

Towards Fair and Comprehensive Comparisons for Image-Based 3D Object Detection

Xinzhu Ma^{1,2} Yongtao Wang³ Yinmin Zhang^{1,2} Zhiyi Xia⁴ Yuan Meng^{4,*}
 Zhihui Wang^{3,*} Haojie Li³ Wanli Ouyang¹

¹Shanghai AI Lab ²University of Sydney ³Dalian University of Technology ⁴Tsinghua University

{maxinzhu, zhangyinmin, ouyangwanli}@pjlab.org.cn

{yongtaowang@mail., zhwang@, hjli@}dlut.edu.cn {yuanmeng, xiazy21}@mail.tsinghua.edu.cn

Abstract

In this work, we build a modular-designed codebase, formulate strong training recipes, design an error diagnosis toolbox, and discuss current methods for image-based 3D object detection. In particular, different from other highly mature tasks, e.g., 2D object detection, the community of image-based 3D object detection is still evolving, where methods often adopt different training recipes and tricks resulting in unfair evaluations and comparisons. What is worse, these tricks may overwhelm their proposed designs in performance, even leading to wrong conclusions. To address this issue, we build a module-designed codebase and formulate unified training standards for the community. Furthermore, we also design an error diagnosis toolbox to measure the detailed characterization of detection models. Using these tools, we analyze current methods in-depth under varying settings and provide discussions for some open questions, e.g., discrepancies in conclusions on KITTI-3D and nuScenes datasets, which have led to different dominant methods for these datasets. We hope that this work will facilitate future research in image-based 3D object detection. Our codes will be released at <https://github.com/OpenGVLab/3dodi>.

1. Introduction

As a new and rapidly developing research field, vision-based 3D object detection [49] shows promising potential in autonomous driving and attracts lots of attention from both academia and industry. Thanks to the unremitting efforts of numerous researchers, lots of advanced technologies, such as model designs [3, 46, 29], detection pipelines [95, 83, 27], and challenging datasets [21, 6, 75], are continuously

proposed, which significantly promotes the development of this research field.

However, although lots of breakthroughs in detection accuracy and inference speed have been achieved, there are still some critical problems to be solved, especially the standardization of evaluation protocols. Specifically, compared with the encouraging developments at the technique level, the conventional rules in model building, training recipes, and evaluation are not well-defined. As shown in Table 1, existing methods generally adopt different settings, e.g. backbones, training strategies, augmentations, etc., to build and evaluate their models, and the commonly used benchmarks (and most of the papers) only record the final accuracy. This makes the comparison of the detectors unfair and may lead to misleading conclusions.

In this paper, we aim to provide a unified platform for image-based 3D object detectors and standardize the protocols in model building and evaluation. Specifically, similar to the advanced codebases in 2D detection [9, 86], we decompose the image-based 3D detection frameworks into several separate components, e.g. backbones, necks, etc., and provide a unified implementation for current methods. Furthermore, we fully investigate existing algorithms and formulate several efficient training protocols, e.g. training schedules, data augmentation, etc. Our summarized training recipes not only provide a fair environment for evaluation but also significantly improve the performance of current methods, particularly in the KITTI-3D dataset. For example, the methods [51, 46] published two years ago trained with our recipes can achieve better (or similar) performance than the recent works [29, 55], emphasizing the importance of establishing a standard training recipe.

Furthermore, to systematically analyze the bottlenecks and issues in image-based 3D object detection, we urgently require a tool to thoroughly examine detection results. Inspired by TIDE [2], we propose an error diagnosis toolbox, named TIDE3D, to measure the detailed characterization of

Corresponding author

	backbone	# epochs	KITTI		nuScenes		
			data aug.	others	# epochs	data aug.	others
GUPNet [46]	DLA34-DLAUp	140	flip, crop	-	-	-	-
FCOS3D [81]	ResNet101-DCN-FPN	-	-	-	12	flip	TTA
PGD [80]	ResNet101-DCN-FPN	48	flip	TTA	24	flip	TTA
BEVDet [27]	ResNet50-FPN-LSS	-	-	-	24	flip, BEV aug.	TTA

Table 1: **Overview of example methods on KITTI-3D and nuScenes benchmarks.** The existing methods generally adopt different settings, *e.g.* backbones, epochs, augmentations, *etc.*, in model training. Besides, KITTI-3D and nuScenes apply different evaluation metrics and are dominated by different detection pipelines. TTA denotes the test-time augmentation.

detection algorithms. Specifically, we decouple the detection errors into seven types, and then independently quantify the impact of each error type by calculating the overall performance improvement of the model after fixing the specified errors. In this way, we can analyze a specific aspect of given algorithms while isolating other factors. In addition to characterizing the detailed features of the models, the proposed TIDE3D also has other useful applications. For example, the effectiveness and action mechanism of a specific design/module can be explored by analyzing the change of error distribution of the baseline model with or without the target design/module.

