

VSRD: Instance-Aware Volumetric Silhouette Rendering for Weakly Supervised 3D Object Detection

Zihua Liu*
 Tokyo Institute of Technology
 zliu@ok.sc.e.titech.ac.jp

Hiroki Sakuma*
 T2 Inc.
 sakuma.h@t2.auto

Masatoshi Okutomi
 Tokyo Institute of Technology
 mxo@sc.e.titech.ac.jp

Abstract

Monocular 3D object detection poses a significant challenge in 3D scene understanding due to its inherently ill-posed nature in monocular depth estimation. Existing methods heavily rely on supervised learning using abundant 3D labels, typically obtained through expensive and labor-intensive annotation on LiDAR point clouds. To tackle this problem, we propose a novel weakly supervised 3D object detection framework named VSRD (Volumetric Silhouette Rendering for Detection) to train 3D object detectors without any 3D supervision but only weak 2D supervision. VSRD consists of multi-view 3D auto-labeling and subsequent training of monocular 3D object detectors using the pseudo labels generated in the auto-labeling stage. In the auto-labeling stage, we represent the surface of each instance as a signed distance field (SDF) and render its silhouette as an instance mask through our proposed instance-aware volumetric silhouette rendering. To directly optimize the 3D bounding boxes through rendering, we decompose the SDF of each instance into the SDF of a cuboid and the residual distance field (RDF) that represents the residual from the cuboid. This mechanism enables us to optimize the 3D bounding boxes in an end-to-end manner by comparing the rendered instance masks with the ground truth instance masks. The optimized 3D bounding boxes serve as effective training data for 3D object detection. We conduct extensive experiments on the KITTI-360 dataset, demonstrating that our method outperforms the existing weakly supervised 3D object detection methods. The code is available at <https://github.com/skmhrk1209/VSRD>.

1. Introduction

3D object detection is one of the most critical components in the perception system for autonomous driving. With the recent success of deep learning in computer vision, numerous 3D object detection methods have been proposed to

*Equal contribution. The order was determined by a coin flip.




Figure 1. Illustration of our proposed weakly supervised 3D object detection framework, which consists of multi-view 3D auto-labeling and subsequent training of monocular 3D object detectors using the pseudo labels generated in the auto-labeling stage.

directly regress 3D bounding boxes from a LiDAR point cloud, multi-view images, or a monocular image. Among them, monocular 3D object detection is the most challenging in principle due to its inherently ill-posed nature in monocular depth estimation. Therefore, existing methods [22, 25, 28, 34, 35] heavily rely on supervised learning using abundant 3D labels manually annotated on LiDAR point clouds. This annotation cost is extremely high, posing a significant barrier to deploying 3D object detectors into autonomous driving systems.

To tackle this problem, we propose a novel weakly supervised 3D object detection framework named VSRD (Volumetric Silhouette Rendering for Detection) to train 3D object detectors without any 3D supervision but only weak 2D supervision. As illustrated in Fig. 1, VSRD consists of multi-view 3D auto-labeling and subsequent training of monocular 3D object detectors using the pseudo labels generated in the auto-labeling stage. In the auto-labeling stage, we represent the surface of each 3D bounding box as a signed distance field (SDF) and render its silhouette as an instance mask through volumetric rendering. Comparing the rendered instance masks with the ground truth instance masks enables us to optimize the 3D bounding boxes directly. We introduce two novel mechanisms in this auto-labeling stage. The first is the instance-aware volumetric silhouette rendering that integrates instance labels along a ray to render the silhouette of each instance rather

than the entire scene. This mechanism enables the silhouette of each instance to be rendered while considering geometric relationships among instances, such as occlusion. The second is the SDF decomposition, whereby the SDF of each instance is decomposed into the SDF of a cuboid and the *residual distance field* (RDF) that represents the residual from the cuboid. This decomposition models the spatial gap between the surfaces of each instance and the 3D bounding box, enabling more accurate silhouette rendering and providing more reliable feedback signals during optimization.

The 3D bounding boxes optimized by the proposed auto-labeling can serve as pseudo labels for training 3D object detectors. However, dynamic objects and inaccurate camera poses lead to low-quality pseudo labels, which negatively impact the training of 3D object detectors. Therefore, we propose a simple but effective algorithm to assign a confidence score representing the label quality to each pseudo label. We demonstrate that using these confidence scores as per-instance loss weights boosts the performance of the 3D object detectors trained on the pseudo labels.

In summary, our main contributions are as follows:

- We propose a weakly supervised 3D object detection framework consisting of multi-view 3D auto-labeling and subsequent training of monocular 3D object detectors.
- We propose a novel instance-aware volumetric silhouette rendering method whereby the silhouette of each instance can be rendered as an instance mask.
- We propose a novel SDF decomposition whereby the SDF of each instance is decomposed into the SDF of a cuboid and the residual distance field (RDF).
- We propose a simple but effective confidence assignment algorithm to incorporate the quality of each pseudo label into the training of 3D object detectors.
- Extensive experiments on the KITTI-360 dataset demonstrate that our method outperforms the existing weakly supervised 3D object detection methods.

2. Related Work

2.1. Monocular 3D Object Detection

Monocular 3D detection is a challenging task due to limited 3D information from monocular imagery. Deep3DBox [19] pioneered this area by regressing relatively stable 3D object properties and combining these estimates with geometric constraints provided by the 2D bounding box. SMOKE [16] estimates 3D bounding boxes by combining keypoint estimates with regressed 3D box parameters. FCOS3D [28] employs a fully convolutional single-stage detector, transforming 7-DoF 3D targets to the image domain and decoupling them as 2D and 3D attributes. To enhance the performance of monocular 3D object detection for truncated objects, MonoFlex [35] explicitly decouples truncated objects and adaptively combines multiple approaches for ob-

ject depth estimation. M3D-RPN [4] and MonoDETR [34] also explored the usage of depth cues to improve monocular 3D object detection. The former designed depth-aware convolutional layers that enable location-specific feature extraction and consequently improved 3D scene understanding, while the latter modified the vanilla Transformer [7] to be depth-aware to guide the whole detection process by contextual depth cues. Despite these advances, the reliance on expensive and labor-intensive manual annotation on LiDAR point clouds remains a significant limitation.

