0.665	4.272	0.171	0.907	0.967	0.984
CADepth [47]	3DV 2021	✗	640 × 192	M	–	0.110	0.812	4.686	0.187	0.882	0.962	0.983
+ Ours	–	✗	640 × 192	M	+ 0ms	0.105	0.745	4.530	0.181	0.888	0.965	0.984

Table 2. Comparisons of existing models with and without our method on KITTI Eigen split [6]. All models are trained with 192×640 images and Resnet18 [17] as backbone. All results are not Post-Processed [9]. Models augmented with our powerful triplet loss are highlighted, achieving better results than their original counterparts on exactly all metrics, while no extra inference computation is needed.

Naive Triplet Loss	Hard Negative	Isolated Triplet	Abs Rel	Sq Rel	RMSE	RMSE log	δ_1	δ_2	δ_3
			0.115	0.903	4.863	0.193	0.877	0.959	0.981
✓			0.113	0.836	4.711	0.187	0.878	0.960	0.982
✓	✓		0.110	0.782	4.622	0.185	0.881	0.963	0.983
✓		✓	0.111	0.823	4.645	0.186	0.881	0.962	0.982
✓	✓	✓	0.108	0.746	4.507	0.182	0.884	0.964	0.983

Table 3. Ablations of our different strategies on KITTI Eigen split [6]. Here, we use Monodepth2 [10] as the baseline. All models are trained with 192×640 monocular videos and Resnet18 [17] as backbone. All results are not post-processed [9].

which is commonly used for comparisons of self-supervised MDE models. The depth ground truth is from the provided sparse Lidar point clouds, which is used to measure all seven standard metrics. As for fair comparison, we adopt the standard depth ceiling of 80 meters [7]. Please refer to [6] for more evaluation details. We report the results in Tab. 1, showing that we produce an unprecedented performance. Specifically, even though we make no use of any stereo pairs during training, our model still outperforms all SoTA by a large margin on every metric. To enable real-time applications, we abandon the post-processing technique [10] which doubles the inference computation, even so, we still show comparable results to previous post-processed SoTA results [45, 30, 24, 47]. Visualization of solving the edge-fattening issue is shown in Fig. 6.

5.3. Generalizability

As mentioned before, our triplet loss is not only powerful, but also versatile and lightweight. After showing our unparalleled state-of-the-art performance, we demonstrate the high generalizability of the proposed triplet loss.

We integrate our triplet loss into a large number of existing open-source methods (some of them are SoTA before ours), and show that our triplet loss can help them to achieve better scores, as reported in Tab. 2. Remarkably, in all integrations, we bring a substantial performance boost to the original model, in exactly all seven metrics. Furthermore, we do not lead in any extra inference computation at all, which shows our lightness. For the benefit of the whole MDE community, we publish our full implementation, since we believe this powerful triplet loss could enhance more subsequent depth estimators.

5.4. Ablation Study

In Tab. 3, we validate the effectiveness of our two design decisions in turn. The original triplet loss, migrated from deep metric learning by Jung *et al.* [24], yields only a slight improvement to the baseline model [10]. Through either of our two contributions, the performance is substantially improved, and not surprisingly, our full model performs the best. It is also worth noting that none of our contributions introduces any computation overhead when inferring.

5.4.1 Consideration on Margin m'

We further compare different values of the margin m' in our redesigned triplet loss, as shown in Tab. 4. We select $m' = 0.65$ because it shows the best result. Either increasing or decreasing m' leads to performance degradation. We point out that selecting the margin of triplet loss in depth estimation is a difficult task, since we have no prior knowledge of whether the semantic boundaries guarantee a large depth difference or not. When m' is too large, the triplet loss would be too strict, *i.e.* leading to some false positives. *E.g.*, let us consider an m' encouraging the depth difference between the intersecting objects to be 100m, then the boundaries with a depth difference of 5m will be wrongly penalized. When m' is too small, the triplet loss will be too tolerant, *i.e.* leading to some false negatives. *E.g.*, let us consider an m' encouraging the depth difference between the intersecting objects just to be 5m, then, boundaries with a depth difference of 100m but the prediction only has a depth difference of 80m, are going to be neglected by the optimization strategy.

m'	Abs Rel	Sq Rel	RMSE	R log	δ_1	δ_2	δ_3
0.50	0.110	0.818	4.646	0.186	0.882	0.962	0.982
0.60	0.109	0.806	4.667	0.184	0.883	0.961	0.983
0.65	0.108	0.746	4.507	0.182	0.884	0.964	0.983
0.70	0.109	0.763	4.613	0.184	0.882	0.963	0.983
0.80	0.110	0.787	4.601	0.185	0.880	0.962	0.983

Table 4. **Ablation on margin m' .** All models are trained with 192×640 monocular videos and Resnet18 [17] as backbone. All results are not post-processed [9].




Figure 6. **Visual comparisons of an existing model with and without our triplet loss.** The edge-fattening problem is significantly alleviated. The depths of thin structures are better aligned to their RGB counterparts, *e.g.* road signs and poles.