Besides, we also find the most concerned two datasets, *i.e.* KITTI-3D [21] and nuScenes [6], are dominated by different detection pipelines, and even gradually become separate research communities. For example, none of the popular Bird’s Eye View (BEV) detection methods provide the KITTI-3D results, and the top-performing methods in the KITTI-3D leaderboard also hard to achieve good results in nuScenes. We apply the cross-metric evaluation (*i.e.* applying the nuScenes-style metrics on the KITTI-3D dataset and vice versa) on the representative models and report the adopted metric is the main factor causing this phenomenon. We also provide TIDE3D analyses for this issue.

To summarize, the contributions of this work are as follows: First, we build a modular-designed codebase for the community of image-based 3D object detection, which can serve as a foundation for future research and algorithm implementation. Second, we investigate the training settings and formulate standard training recipes for this task. Third, we provide an error diagnosis toolbox that can quantitatively analyze the detection models at a fine-grained level. Last, we discuss some open problems in this field, which may provide insights for future research. We hope our codebase, training recipes, error diagnosis toolbox, and discussions will promote better and more standardized research practices within the image-based 3D object detection community.

2. Related Work

Image-based 3D detection. Image-based 3D detection is a rapidly developing research direction, and there are lots of

detectors are proposed. Here we review these works based on the taxonomy proposed in [49]: methods based on result-lifting, methods based on feature-lifting, and methods based on data-lifting. Refer to [49] for the detailed and comprehensive literature review.

Result-lifting. The methods in this branch first estimate the 2D projections and other items (*e.g.* depth, orientation, *etc.*) of the 3D bounding boxes and then lift the results from the 2D image plane to the 3D world space. In particular, pioneering works [11, 10, 3, 95, 81, 32, 84] introduce the popular 2D detection paradigms, *e.g.* faster R-CNN [64] and FCOS [77], into this research field, and lots of following works improve these baseline models in several aspects, including backbone designs [3, 29], loss functions [70, 51, 22, 52], depth estimation [46, 93, 71, 80, 35, 37, 56], geometric constraints [60, 41, 39, 8, 44, 28, 62, 31, 61], feature embedding [42, 94], key-point constrains [7, 33, 45], depth augmented learning [87, 18, 15, 23, 14, 78], temporal sequences [4, 79], semi-supervised learning [54, 48, 40, 26], NMS [68, 30], data integration [5], *etc.*

Feature-lifting. OFT [65] and DSGN [13] are the pioneering works in this research line, which lift the 2D features into 3D features (generally represented by BEV map) and then directly estimates the 3D bounding boxes with the resulting features. CaDDN [63] uses a better transformation strategy [57] and achieves promising results. BEVDet [27] introduces this paradigm into the multi-camera setting, and this BEV pipeline gradually dominates the nuScenes [6] benchmark with effects of lots of follow-up works, such as [38, 36, 34].

Data-lifting. The data-lifting-based methods first estimate dense depth maps for the input images and then lift the pixels into 3D points using the estimated depth and camera parameters. After that, they generally leverage the LiDAR-based 3D detectors to predict the results from these ‘pseudo-LiDAR’ signals [83, 50]. The followers in this group are involved in the following aspects: improving the quality of pseudo-LiDAR [89, 20], focusing on foreground objects [82, 74, 88], end-to-ending training [59], feature representation [47, 53], geometric constraints [85] or confidence refinement [69].

Codebases. At present, there are two public codebases that

support vision-based 3D object detection: MMDetection3D [16] and OpenPCDet [76]. However, these two codebases are initially designed for LiDAR-based 3D object detection, and the supported vision-based methods are limited. In particular, MMDetection3D supports four vision-based detectors [66, 44, 81, 80], and OpenPCDet only support CaDDN [63][†] (also note some methods are developed based on these two codebases although they are not included in the official ones). In contrast, our codebase is designed for vision-based methods, excluding redundant codes and dependencies for LiDAR-based methods.

Error diagnosis for object detection. To obtain the detailed characterization under the overall evaluation metric and show the strengths and weaknesses of the given models, some works [25, 1, 2] exist to analyze the errors in 2D detectors. In particular, [25] divides the errors of false positives into several categories and selects the top N most confident detections for errors analysis. [1] further adds statistics for false negatives. More specifically, it progressively replaces the predicted items with corresponding ground-truth values and uses the Δ AP metric to quantify the importance of each item. [2] points out that this iterative approach incorrectly amplifies the impact of later error types, thus proposing a strategy where each error type is fixed independently. These works contribute to model analysis in the field of 2D detection, and to our best knowledge, there is no previous attempt to provide such a generic toolbox for 3D object detection. In this work, we provide a costumed error diagnosis toolbox to fill this gap.