2.2. Weakly Supervised 3D Object Detection

Many weakly supervised methods have been proposed to mitigate the high annotation cost in 3D object detection. WS3D [17] introduced a LiDAR-based two-stage pipeline where cylindrical object proposals are first generated under weak supervision and then refined using a few labeled object instances. VS3D [24] introduced an unsupervised 3D proposal module that generates object proposals by leveraging normalized point cloud densities. WeakM3D [23] introduced a weakly supervised monocular 3D object detection method that leverages the 3D alignment loss between each predicted 3D bounding box and corresponding RoI LiDAR points. Furthermore, it introduced a method to estimate the orientation from RoI LiDAR points based on its statistics. Recently, WeakMono3D [26] leverages projection loss with multi-view and direction consistency, achieving a weakly supervised monocular object detection that relies only on 2D supervision. However, its reliance on 2D direction annotations restricts its applicability to large-scale datasets. Zakharov et al. [33] proposed an auto-labeling pipeline that integrates an SDF-based differentiable shape renderer and normalized object coordinate spaces (NOCS). However, additional training on synthetic data is still required for shape and coordinate estimation. In contrast, our method purely relies on 2D supervision, eliminating the necessity of synthetic data or 3D supervision.

2.3. 3D Object Detection with Neural Fields

Neural Radiance Fields (NeRF) [18] introduced a new perspective for implicit learning of 3D geometry from posed multi-view images by volume rendering. Building upon the vanilla NeRF, subsequent research has focused on enhancing novel view synthesis [1–3, 21] or accelerating volume rendering [6, 8, 20]. NeuS [27] introduced a novel volume rendering scheme to learn a neural SDF representation by introducing a density distribution induced by the SDF. Similarly, VolSDF [32] defined the volume density function as Laplace’s cumulative distribution function applied to an SDF representation. Many works [12, 30, 31] recently have attempted to utilize neural fields for 3D object detection. Notably, NeRF-RPN [12] demonstrated that the 3D bounding boxes of objects in NeRF can be directly regressed with-

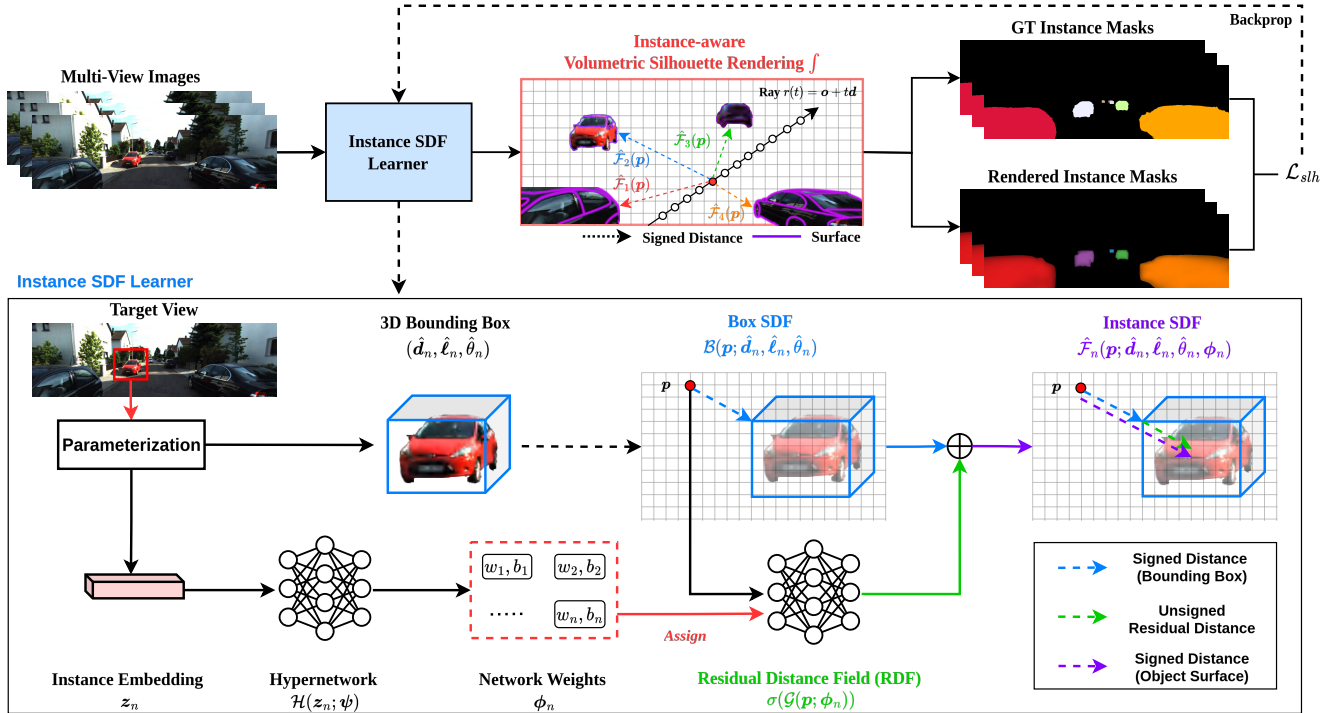


Figure 2. Illustration of the pipeline of our proposed multi-view 3D auto-labeling. We represent the surface of each instance as an SDF and decompose it into the SDF of a 3D bounding box and the *residual distance field* (RDF), which is learned via a hypernetwork. The composed instance SDF is used to render the silhouette of the instance through our proposed instance-aware volumetric silhouette rendering. All the 3D bounding boxes are optimized based on the loss between the rendered and ground truth instance masks.

out rendering. NeRF-Det [30] connects the detection and NeRF branches through a shared MLP, enabling an efficient adaptation of NeRF to detection and yielding geometry-aware volumetric representations. MonoNeRD [31] models scenes with SDFs and renders RGB images and depth maps through volume rendering to obtain intermediate 3D representations for detection. However, these approaches still rely on ground truth 3D labels for supervision. In contrast, we propose the first volume rendering-based weakly supervised 3D object detection framework that relies on multi-view geometry and 2D supervision without any 3D supervision, such as 3D bounding boxes or LiDAR point clouds.