5.5. Can Current MDE Go into Production?

Apart from solving the ubiquitous edge-fattening problem, we also reveal new problems for our whole MDE community as the future works. We hope our community could focus not only on the performance numbers, but also on the barriers that prevent our practical application.

The scenario shown in Fig. 7 is not from any open datasets, but is captured by one of our authors. As can be




Figure 7. **A fatal failure case for autonomous driving.** The depth predictions of the crossbar in the middle of the input image are totally wrong, which may lead to a fatal error in autonomous driving. When vertically flipping the input image, the depth of the crossbar is correctly estimated, while other objects, *e.g.*, people, are wrong again.

seen, quantities of self-supervised MDE models fail completely in predicting the depth of the crossbar. We speculate that there are two probable reasons:

- Existing methods rely too heavily on the images' shape features in the dataset, and by chance, there are no horizontal bars in KITTI [8] dataset.
- When training with stereo pairs, for the horizontal bar, the photometric loss for all disparities is small, thus no clear global optimum exists. The network therefore tends to mix the crossbar and the background road together.

We expect future MDE models to better generalize to real-life scenes via better feature representation techniques, *e.g.* deformable convolution [4] or data augmentation in feature space [5].

6. Conclusion

In this paper, we solve the notorious and ubiquitous edge-fattening issue in self-supervised MDE by introducing a well-designed triplet loss. We tackle two drawbacks of the original patch-based triplet loss in MDE. First, we propose to apply *min.* operator on computing anchor-negative distance, to prevent the error of the edge-fattening negatives from being masked. Second, we split the anchor-positive distance and anchor-negative distance from within the original triplet loss. This strategy provides more direct optimizations to the positives, without the mutual effect with the negatives. Our triplet loss is highly powerful, versatile and lightweight. Experiments show that it not only brings our *TriDepth* model an unprecedented SoTA performance, but also provides substantial performance boosts to a large number of existing models, without introducing any extra inference computation at all.

Acknowledgement. This work was supported by National Natural Science Foundation of China (No. 62172021).

References

- [1] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah. Signature verification using a” siamese” time delay neural network. *Advances in neural information processing systems*, 6, 1993.
- [2] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8001–8008, 2019.
- [3] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 539–546. IEEE, 2005.
- [4] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.
- [5] Terrance DeVries and Graham W Taylor. Dataset augmentation in feature space. *arXiv preprint arXiv:1702.05538*, 2017.
- [6] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *arXiv preprint arXiv:1406.2283*, 2014.
- [7] Ravi Garg, Vijay Kumar Bg, Gustavo Carneiro, and Ian Reid. Unsupervised cnn for single view depth estimation: Geometry to the rescue. In *European conference on computer vision*, pages 740–756. Springer, 2016.
- [8] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012.
- [9] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 270–279, 2017.
- [10] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3828–3838, 2019.
- [11] Juan Luis GonzalezBello and Munchurl Kim. Forget about the lidar: Self-supervised depth estimators with med probability volumes. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12626–12637. Curran Associates, Inc., 2020.
- [12] Ariel Gordon, Hanhan Li, Rico Jonschkowski, and Anelia Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8977–8986, 2019.
- [13] Brent Griffin, Victoria Florence, and Jason Corso. Video object segmentation-based visual servo control and object depth estimation on a mobile robot. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1647–1657, 2020.
- [14] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Rantos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2485–2494, 2020.
- [15] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-guided representation learning for self-supervised monocular depth. *arXiv preprint arXiv:2002.12319*, 2020.
- [16] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 2, pages 1735–1742. IEEE, 2006.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [18] Philipp Heise, Sebastian Klose, Brian Jensen, and Alois Knoll. Pm-huber: Patchmatch with huber regularization for stereo matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2360–2367, 2013.
- [19] Alexander Hermans, Lucas Beyers, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017.
- [20] Noriaki Hirose, Satoshi Koide, Keisuke Kawano, and Ruho Kondo. Plg-in: Pluggable geometric consistency loss with wasserstein distance in monocular depth estimation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12868–12874. IEEE, 2021.
- [21] Pan Ji, Runze Li, Bir Bhanu, and Yi Xu. Monoindoor: Towards good practice of self-supervised monocular depth estimation for indoor environments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12787–12796, October 2021.
- [22] Jianbo Jiao, Ying Cao, Yibing Song, and Rynson Lau. Look deeper into depth: Monocular depth estimation with semantic booster and attention-driven loss. In *Proceedings of the European conference on computer vision (ECCV)*, pages 53–69, 2018.
- [23] Adrian Johnston and Gustavo Carneiro. Self-supervised monocular trained depth estimation using self-attention and discrete disparity volume. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4756–4765, 2020.
- [24] Hyunyoung Jung, Eunhyeok Park, and Sungjoo Yoo. Fine-grained semantics-aware representation enhancement for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12642–12652, 2021.
- [25] Sungyeon Kim, Dongwon Kim, Minsu Cho, and Suha Kwak. Proxy anchor loss for deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3238–3247, 2020.