3. Codebase

The objective of this work is to build a fair platform and provide experimental analysis for image-based 3D detection algorithms. For this purpose, we first build the general codebase which can support this work and facilitate future research. Specifically, we decompose the image-based 3D detection frameworks into different components, *e.g.* backbones, necks, heads, *etc.*, and we provide common choices for each component, *e.g.* ResNet [24], DLA [90], or DCN [96] for the backbones. In addition to the CNN modules, we also provide the implementation of other parts, such as augmentations, post-processing, *etc.* Currently, the proposed codebase supports more than ten methods such as [51, 46, 91, 43, 44, 95, 81, 48, 54, 27, 80, 15] and two commonly used datasets [21, 6]. Our codebase will be publicly available and continuously maintained.

Comparison with official implementation. Here we provide a comparison of our implementation and the official implementation for some representative works. In particular, here we take MonoDLE and GUPNet as examples for KITTI-3D dataset, and FCOS3D as example for nuScenes

	KITTI			nuScenes	
	Easy	Mod.	Hard	mAP	NDS
MonoDLE [51]	17.89	13.87	12.03	-	-
MonoDLE - ours	18.37	13.79	12.24	-	-
GUPNet [46]	21.48	15.22	12.79	-	-
GUPNet - ours	21.66	15.45	12.51	-	-
FCOS3D [81]	-	-	-	30.6	38.1
FCOS3D - ours	-	-	-	30.4	38.5

Table 2: **Comparison of the implementations** on KITTI and nuScenes *validation* sets. We report the average performance of the last epoch for both the official codes and our implementations over five runs.

dataset. Based on the results shown in Table 2, we can find our implementation gets consistent or better results than the official numbers. More details and the results for other models can be found in our public codebase.

4. Diagnosis Toolkit

We propose a general toolkit, named TIDE3D, to diagnose the cause of errors for 3D object detection models. We first briefly review the computing process of mean Average Precision (mAP) [19] and then introduce how to divide the detection errors into fine-grained types. Finally, we present how to weigh these errors and the implementation details. We also recommend readers refer to TIDE [2] for preliminary knowledge.

Review of mAP. As the most commonly used metric for evaluating object detection methods, mAP provides a comprehensive overview of a detector’s performance. Given the predictions and corresponding ground truths, each ground truth is matched to at most one prediction according to a specified metric, *e.g.* KITTI-3D uses 3D Intersection over Union (IoU) and nuScenes adopts center distance as the metric. If multiple predictions meet the constraint, the ground truth only matches the prediction which has the maximum confidence score. The matched predictions are true positives (TP), and the remaining ones are false positives (FP). After sorting the predictions by descending confidence, the number of true positives and false positives in the subset of predictions with confidence scores greater than c is counted as N_{TP} and N_{FP} . Then, the precision and recall of the predictions subset are obtained by:

$$P_c = \frac{N_{TP}}{N_{TP} + N_{FP}} \quad R_c = \frac{N_{TP}}{N_{GT}} \quad (1)$$

where N_{GT} is the number of ground-truth objects. As the confidence threshold c changes, the precision-recall curve is plotted and the Average Precision (AP) is obtained by calculating the area under the curve. Finally, mAP is obtained by averaging the AP values of all categories.

Error definitions. Due to the complexity of calculating

[†]data collected at 24-02-2023

mAP, it prevents researchers from further analyzing the performance, and two detectors with the same mAP value may have different strengths and weaknesses. To address this issue, we segment the detection errors into several separate groups and measure the contribution of each error type. Specifically, based on the taxonomy of 2D detection error proposed in TIDE [2] (including classification error, localization error, both classification and localization error, duplication error, background error, and missing error. See Table 3 for details), we further propose two modifications for 3D object detection, including the sub-error types of localization and the ranking error.

Sub-error types of localization. Since the localization of the 3D bounding box is determined by many factors, we further divide the localization error into three sub-errors to explore the effect of predicted location, dimension, and orientation on inaccurate localization.

Ranking error. To describe each object in the scene, the detection model needs to output the 3D bounding box and its confidence. However, some detections may have higher confidence but lower accuracy. In the process of calculating AP, the detected boxes are sorted by *descending confidence* and the precision of the high-confidence detections is calculated first. Therefore, the misalignment between confidence and the quality of the box naturally leads to the decline of AP and brings about possible room for improvement.

Weighting the errors. We use ΔAP to quantify the impact of each error type. In particular, similar to TIDE [2], we independently fix the errors (called “oracle”, see Table 3 for the details), and compute the change of AP to measure the impact of this error type by:

$$\Delta AP_o = AP_o - AP, \quad (2)$$

where AP_o is the AP after applying the oracle o . In this way, we can capture the effect of each error type on the final metric, and a lower ΔAP of the given error type indicates the model performs better at this part.

Due to current datasets generally adopting different evaluation metrics, the detailed implementations of TIDE3D are also slightly different for these datasets [21, 6, 17]. More details of the specific implementations are provided in Appendix A.3.