3. Method

3.1. Multi-View 3D Auto-Labeling

The pipeline of our proposed multi-view 3D auto-labeling is illustrated in Fig. 2. In Sec. 3.1.1, we review the SDF-based volume rendering formulation, which is the foundation of our method. In Sec. 3.1.2, we define the optimization problem for the multi-view 3D auto-labeling. In Sec. 3.1.3, we introduce the SDF for cuboids, which is used to optimize 3D bounding boxes through rendering. In Sec. 3.1.4, we introduce a novel neural field named *residual distance*

field to fill the spatial gap between the surfaces of each instance and the 3D bounding box. In Sec. 3.1.5, we introduce instance-aware volumetric silhouette rendering to render the silhouette of each instance based on its SDF. In Sec. 3.1.6, we introduce loss functions using instance masks and 2D bounding boxes as weak supervision.

3.1.1 Preliminaries

SDF-based Volumetric Rendering NeRF [18] represents a 3D scene with neural density and color fields. Given a camera position $\mathbf{o} \in \mathbb{R}^3$ and a ray direction $\mathbf{d} \in \mathbb{R}^3$ emitted from a pixel, the volume rendering scheme integrates the colors of sampled points along the ray as follows:

$$\hat{C}(\mathbf{o}, \mathbf{d}) = \int_0^\infty w(t) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt, \quad (1)$$

$$w(t) = \exp\left(-\int_0^t \sigma(\mathbf{r}(u)) du\right) \sigma(\mathbf{r}(t)), \quad (2)$$

where $\hat{C}(\mathbf{o}, \mathbf{d}) \in \mathbb{R}^3$ denotes the rendered color of the pixel, $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ denotes the ray, $\mathbf{c}(\mathbf{p}, \mathbf{d})$ denotes the color at the position \mathbf{p} and view direction \mathbf{d} , and $\sigma(\mathbf{p})$ denotes the *volume density* at the position \mathbf{p} . Since the density field cannot represent the surfaces explicitly, novel SDF-

based volume rendering formulation has been introduced in NeuS [27], where a surface represented by an SDF is reinterpreted as a participating medium represented by a density field, enabling the surface to be rendered through volume rendering. In this formulation, the weight $w(t)$ in Eq. (2) is re-written by introducing *opaque density* $\rho(t)$ as follows:

$$w(t) = \exp\left(-\int_0^t \rho(u)du\right)\rho(t), \quad (3)$$

$$\rho(t) = \max\left(-\frac{d\Phi}{dt}(\hat{\mathcal{F}}(\mathbf{r}(t))), 0\right), \quad (4)$$

where $\hat{\mathcal{F}}(\cdot)$ denotes the SDF for the entire scene and $\Phi(\cdot)$ denotes the Sigmoid function. Our proposed instance-aware volumetric silhouette rendering introduced in Sec. 3.1.5 is based on the same SDF-based weight formulation as Eq. (3) but integrates instance labels instead of colors along a ray to render instance masks.

3.1.2 Problem Definition

Given a monocular video consisting of posed frames, each frame annotated with instance masks, our goal is to optimize the 3D bounding box frame by frame without 3D supervision. More specifically, for each *target* frame t in the video, we sample multiple *source* frames \mathcal{S} and optimize the N 3D bounding boxes in the target frame using the instance masks of the source frames as weak supervision, where N denotes the number of instances in the target frame. Please refer to the supplementary material for how to sample \mathcal{S} . We parameterize the n -th 3D bounding box $\hat{\mathbf{B}}_n \in \mathbb{R}^{8 \times 3}$ in the target frame with a dimension $\hat{\mathbf{d}}_n \in \mathbb{R}_+^3$, location $\hat{\ell}_n \in \mathbb{R}^3$, and orientation $\hat{\theta}_n \in \mathbb{R}$, which is the rotation angle in the bird’s-eye-view. In addition to these parameters for each bounding box, we prepare a learnable instance embedding $\mathbf{z}_n \in \mathbb{R}^D$ for each instance and a shared *hypernetwork* parameterized by ψ for the *residual distance field* introduced in Sec. 3.1.4. We stack each parameter group into a single tensor over all the instances, yielding dimensions $\hat{\mathbf{D}} \in \mathbb{R}_+^{N \times 3}$, locations $\hat{\mathbf{L}} \in \mathbb{R}^{N \times 3}$, orientations $\hat{\Theta} \in \mathbb{R}^{N \times 1}$, and instance embeddings $\mathbf{Z} \in \mathbb{R}^{N \times D}$. Given the loss function \mathcal{L} explained in Sec. 3.1.6, we optimize $\hat{\mathbf{D}}$, $\hat{\mathbf{L}}$, $\hat{\Theta}$, \mathbf{Z} , and ψ via stochastic gradient descent as follows:

$${}^* \mathbf{D}, {}^* \mathbf{L}, {}^* \Theta, {}^* \mathbf{Z}, \psi = \underset{\hat{\mathbf{D}}, \hat{\mathbf{L}}, \hat{\Theta}, \mathbf{Z}, \psi}{\operatorname{argmin}} \mathcal{L}(\hat{\mathbf{D}}, \hat{\mathbf{L}}, \hat{\Theta}, \mathbf{Z}, \psi). \quad (5)$$

The optimized 3D bounding boxes ${}^* \mathbf{B}$ decoded from ${}^* \mathbf{D}$, ${}^* \mathbf{L}$, and ${}^* \Theta$ can be used as pseudo labels for training 3D object detectors, as explained in Sec. 3.2.