- [26] Hanhan Li, Ariel Gordon, Hang Zhao, Vincent Casser, and Anelia Angelova. Unsupervised monocular depth learning in dynamic scenes. *arXiv preprint arXiv:2010.16404*, 2020.
- [27] Yawen Lu and Guoyu Lu. An alternative of lidar in night-time: Unsupervised depth estimation based on single thermal image. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3833–3843, January 2021.
- [28] Chenxu Luo, Zhenheng Yang, Peng Wang, Yang Wang, Wei Xu, Ram Nevatia, and Alan Yuille. Every pixel counts++: Joint learning of geometry and motion with 3d holistic understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2624–2641, 2019.
- [29] Xuan Luo, Jia-Bin Huang, Richard Szeliski, Kevin Matzen, and Johannes Kopf. Consistent video depth estimation. *ACM Transactions on Graphics (ToG)*, 39(4):71–1, 2020.
- [30] Xiaoyang Lyu, Liang Liu, Mengmeng Wang, Xin Kong, Lina Liu, Yong Liu, Xinxin Chen, and Yi Yuan. HR-Depth: High resolution self-supervised monocular depth estimation. *arXiv preprint arXiv:2012.07356*, 6, 2020.
- [31] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [32] Vaishakh Patil, Wouter Van Gansbeke, Dengxin Dai, and Luc Van Gool. Don’t forget the past: Recurrent depth estimation from monocular video. *IEEE Robotics and Automation Letters*, 5(4):6813–6820, 2020.
- [33] Sudeep Pillai, Rareş Ambruş, and Adrien Gaidon. Superdepth: Self-supervised, super-resolved monocular depth estimation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9250–9256. IEEE, 2019.
- [34] Andrea Pilzer, Stephane Lathuiliere, Nicu Sebe, and Elisa Ricci. Refine and distill: Exploiting cycle-inconsistency and knowledge distillation for unsupervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9768–9777, 2019.
- [35] Michaël Ramamonjisoa, Michael Firman, Jamie Watson, Vincent Lepetit, and Daniyar Turmukhambetov. Single image depth prediction with wavelet decomposition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11089–11098, 2021.
- [36] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J Black. 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*, pages 12240–12249, 2019.
- [37] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [38] Fabio Tosi, Filippo Aleotti, Matteo Poggi, and Stefano Mattoccia. Learning monocular depth estimation infusing traditional stereo knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [39] Mikaela Angelina Uy and Gim Hee Lee. PointNetVLAD: Deep point cloud based retrieval for large-scale place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4470–4479, 2018.
- [40] Madhu Vankadari, Sourav Garg, Anima Majumder, Swagat Kumar, and Ardhendu Behera. Unsupervised monocular depth estimation for night-time images using adversarial domain feature adaptation. In *European Conference on Computer Vision*, pages 443–459. Springer, 2020.
- [41] Jiang Wang, Yang Song, Thomas Leung, Chuck Rosenberg, Jingbin Wang, James Philbin, Bo Chen, and Ying Wu. Learning fine-grained image similarity with deep ranking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1386–1393, 2014.
- [42] Jianrong Wang, Ge Zhang, Zhenyu Wu, XueWei Li, and Li Liu. Self-supervised joint learning framework of depth estimation via implicit cues. *arXiv preprint arXiv:2006.09876*, 2020.
- [43] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8445–8453, 2019.
- [44] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [45] Jamie Watson, Michael Firman, Gabriel J Brostow, and Daniyar Turmukhambetov. Self-supervised monocular depth hints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2162–2171, 2019.
- [46] Jamie Watson, Oisín Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1164–1174, 2021.
- [47] Jiaying Yan, Hong Zhao, Penghui Bu, and YuSheng Jin. Channel-wise attention-based network for self-supervised monocular depth estimation. In *2021 International Conference on 3D Vision (3DV)*, pages 464–473. IEEE, 2021.
- [48] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Deep metric learning for person re-identification. In *2014 22nd international conference on pattern recognition*, pages 34–39. IEEE, 2014.
- [49] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. *arXiv preprint arXiv:1906.06310*, 2019.
- [50] Baosheng Yu, Tongliang Liu, Mingming Gong, Changxing Ding, and Dacheng Tao. Correcting the triplet selection bias for triplet loss. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 71–87, 2018.
- [51] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1851–1858, 2017.

- [52] Shengjie Zhu, Garrick Brazil, and Xiaoming Liu. The edge of depth: Explicit constraints between segmentation and depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13116–13125, 2020.
- [53] Yi Zhu, Karan Sapra, Fitsum A Reda, Kevin J Shih, Shawn Newsam, Andrew Tao, and Bryan Catanzaro. Improving semantic segmentation via video propagation and label relaxation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8856–8865, 2019.
- [54] Bohan Zhuang, Guosheng Lin, Chunhua Shen, and Ian Reid. Fast training of triplet-based deep binary embedding networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5955–5964, 2016.