5. Approach

We first discuss the effects of common training settings on the final accuracy of current methods and then provide efficient and fair training recipes for further work. Furthermore, we also show the state-of-the-art (SOTA) methods can be further promoted in performance with the summarized training recipes and some existing techniques in the KITTI-3D benchmark. Besides, the detailed features of the detection algorithms are evaluated with the proposed

TIDE3D, and we also show our diagnosis tool can be used to explore some open issues. Finally, we discuss the conflicts between KITTI-3D and nuScenes datasets.

Training recipes. We observe that current methods generally adopt different training recipes, *e.g.* the training schedule in KITTI-3D ranges from 60 epochs [44] to 200 epochs [55]. We show that training differences like this greatly affect the performance of the algorithms and may mislead the latecomers to some extent, and it is necessary to formulate standard training recipes. Specifically, in KITTI-3D, we take the 70-epoch schedule as baseline ($1\times$) and also evaluate the models under $2\times$ and $3\times$ schedules. Similarly, the corresponding settings in nuScenes are 12, 24, and 36 epochs. We recommend the $2\times$ schedule because this is the closest one to the common choice of other algorithms, and use other schedulers for further evaluation. Furthermore, we also ablate other choices in model training and provide a general training recipe for image-based 3D detectors. We report, in addition to some basic choices such as random crop or horizontal flip, the photometric distortion and one-cycle learning rate [73] are also highly effective for our task and should be included in the conventional setting.

Promoted baselines in KITTI-3D. We also find, although KITTI-3D is released for a longer time and more methods are proposed on KITTI-3D benchmark, the baseline models in KITTI-3D are still *under-tuned* compared to these on nuScenes. Here we show current baseline models can be greatly promoted. We report that a two-year-ago model [46] (re-trained with pre-trained backbone, pseudo-labels, and the summarized training recipe) can achieve 17.65 AP_{40} on the KITTI-3D validation set, surpassing the SOTA by 1.11 points. Even so, we believe this number can be improved with further tuning or other tricks (*e.g.* test-time augmentation and model-ensemble [81]).

Dataset and metrics. KITTI-3D [21] and nuScenes [6] are the most concerned datasets in the image-based 3D detection field, however, they are dominated by different detection pipelines. We find that they are different in many aspects, and nuScene is generally considered to be more friendly to image-based algorithms. We adopt nuScenes-style evaluation on the KITTI-3D (and vice versa) and provide our discussion based on the cross-metric evaluation. More details are given in Section 6.4.

Detailed analysis with TIDE3D. With the TIDE3D introduced in Section 4, we can: (i) Analyze the detailed characteristics of 3D object detection models, whether it is an image-based model or LiDAR-based model, evaluated with KITTI-3D-style or nuScene-style metrics, and designed for outdoor or indoor scenes (see Appendix A.5 for the analyses for indoor dataset [17] and LiDAR-based models). (ii) Conduct solid ablation studies and explore the action mechanisms for specific designs/modules.

Error Type	Definition	Oracle
Classification	Incorrect classification and correct localization.	Correct the category classification / location of detection, and delete this detection if it becomes duplicated.
Localization	Correct classification and incorrect localization.	
Both Cls. and Loc.	Incorrect classification and incorrect localization.	Delete the inaccurate / duplicate detection.
Duplication	Correct classification and correct localization, but the corresponding GT has been matched by another higher-scoring detection.	
Background	Detection from the background area.	
Missing	Undetected GT not covered by classification and localization error.	Delete the missed ground truth.
Ranking	Inconsistency between confidence and localization quality of detections.	Sort the detections by descending IoU score and calculate the precision of accurate detections first.

Table 3: Definition and oracle of detection errors.