3.1.3 3D Bounding Box Represented as an SDF

To optimize the 3D bounding boxes in each target frame by solving Eq. (5) through rendering, we represent the surface of each 3D bounding box as a signed distance field (SDF), which is one of the most common surface representations. An SDF is represented as a function $\mathcal{F} : \mathbb{R}^3 \rightarrow \mathbb{R}$ that maps a spatial position $\mathbf{p} \in \mathbb{R}^3$ to its signed distance to the closest point on the surface, indicating that the zero-level set $\{\mathbf{p} \in \mathbb{R}^3 \mid \mathcal{F}(\mathbf{p}) = 0\}$ represents the surface itself. The SDF for a cuboid parameterized by a dimension $\mathbf{d} \in \mathbb{R}_+^3$, location $\ell \in \mathbb{R}^3$, and orientation $\mathbf{R} \in \text{SO}(3)$ can be derived theoretically and is denoted by $\mathcal{B}(\cdot; \mathbf{d}, \ell, \mathbf{R})$. Please refer to the supplementary material for how to derive this formula.

3.1.4 Residual Distance Field

In general, the shape of each instance is not a cuboid. Therefore, if we leverage only the cuboid SDF introduced in Sec. 3.1.3 to render the surface of each instance, it is no longer possible to render an accurate silhouette due to the spatial gap between the surfaces of the instance and the 3D bounding box, leading to unreliable feedback signals during optimization. Therefore, we propose a novel neural field named *residual distance field* (RDF), which models the residual between the signed distances to the surfaces of the instance and the 3D bounding box. Let $\mathcal{F}_n(\cdot)$ be the *true* SDF of the surface bounded by the n -th 3D bounding box $\hat{\mathbf{B}}_n$ whose SDF is given by $\hat{\mathcal{B}}_n(\cdot) = \mathcal{B}(\cdot; \hat{\mathbf{d}}_n, \hat{\ell}_n, \mathbf{R}_y(\hat{\theta}_n))$, where $\mathbf{R}_y(\theta)$ denotes the rotation matrix around the y -axis by an angle θ . For any point $\mathbf{p} \in \mathbb{R}^3$, we define the RDF as $\hat{\mathcal{R}}_n(\mathbf{p}) := \mathcal{F}_n(\mathbf{p}) - \hat{\mathcal{B}}_n(\mathbf{p})$. Here, based on the definition of a 3D bounding box that it encloses the corresponding instance, $\forall \mathbf{p} \in \mathbb{R}^3 : \hat{\mathcal{R}}_n(\mathbf{p}) \geq 0$ is required. One straightforward way to model the RDF for each instance is to train N individual networks. However, as objects belonging to the same semantic class often have similar geometric shapes, we train a single *hypernetwork* [10] that regresses the weights of the neural RDF directly from an instance embedding instead. Given the n -th instance embedding $\mathbf{z}_n \in \mathbb{R}^D$, where D denotes the number of dimensions, the n -th neural RDF is given by:

$$\hat{\mathcal{R}}_n(\mathbf{p}) = \sigma(\mathcal{G}(\mathbf{p}; \phi_n)), \quad (6)$$

$$\phi_n = \mathcal{H}(\mathbf{z}_n; \psi), \quad (7)$$

where $\mathcal{G}(\cdot; \phi_n)$ denotes the n -th neural network parameterized by ϕ_n , $\mathcal{H}(\cdot; \psi)$ denotes the shared hypernetwork parameterized by ψ , and $\sigma(\cdot)$ denotes the Softplus function to force the residual distance to be positive based on the definition of a bounding box. For the instance-aware volumetric silhouette rendering introduced in Sec. 3.1.5, we employ $\hat{\mathcal{F}}_n(\mathbf{p}) = \hat{\mathcal{B}}_n(\mathbf{p}) + \hat{\mathcal{R}}_n(\mathbf{p})$ as the SDF of each instance that represents an arbitrary surface bounded by the 3D bounding box $\hat{\mathbf{B}}_n$ for more accurate silhouette rendering.

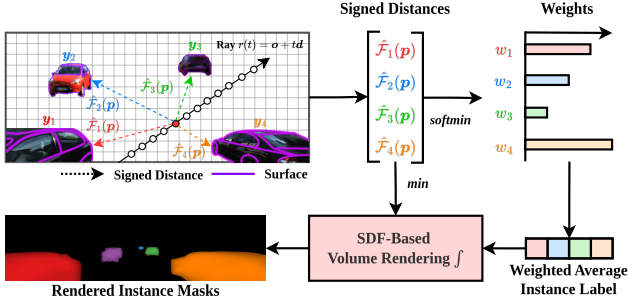


Figure 3. Illustration of our proposed instance-aware volumetric silhouette rendering. The instance labels are averaged for each sampled point along a ray based on the signed distance to each instance. The averaged instance labels are integrated along the ray based on the SDF-based volume rendering formulation [27].

3.1.5 Instance-Aware Volumetric Silhouette Rendering

This section details our approach to optimizing 3D bounding boxes through volumetric rendering. The core idea is to render instance masks and compare them with ground truth instance masks. To achieve this, we propose a novel SDF-based *instance-aware* volumetric silhouette rendering, where instance labels instead of colors are integrated along a ray based on the same SDF-based volume rendering formulation as Eq. (3) as follows:

$$\hat{S}(\mathbf{o}, \mathbf{d}) = \int_0^\infty w(t) \mathbf{s}(\mathbf{r}(t)) dt, \quad (8)$$

$$\mathbf{s}(\mathbf{p}) = \sum_{n=1}^N \text{softmin}([\hat{\mathcal{F}}_n(\mathbf{p})]_{n=1}^N) \cdot \mathbf{y}_n, \quad (9)$$

where, $\hat{S}(\mathbf{o}, \mathbf{d}) \in [0, 1]^N$ denotes the rendered soft instance label, $\mathbf{s}(\mathbf{p}) \in [0, 1]^N$ denotes the weighted average instance label at the position \mathbf{p} indicating how relatively close the position \mathbf{p} is to each instance, and $\mathbf{y}_n \in \{0, 1\}^N$ denotes the one-hot instance label of the n -th instance. $\hat{\mathcal{F}}_n(\cdot)$ denotes the SDF of the n -th instance introduced in Sec. 3.1.4. $w(t)$ in Eq. (8) is derived from Eq. (3). To compute Eq. (3), we model the entire scene as the union of all the object surfaces, i.e., $\hat{\mathcal{F}}(\mathbf{p}) = \min(\hat{\mathcal{F}}_1(\mathbf{p}), \dots, \hat{\mathcal{F}}_N(\mathbf{p}))$. This mechanism enables us to render instance masks considering the geometric relationships among instances, such as occlusions.