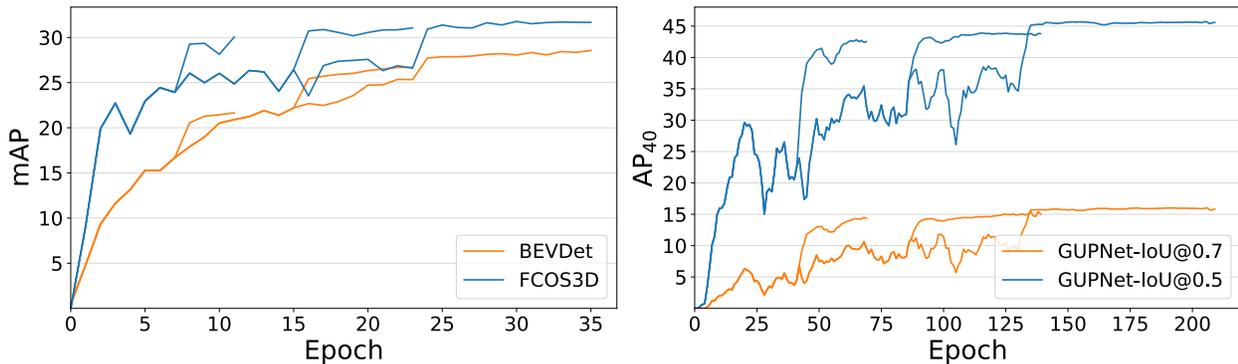


Figure 1: Learning curves on nuScenes (**Left**) and KITTI-3D (**Right**) with different detection models. With the $3\times$ scheduler, BEVDet with R50-FPN optimized by AdamW and FCOS3D with R101-DCN-FPN optimized by SGD achieve 28.56 and 31.67 mAP on nuScenes, respectively, while GUPNet with DLA-34 optimized by Adam achieves 15.64 $AP_{40}@IoU=0.7$ on KITTI-3D (*validation* set). For better presentation, we present the curves by moving the average with windows size = 5. The original curves are shown in Appendix A.4.

6. Experiments and Analysis

Dataset and metrics. We conduct experiments on KITTI-3D and nuScenes datasets, and the detailed introduction for these two datasets is provided in Appendix A.1. Due to access restrictions of testing servers, the evaluation is conducted on *validation* set, and we further apply cross-dataset metrics for better evaluation, *e.g.* applying nuScenes-style metrics on the KITTI-3D dataset.

Baseline models. In this work, we mainly focus on the baselines based on the popular detection pipeline, including the CenterNet pipeline [46, 29], the BEV pipeline [27], and the FCOS3D pipeline [81, 80].

Other settings. To replicate the reported performance and minimize the impact of complex architectures and additional hyper-parameters, we adopt the simplest default configurations provided in the official code for most methods. Specifically, we use basic training schedules, simple architectures, and standard data augmentation techniques as the default settings. Further details beyond to these settings are provided in according sections.

6.1. Training Recipes

Training schedules. We first present the effects of training schedules. Specifically, we experiment with three different schedules, ranging from $1\times$ to $3\times$, for both KITTI-3D and nuScenes. Notably, we plot the training curves of all three schedules in the same figure to facilitate comparison. The AP curves shown in Figure 1 indicate that schedules have a significant impact on the final performance, highlighting the importance of a standard scheduler for a fair comparison. Furthermore, the following observations can be summarized from these curves:

- i. Different detection models require varying numbers of epochs as the optimal settings. For example, BEVDet exhibits significant and smooth performance improvements as the number of training epochs increases, while FCOS3D achieves near-optimal performance faster with perturbation after learning rate decay. This discrepancy may stem from the difference in optimizers, iterations, and detection pipelines. Intuitively, required training time can substantially vary across different settings, which may reduce the

	schedule	pre-training	AP ₄₀ @IoU=0.7			AP ₄₀ @IoU=0.5		
			easy	mod.	hard	easy	mod.	hard
GUPNet [46]	2×	ImageNet	21.48	15.22	12.79	59.30	43.83	38.40
DID-M3D [55]	≈ 3×	ImageNet	22.98	16.12	14.03	-	-	-
DEVIANT [29]	2×	ImageNet	24.63	16.54	14.52	61.00	46.00	40.18
GUPNet	2×	None	18.74	13.40	10.62	53.80	40.69	35.73
GUPNet	2×	ImageNet	21.51	15.82	13.30	60.18	44.35	39.27
GUPNet	2×	DD3D	24.25	15.82	13.26	63.16	44.60	40.08
GUPNet + PL	2×	DD3D	24.71	17.25	15.21	63.73	46.56	42.45

Table 4: **Promoted baselines on KITTI-3D validation set.** We report the 3D object detection performance of GUPNet with various pretrained backbones, evaluated by AP₄₀@IoU=0.7/0.5. We also present SOTA methods for reference. Note that both DID-M3D and DEVIANT are built based on the GUPNet baseline, and our modifications are only involved in the training phase, which shows the importance of building fair training recipes again.

fairness of comparisons between them.

ii. Increasing the training procedure enhances the stability and precision of the fine-tuned model. Even after completing the 3× training, the trend of mAP/AP₄₀ improvement continues to exist. The commonly accepted training schedule (mostly 1× or 2× on nuScenes) may be insufficient to unlock the full potential of a model. Therefore, the standard training procedure remains further discussed.

iii. We observe a significant drop and a recovery at the beginning of the finetuning stage of 1× training in the experiments of BEVDet and GUPNet. With training time increasing from 2× to 3×, this phenomenon alleviates and even disappears. We suppose that longer training time helps the network learn more robust features, which prevents the overfitting brought by the sudden decay of the learning rate. More training steps at a small learning rate also help counter this issue.

iv. Even when the training is about to be completed, we still observe performances are unstable (see Appendix A.4 for the un-smoothed curves). This phenomenon is particularly evident for the KITTI-3D-style metrics. Based on this, the comparison based on KITTI-3D metrics should be carefully conducted.

Other training recipes. Due to strict geometric relations between the 2D image plane and the 3D world space, some complicated data augmentations are hard to be applied, and the horizontal flip is the common choice for existing methods. Note some simple geometric augmentations, like center crop and resize, are also commonly used in the KITTI-3D dataset, and see Appendix A.6 for more discussions on such data augmentations. Here we show the photometric distortion and the one-cycle learning rate schedule are two cost-free training tricks to improve the detection accuracy (especially in KITTI-3D) and show their effectiveness with two representative baselines in Table 5. We can find that both of them improve the GUPNet by a significant margin and can work together. Although nuScenes is a large-scale

dataset, applying photometric distortion still brings a modest mAP improvement. Besides, we find the effects of one-cycle learning rate on nuScenes baselines are unstable, so we omit the corresponding cells to avoid misleading.

	baseline	w/ distortion	w/ one-cycle lr	full
GUPNet [AP ₄₀]	15.26	16.54	15.87	16.83
BEVDet [mAP]	25.12	25.20	-	-

Table 5: The photometric distortion and one-cycle learning rate are also effective training choices.

6.2. Promoted KITTI-3D Baselines.

Leveraging pre-training weights. We show the impact of weight initialization on detection models in Table 4. In particular, we compare three initial weights, including random initialization, ImageNet pretraining, and DD3D pertaining [53]. We can find that: **i.** Similar to other computer vision tasks, model pre-training provides useful priors to the models and improve the performance significantly. **ii.** As the data domains and proxy tasks become more similar (see Table 6), the pre-training weights also become more effective. Based on the above observations, designing custom pre-training algorithms, especially in an unsupervised manner, should be a promising research direction for future work.

	domain	proxy task	# img. × # ep.
ImageNet	daily scenes	classification	1.3M × 100
DD3D	driving scenes	depth estimation	15M × 12.8

Table 6: Information of different pre-training models. Ep. denotes the pretraining epochs.

Pseudo-labeling. According to [54, 48], we generate the pseudo-labels (only for the 3,712 training images) from a LiDAR-based model [67]. Then we use them to train GUPNet and report the performance in Table 4. According to the results, the final result of our GUPNet can achieve 17.65

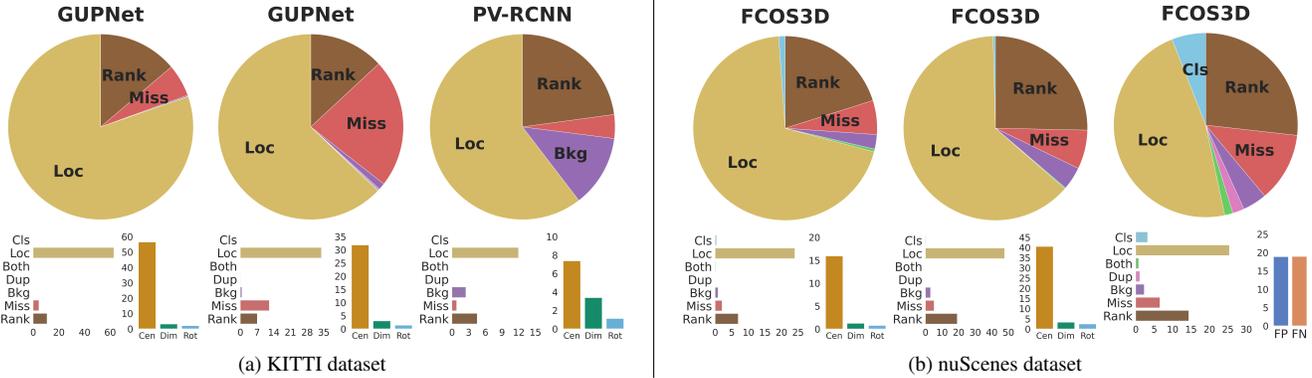


Figure 2: **Example error diagnosis results.** (a) Error analyses on KITTI-3D *validation* set under the moderate setting. From left to right, we show the results of GUPNet [46] for the Car category with 0.7 IoU threshold, 0.5 IoU threshold, and the results of a LiDAR-based method [67] for the Car category with 0.7 IoU threshold. (b) Error diagnosis results of FCOS3D [81] on nuScenes *validation* set. From left to right, we show error diagnosis results based on KITTI-3D-style metrics with IoU@0.5 for all categories, the Car category, and the results based on nuScenes-style metrics for all categories.

AP_{40} under the moderate setting, surpassing the previous SOTA (DEVIANT) by 0.71 points. We emphasize that we do not make any changes to the detection model, and the modifications are only involved in the training phase. This result shows the urgency and importance of establishing a fair and efficient training recipe.