3.1.6 Loss Functions

We optimize the parameters of the 3D bounding boxes $\hat{\Omega} = \{\hat{D}, \hat{L}, \hat{\Theta}\}$ in the target frame along with the instance embeddings \mathbf{Z} and parameter ψ of the hypernetwork $\mathcal{H}(\cdot; \psi)$. The final loss \mathcal{L} is defined as the combination of the *multi-view projection loss* $\mathcal{L}_{\text{proj}}$, *multi-view silhouette loss* \mathcal{L}_{slh} , and *Eikonal regularization* \mathcal{L}_{reg} [27] as follows:

$$\begin{aligned} \mathcal{L}(\hat{\Omega}, \mathbf{Z}, \psi) = & \lambda_{\text{proj}} \mathcal{L}_{\text{proj}}(\hat{\Omega}) + \lambda_{\text{slh}} \mathcal{L}_{\text{slh}}(\hat{\Omega}, \mathbf{Z}, \psi) \\ & + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}(\hat{\Omega}, \mathbf{Z}, \psi), \end{aligned} \quad (10)$$

where λ_{proj} , λ_{slh} , and λ_{reg} denote the loss weights. However, although the multi-view projection loss $\mathcal{L}_{\text{proj}}$ and multi-view silhouette loss \mathcal{L}_{slh} are based on 2D supervision, the 3D-2D correspondence to define the losses is not obvious. Therefore, we find an optimal bipartite matching between the optimized 3D bounding boxes and ground truth 2D bounding boxes by the Hungarian algorithm [14]. Each pair-wise matching cost is defined as the projection loss introduced in this section, but it is computed for not all the source frames but only the target frame. We assume that the ground truth 2D bounding boxes and instance masks have already been reordered based on the optimal permutation for simplicity.

Multi-View Projection Loss Although introducing the residual distance field enables us to model an arbitrary surface bounded by a 3D bounding box, it introduces another problem: the 3D bounding box can grow indefinitely without any constraint. To address this problem, we employ the ground truth 2D bounding box as a constraint to keep the 3D bounding box tight. The multi-view projection loss $\mathcal{L}_{\text{proj}}$ is defined as the average distance between the *projected* and ground truth 2D bounding boxes as follows:

$$\begin{aligned} \mathcal{L}_{\text{proj}}(\hat{\Omega}) = & \alpha \sum_{i \in \mathcal{S}} \sum_{n=1}^N \|\hat{B}_{in}^{2D} - B_{in}^{2D}\|_{\text{H}} \\ & - \beta \sum_{i \in \mathcal{S}} \sum_{n=1}^N \text{DIoU}(\hat{B}_{in}^{2D}, B_{in}^{2D}), \end{aligned} \quad (11)$$

where $\|\cdot\|_{\text{H}}$ denotes the Huber loss, $\text{DIoU}(\cdot, \cdot)$ denotes the Distance-IoU [36], and α, β denote balancing coefficients. The projected 2D bounding box \hat{B}_{in}^{2D} is defined as the rectangle with minimal area enclosing the projected vertices $\hat{V}_{in}^{2D} \propto \hat{B}_n \mathbf{E}_i^T \mathbf{K}_i^T$, where \mathbf{E}_i and \mathbf{K}_i denote the extrinsic and intrinsic matrices for frame i , respectively.

Multi-View Silhouette Loss The multi-view silhouette loss is defined as the cross entropy between the rendered and ground truth instance masks. Since the spatial gap between the surfaces of each instance and the 3D bounding box is modeled by the residual distance field introduced in Sec. 3.1.4, the 3D bounding box constrained to tightly fit the ground truth 2D bounding box by the multi-view projection loss $\mathcal{L}_{\text{proj}}$ is further refined by the multi-view silhouette loss \mathcal{L}_{slh} , which is given by:

$$\mathcal{L}_{\text{slh}}(\hat{\Omega}, \mathbf{Z}, \psi) = \sum_{i \in \mathcal{S}} \sum_{j=1}^{R_i} \text{CE}(\hat{S}(\mathbf{o}_i, \mathbf{d}_{ij}), \mathbf{S}_{ij}), \quad (12)$$




Figure 4. Comparison of the confidence scores in static and dynamic scenes. It can be seen that the confidence scores are lower for dynamic, occluded, or truncated objects, indicating less influence on the subsequent training of 3D object detectors.

where $\mathbf{o}_i \in \mathbb{R}^3$, $\mathbf{d}_{ij} \in \mathbb{R}^3$, and $\mathbf{S}_{ij} \in \{0, 1\}^N$ denote the camera position, ray direction, and ground truth instance label at the j -th sampled pixel in frame i , respectively. $\hat{\mathbf{S}}(\mathbf{o}_i, \mathbf{d}_{ij})$ denotes the rendered instance label based on Eq. (8). $\text{CE}(\cdot, \cdot)$ denotes the cross entropy loss. R_i denotes the number of rays sampled for frame i . Please refer to the supplementary material for how to determine R_i .

3.2. Training of 3D Object Detectors

Once the 3D bounding boxes are optimized by the proposed multi-view 3D auto-labeling, they can serve as pseudo labels for training 3D object detectors. To incorporate the quality of each pseudo label into the training of 3D object detectors, we introduce a simple but effective confidence assignment method in Sec. 3.2.1 and confidence-based weighted loss for bounding box regression in Sec. 3.2.2.