6.3. TIDE3D Analysis

Error diagnosis. Here we present the error diagnosis results of some representative methods, settings, and evaluation metrics in Figure 2. According to these error distributions, we can get the following observations:

- i. Similar to [51, 29, 53, 80], we find localization error is the bottleneck of current methods and it is mainly caused by the inaccurate location (center). Even for the LiDAR-based method [67] which estimates the results from the data with accurate spatial information, localization error still ranks first among all error types. Differently, for [67], the dimension and the orientation also have a significant impact on the quality of localization.
- ii. We find the misalignment between the confidence and the bounding box’s quality is a serious problem in this field, suggesting that giving more accurate confidence is an effective way to improve detection accuracy. An empirical study of this problem is given in the following part.
- iii. Under a lower IoU threshold, the KITTI-3D-style metric is still dominated by the localization error. Meanwhile, this issue has been alleviated on nuScenes-style metric. This indicates 3D IoU-based AP is a very sensitive metric, and center distance-based AP is more friendly to image-based methods. More TIDE3D results and analyses, *e.g.* indoor scenes, are provided in Appendix A.5.

Validating design choices. We show the proposed TIDE3D

	$AP_{40} \uparrow$	$E_{loc} \downarrow$	$E_{bkg} \downarrow$	$E_{miss} \downarrow$	$E_{rank} \downarrow$
GUPNet w/ 2D conf.	14.81	69.87	0.46	1.73	15.10
GUPNet w/ 3D conf.	17.02	62.84	0.15	4.41	10.81
Improvement	+2.21	-7.03	-0.31	+2.68	-4.29

Table 7: **Effects of 3D confidence.** Experiments are conducted on KITTI-3D *validation* set. ΔAP_{40} is denoted as E for brevity.

can validate whether a given design choice supports its motivation/claim. For instance, some works, such as [46, 69, 5] proposed their designs to get better confidence scores of detection results. Here we use TIDE3D to validate the design of GUPNet. Particularly, GUPNet models the uncertainty of estimated depth (σ) and further gets the confidence of depth with this uncertainty by: $p_{depth} = exp(-\sigma)$. Finally, the 3D confidence (p_{3D}) can be obtained by correcting 2D confidence (p_{2D}) with depth confidence: $p_{3D} = p_{depth} \cdot p_{2D}$. Table 7 gives the results of GUPNet equipped with 2D confidence and 3D confidence respectively. We observe that the AP_{40} improved by 2.21 and the ranking error reduces to 10.81 from 15.10, which confirms the effectiveness of this design. Meanwhile, we find the localization error also reduced, which is caused by the high-quality results being improved and corresponding results having a higher priority in AP computing. See Appendix A.7 for more analysis about the ranking error and localization error in TIDE3D.

6.4. Metric Analysis for KITTI-3D and nuScenes

Additionally, we have observed interesting discrepancies in evaluation metrics between KITTI-3D and nuScenes. Specifically, as shown in Table 8, GUPNet and FCOS3D get similar AP of Car category under the nuScenes-style met-

	Dataset	Class	KITTI Metric			nuScenes Metric	
			@0.7	@0.5	@0.25	mAP	NDS
GUPNet	KITTI-3D	Car	14.85	40.85	62.87	47.57	61.85
PV-RCNN	KITTI-3D	Car	82.41	94.31	94.47	92.74	92.61
FCOS3D	nuScenes	Car	1.83	20.87	50.74	47.96	54.32
FCOS3D	nuScenes	All	0.26	6.06	17.57	31.05	40.31
BEVDet	nuScenes	All	0.10	6.55	19.25	30.75	38.22

Table 8: **Evaluation Metric.** We exchange the metrics between KITTI-3D and nuScenes for the evaluation of 3D object detectors. This provides a standardized way to compare the performance of 3D object detectors across the two datasets. Dataset represents the source of training data for the detector. For detectors trained on KITTI-3D, we only compute translation, scale, and orientation for NDS under nuScenes metric. For KITTI metric, we evaluate the results of cars under the hard strict with varying IoU thresholds.

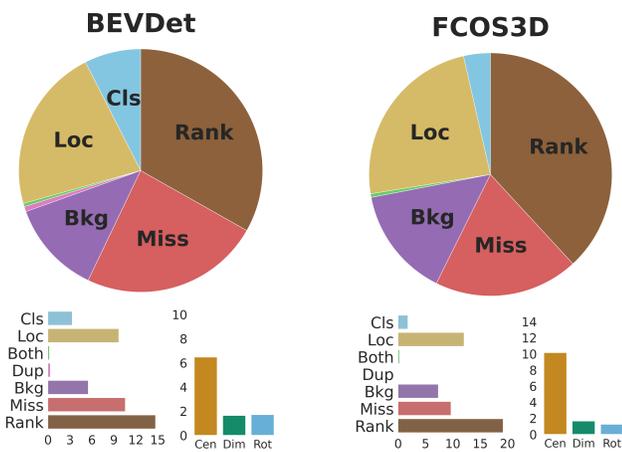


Figure 3: **TIDE3D analyses** of BEVDet [27] and FCOS3D [81] for all categories on nuScenes *validation* set based on the KITTI-style metric with 0.25 IoU threshold. Our results indicate that FCOS3D is impacted by various errors, while BEVDet is mainly dominated by localization errors.