3.2.1 Confidence Assignment

As shown in Fig. 4, the 3D bounding boxes optimized by the proposed auto-labeling are not reliable for dynamic, occluded, or truncated objects. Therefore, we propose a confidence assignment method based on the multi-view projection loss. First, for each target frame t , we identify a set of source frames \mathcal{I} such that all the instances in the target frame are visible from every frame in the set. Next, we find an optimal bipartite matching between the optimized 3D bounding boxes and ground truth 2D bounding boxes by the Hungarian algorithm. The cost matrix $\mathbf{Q} \in \mathbb{R}^{N \times N}$ is defined as the pairwise IoUs between the projected and ground truth 2D bounding boxes averaged over all the source frames, as follows:

$$\mathbf{Q}_{nm} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} 1 - \text{IoU}(*\mathbf{B}_{in}^{2D}, \mathbf{B}_{im}^{2D}), \quad (13)$$

where $*\mathbf{B}_{in}^{2D}$ denotes the n -th projected 2D bounding box in frame i , which is obtained by projecting the n -th optimized 3D bounding box $*\mathbf{B}_n$ onto frame i , and \mathbf{B}_{im}^{2D} denotes the m -th ground truth 2D bounding box in frame i . Once an optimal permutation matrix $\mathbf{P} \in \{0, 1\}^{N \times N}$ is obtained, the confidence scores $*\mathbf{C} \in [0, 1]^N$ for the optimized 3D

bounding boxes $*\mathbf{B}$ are given as follows:

$$*\mathbf{C} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \text{IoU}(*\mathbf{B}_i^{2D}, \mathbf{P}\mathbf{B}_i^{2D}). \quad (14)$$

3.2.2 Confidence-based Weighted Loss

The 3D bounding boxes optimized by the proposed multi-view 3D auto-labeling serve as pseudo labels for most 3D object detectors, which leverage 3D bounding boxes and semantic class labels as supervision, without any modification of the architectures, loss functions, and training procedures, except for the confidence-based loss weighting. We incorporate the confidence scores computed by Eq. (14) into only the regression loss. Given a regression loss function \mathcal{L}_{box} , its confidence-based weighted version $\tilde{\mathcal{L}}_{\text{box}}$ is given by:

$$\tilde{\mathcal{L}}_{\text{box}}(\hat{\mathbf{B}}, *\mathbf{B}, *\mathbf{C}) = \sum_{m=1}^M *\mathbf{C}_{\pi(m)} \mathcal{L}_{\text{box}}(\hat{\mathbf{B}}_m, *\mathbf{B}_{\pi(m)}), \quad (15)$$

where $\hat{\mathbf{B}}$ denotes the predicted 3D bounding boxes, $*\mathbf{B}$ denotes the 3D bounding boxes optimized by the proposed auto-labeling, and $*\mathbf{C}$ denotes the corresponding confidence scores. M denotes the number of positive anchors, and $\pi(\cdot)$ denotes the label assigner that maps the index of an anchor to that of the matched ground truth.

4. Experiments

4.1. Dataset

We use the KITTI-360 [15] dataset for our experiments, splitting it into training (43,855 images), validation (1,173 images), and test sets (2,531 images). We follow the same evaluation protocol as the KITTI dataset [9]. However, as occlusion and truncation labels are not available for the KITTI-360 dataset, unlike the KITTI dataset, we consider only two difficulty levels, namely *Easy* and *Hard*, based on whether the height of each ground truth 2D bounding box is greater than 40 and 25 in pixels, respectively. Following prior works, we evaluate our method on only category *Car*.

4.2. Implementation Details

4.2.1 Multi-View 3D Auto-Labeling

We sample 16 source frames for each target frame. We sample 1000 rays across all the source frames, i.e., $\sum_{i \in \mathcal{S}} R_i = 1000$, at each iteration based on the ground truth instance masks. Please refer to the supplementary material for more details. We employ the same hierarchical volume sampling as NeRF [18] and sample 100 query points for both *coarse* and *fine* sampling. Unlike NeRF, both coarse and fine samples are drawn from the single scene SDF rather than two distinct ones. We set the number of dimensions of each instance embedding as $D = 256$. The neural RDF \mathcal{G} and hypernetwork \mathcal{H} are implemented as MLPs with four hidden

Table 1. Ablation study on the KITTI-360 training set to verify the effectiveness of each component in our proposed multi-view 3D auto-labeling. Just one randomly selected sequence is used.

Components			AP _{BEV} /AP _{3D} @0.3		AP _{BEV} /AP _{3D} @0.5@0.5	
$\mathcal{L}_{\text{proj}}$	\mathcal{L}_{slh}	RDF	Easy	Hard	Easy	Hard
✓			60.77/54.88	63.99/57.66	37.38/23.33	37.44/24.82
✓	✓		63.84/56.11	60.86/56.87	41.73/26.84	39.22/25.58
✓	✓	✓	73.84/66.64	73.22/66.32	46.35/31.11	43.07/30.16

Table 2. Ablation study on the KITTI-360 test set to verify the effectiveness of the confidence scores against monocular 3D object detection. S-WeakM3D is used as a monocular 3D object detector.

Conf.	AP _{BEV} /AP _{3D} @0.3		AP _{BEV} /AP _{3D} @0.5	
	Easy	Hard	Easy	Hard
	38.67/31.62	31.25/23.88	10.95/8.25	5.37/4.41
✓	51.09/42.94	41.27/33.97	19.50/11.91	14.39/8.46

layers, each of which has 256 and 16 channels, respectively. We use the Adam optimizer [13] and the learning rates are decayed exponentially from $1e^{-2}$, $1e^{-3}$, and $1e^{-4}$ to $1e^{-4}$, $1e^{-5}$, and $1e^{-6}$ for the box parameters $\hat{\Omega}$, instance embeddings \mathbf{Z} , and parameter ψ of the hypernetwork $\mathcal{H}(\cdot; \psi)$ over 3000 iterations, respectively. For the loss weights, we set $\alpha = 1.0$, $\beta = 0.1$, $\lambda_{\text{proj}} = 1.0$, $\lambda_{\text{slh}} = 1.0$, and $\lambda_{\text{reg}} = 0.01$.

4.2.2 Monocular 3D Object Detection

To compare our method and Autolabels [33] with WeakM3D [23], we modify the architecture of WeakM3D so that it can be trained in a supervised manner using pseudo labels. More specifically, we train dimension and confidence heads in addition to the existing location and orientation heads using the same supervised loss as MonoDIS [25]. We call this model S-WeakM3D. For both WeakM3D and S-WeakM3D, ground truth 2D bounding boxes and instance masks are used for RoIAlign [11] and generating LiDAR points on each object surface during training, respectively. During inference, we employ Cascade Mask R-CNN [5] with InternImage-XL [29] as the off-the-shelf 2D detector to provide 2D bounding boxes for RoIAlign. Please refer to the supplementary material for more details.

4.3. Ablation Study

4.3.1 Multi-View 3D Auto-Labeling

We conduct an ablation study to demonstrate the effectiveness of each component in our proposed multi-view 3D auto-labeling. As can be seen from Tab. 1, the multi-view projection loss $\mathcal{L}_{\text{proj}}$ serves as a practical baseline on its own. The multi-view silhouette loss \mathcal{L}_{slh} further boosts the quality of the pseudo labels a little bit, while the spatial gap




Figure 5. Visualization results of the optimized 3D bounding boxes (1st row) and rendered instance masks (2nd row). We assign a unique color to each instance, and each pixel is colored as the weighted summation based on the rendered soft instance label.

Table 3. Evaluation results of our proposed multi-view 3D auto-labeling on the KITTI-360 training set compared with the LiDAR-based monocular 3D auto-labeling proposed in Autolabels [33]. *Reproduced with the official code.

Conf.	Method	AP _{BEV} /AP _{3D} @0.3		AP _{BEV} /AP _{3D} @0.5	
		Easy	Hard	Easy	Hard
≥ 0.0	Autolabels*	71.24/15.09	67.33/11.42	51.85/4.65	46.10/2.92
	VSRD	75.03/68.53	72.11/65.64	47.12/35.25	43.91/32.64
≥ 0.8	Autolabels*	75.56/17.18	72.23/12.19	58.64/5.77	52.99/3.75
	VSRD	84.54/80.25	81.66/77.37	58.57/45.76	55.09/44.17

between the surfaces of each instance and the 3D bounding box limits further improvements. However, the residual distance field (RDF) can significantly boost the quality of the pseudo labels by addressing this problem. The systematic improvements with the inclusion of each component indicate their effectiveness, resulting in a precise auto-labeling system for weakly supervised 3D object detection.

4.3.2 Confidence Assignment

Due to the unreliability of the 3D bounding boxes optimized by the proposed auto-labeling for dynamic, occluded, or truncated objects, we conduct an ablation study to demonstrate the effectiveness of our proposed confidence assignment. Tab. 2 highlights substantial enhancement in detection performance when using the confidence-incorporated pseudo labels compared with the baseline, demonstrating the effectiveness of the proposed confidence assignment.

4.4. Evaluation Results

4.4.1 Multi-View 3D Auto-Labeling

Autolabels [33] employs a similar two-stage framework consisting of auto-labeling and subsequent training of 3D object detectors using the pseudo labels. We compare our method with Autolabels to evaluate the quality of the pseudo labels. As our pseudo labels are unreliable for dy-

Table 4. Evaluation results of monocular 3D object detection on the KITTI-360 test set. *Reproduced with the official code. †CAD models are used as extra data. ‡ $M(D)$ indicates that detection model D is employed for model-agnostic method M .

Method	Weak Supervision		Full Supervision	AP _{BEV} /AP _{3D} @0.3		AP _{BEV} /AP _{3D} @0.5	
	LiDAR	Masks	3D Boxes	Easy	Hard	Easy	Hard
WeakM3D* [23]	✓	✓		29.89/21.25	24.01/15.34	8.10/2.96	2.96/2.01
Autolabels*†‡ [33] (S-WeakM3D)	✓	✓		48.16/12.92	37.34/9.94	20.18/4.69	14.33/2.79
VSRD‡ (S-WeakM3D)		✓		51.09/42.94	41.28/33.78	19.51/11.91	14.39/8.46
VSRD‡ (MonoFlex)		✓		54.40/48.16	45.67/40.04	29.18/18.53	22.31/13.60
VSRD‡ (MonoDETR)		✓		58.40/50.86	50.61/43.45	29.07/ 21.77	22.83/16.46
MonoFlex* [35]			✓	69.70/67.07	59.86/57.26	50.82/43.11	41.78/34.43
MonoDETR* [34]			✓	63.07/60.49	54.04/50.03	47.21/41.01	36.05/30.38

Table 5. Evaluation results of semi-supervised monocular 3D object detection on the KITTI validation set.

Method	Ratio	AP _{BEV} /AP _{3D} @0.7		
		Easy	Moderate	Hard
MonoDETR [34]	1.00	37.99/29.36	26.76/20.64	23.02/17.30
	0.00	0.002/0.001	0.004/0.001	0.005/0.002
	0.25	31.72/21.76	22.32/15.43	18.86/12.55
VSRD (MonoDETR)	0.50	43.44/31.05	31.54/21.48	27.17/17.93
	0.75	42.58/32.95	31.08/24.68	27.19/21.38

dynamic, occluded, or truncated objects, we group the dataset based on whether the average confidence score for each image is greater than a certain threshold. The evaluation results with the confidence thresholds of 0.0 and 0.8 are shown in Tab. 3. Our method exhibits superior performance in terms of AP_{BEV}@0.3 and AP_{3D}. For AP_{BEV}@0.5, our method exhibits slightly lower performance for the confidence threshold of 0.0. However, by raising the confidence threshold to 0.8, our method demonstrates superior performance, indicating that our method can generate more high-quality pseudo labels while abandoning low-quality pseudo labels. The optimized 3D bounding boxes and rendered in-stance masks are visualized in Fig. 5.

4.4.2 Monocular 3D Object Detection

The pseudo labels generated by the proposed auto-labeling serve as 3D supervision. We further investigate their applicability using the existing monocular 3D object detectors.

Weakly Supervised Setting Tab. 4 shows the evaluation results of our method compared with the existing weakly supervised and fully supervised methods. Our method demonstrates a significant superiority over WeakM3D [23] across all the metrics while eliminating the need for LiDAR points for 3D supervision. Moreover, the detector trained on the pseudo labels generated by the proposed auto-labeling outperforms that trained on the pseudo labels generated by Autolabels [33]. Furthermore, employing more sophisticated

monocular 3D object detectors such as MonoFlex [35] and MonoDETR [34] further improves detection performance, demonstrating the broad versatility of our method, which is not limited to a specific detection model. It is noteworthy that the detectors trained on the pseudo labels generated by our method demonstrate competitive performance compared with those trained in a fully supervised manner.

Semi-Supervised Setting The essential advantage of our method is that it avoids costly 3D annotations, making more data available for training. Therefore, we investigate a realistic scenario where a detector pre-trained on a large amount of unlabeled data from a source domain is fine-tuned on a small amount of labeled data from a target domain. We select the KITTI-360 and KITTI datasets as source and target domains, respectively. Tab. 5 shows the performance of the detector pre-trained on the KITTI-360 dataset in a weakly supervised manner with the proposed auto-labeling and then fine-tuned on a subset of the KITTI dataset in a supervised manner. The zero-shot performance is quite low due to the characteristic that monocular depth estimation is greatly affected by the differences in camera parameters, but the performance of the detector fine-tuned on only 50% of the labeled data significantly outperforms that trained on the whole data from scratch, highlighting the broad applicability of our method.

5. Conclusion

In this paper, we propose a novel weakly supervised 3D object detection framework named VSRD, which consists of multi-view 3D auto-labeling and subsequent training of monocular 3D object detectors using the pseudo labels generated in the auto-labeling stage. Our method demonstrates superior performance compared with the existing weakly supervised 3D object detection methods. Moreover, it exhibits remarkable scalability using partially labeled data for semi-supervised learning. Our proposed method allows leveraging abundant 2D annotations to enhance 3D object detection without explicit 3D supervision, providing a promising avenue for further advancements in the field.

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 2
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022.
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. *arXiv preprint arXiv:2304.06706*, 2023. 2
- [4] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9287–9296, 2019. 2
- [5] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: High quality object detection and instance segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 43(5):1483–1498, 2019. 7
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. 2
- [7] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [8] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5501–5510, 2022. 2
- [9] 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. 6
- [10] David Ha, Andrew Dai, and Quoc V Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016. 4
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 7
- [12] Benran Hu, Junkai Huang, Yichen Liu, Yu-Wing Tai, and Chi-Keung Tang. Nerf-rpn: A general framework for object detection in nerfs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23528–23538, 2023. 2
- [13] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [14] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 5
- [15] Yiyi Liao, Jun Xie, and Andreas Geiger. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3292–3310, 2022. 6
- [16] Zechen Liu, Zizhang Wu, and Roland Tóth. Smoke: Single-stage monocular 3d object detection via keypoint estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 996–997, 2020. 2
- [17] Qinghao Meng, Wenguan Wang, Tianfei Zhou, Jianbing Shen, Luc Van Gool, and Dengxin Dai. Weakly supervised 3d object detection from lidar point cloud. In *European Conference on computer vision*, pages 515–531. Springer, 2020. 2
- [18] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2, 3, 6
- [19] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7074–7082, 2017. 2
- [20] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 2
- [21] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022. 2
- [22] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3142–3152, 2021. 1
- [23] Liang Peng, Senbo Yan, Boxi Wu, Zheng Yang, Xiaofei He, and Deng Cai. Weakm3d: Towards weakly supervised monocular 3d object detection. *arXiv preprint arXiv:2203.08332*, 2022. 2, 7, 8
- [24] Zengyi Qin, Jinglu Wang, and Yan Lu. Weakly supervised 3d object detection from point clouds. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4144–4152, 2020. 2
- [25] Andrea Simonelli, Samuel Rota Buló, Lorenzo Porzi, Manuel López-Antequera, and Peter Kotschieder. Disentangling monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1991–1999, 2019. 1, 7
- [26] Runzhou Tao, Wencheng Han, Zhongying Qiu, Chengzhong Xu, and Jianbing Shen. Weakly supervised monocular

- 3d object detection using multi-view projection and direction consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17482–17492, 2023. 2
- [27] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 2, 4, 5
- [28] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 913–922, 2021. 1, 2
- [29] Wenhai Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, et al. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14408–14419, 2023. 7
- [30] Chenfeng Xu, Bichen Wu, Ji Hou, Sam Tsai, Ruilong Li, Jialiang Wang, Wei Zhan, Zijian He, Peter Vajda, Kurt Keutzer, et al. Nerf-det: Learning geometry-aware volumetric representation for multi-view 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23320–23330, 2023. 2, 3
- [31] Junkai Xu, Liang Peng, Haoran Cheng, Hao Li, Wei Qian, Ke Li, Wenxiao Wang, and Deng Cai. Mononerf: Nerf-like representations for monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6814–6824, 2023. 2, 3
- [32] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 2
- [33] Sergey Zakharov, Wadim Kehl, Arjun Bhargava, and Adrien Gaidon. Autolabeling 3d objects with differentiable rendering of sdf shape priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12224–12233, 2020. 2, 7, 8
- [34] Renrui Zhang, Han Qiu, Tai Wang, Ziyu Guo, Ziteng Cui, Yu Qiao, Hongsheng Li, and Peng Gao. Monodetr: Depth-guided transformer for monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9155–9166, 2023. 1, 2, 8
- [35] Yunpeng Zhang, Jiwen Lu, and Jie Zhou. Objects are different: Flexible monocular 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3289–3298, 2021. 1, 2, 8
- [36] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In *Proceedings of the AAAI conference on artificial intelligence*, pages 12993–13000, 2020. 5