ric, but they show a significant performance gap under the KITTI-style metric. As we know, KITTI-3D primarily uses the strict 3D IoU metric with a certain threshold as the metric, while nuScenes measures the center distance and a series of decoupled error attributes such as translation, scale, orientation, *etc.*, separately (Appendix A.1 provides the details of metrics). We find that translation, scale, and orientation jointly affect the result of IoU, which means that methods achieving competitive performance on the KITTI-3D metric require comprehensive development, while methods with special strengths in nuScenes can also get a good performance. According to our experience, a model with good performance under the KITTI-style metric generally performs well under the nuScenes-style metric, but the opposite is not always true.

Furthermore, we present TIDE3D analyses of BEVDet [27] and FCOS3D [81] for all categories on nuScenes *val-*

idation set based on the KITTI-style metric with 0.25 IoU threshold in Figure 3. Based on the distribution of localization sub-errors (6.5-1.6-1.7 *v.s.* 10.1-1.7-1.2), we can find that FCOS3D makes more accurate predictions for orientation and dimension than BEVDet. One of the reasons for this is the loss of height information of BEVDet in the LSS (lift, splat, shoot) process, which serves to transform the image feature into BEV space. Consequently, the model may struggle to estimate the height of objects accurately and the center height of objects. See Appendix A.2 for more results and analyses.

The above two reasons explain why the mainstream methods for KITTI and nuScenes are distinct, and it is reasonable and necessary to separately design models to match the preferences of these two kinds of metrics.

7. Conclusion

In this paper, we aim to standardize the evaluation protocols for image-based 3D object detection, an ever-evolving research field. We build a modular-designed codebase and provide efficient training recipes for this task, leading to significant improvements in the performance of current methods. Besides, we offer an error diagnosis toolbox to measure the detailed characterization of detection algorithms and highlight some open problems in the field. We hope our codebase, training recipes, error diagnosis toolbox, and discussions will foster better and more standardized research practices within the image-based 3D object detection community.

Acknowledgments

This work was supported by the Australian Medical Research Future Fund MRFAI000085, CRC-P Smart Material Recovery Facility (SMRF) – Curby Soft Plastics, CRC-P ARIA - Bionic Visual-Spatial Prosthesis for the Blind, and National Natural Science Foundation of China (NSFC Grants NO. 61976038).

References

- [1] Coco analysis toolkit. <http://cocodataset.org/#detection-eval>. Accessed: 2020-03-01. **3**
- [2] Daniel Bolya, Sean Foley, James Hays, and Judy Hoffman. Tide: A general toolbox for identifying object detection errors. In *ECCV*, 2020. **1, 3, 4, 15**
- [3] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In *ICCV*, 2019. **1, 2**
- [4] Garrick Brazil, Gerard Pons-Moll, Xiaoming Liu, and Bernt Schiele. Kinematic 3d object detection in monocular video. In *ECCV*, 2020. **2**
- [5] Garrick Brazil, Julian Straub, Nikhila Ravi, Justin Johnson, and Georgia Gkioxari. Omni3D: A large benchmark and model for 3D object detection in the wild. *arXiv:2207.10660*, 2022. **2, 7**
- [6] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020. **1, 2, 3, 4, 12, 14**
- [7] Florian Chabot, Mohamed Chaouch, Jaonary Rabarisoa, Céline Teuliere, and Thierry Chateau. Deep manta: A coarse-to-fine many-task network for joint 2d and 3d vehicle analysis from monocular image. In *CVPR*, 2017. **2**
- [8] Hansheng Chen, Yuyao Huang, Wei Tian, Zhong Gao, and Lu Xiong. Monorun: Monocular 3d object detection by reconstruction and uncertainty propagation. In *CVPR*, 2021. **2**
- [9] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. **1**
- [10] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *CVPR*, 2016. **2**
- [11] Xiaozhi Chen, Kaustav Kundu, Yukun Zhu, Andrew G Berneshawi, Huimin Ma, Sanja Fidler, and Raquel Urtasun. 3d object proposals for accurate object class detection. In *NeurIPS*, 2015. **2, 12**
- [12] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *CVPR*, 2017. **12**
- [13] Yilun Chen, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Dsgn: Deep stereo geometry network for 3d object detection. In *CVPR*, 2020. **2**
- [14] Yi-Nan Chen, Hang Dai, and Yong Ding. Pseudo-stereo for monocular 3d object detection in autonomous driving. In *CVPR*, 2022. **2**
- [15] Zhiyu Chong, Xinzhu Ma, Hong Zhang, Yuxin Yue, Haojie Li, Zhihui Wang, and Wanli Ouyang. Monodistill: Learning spatial features for monocular 3d object detection. In *ICLR*, 2022. **2, 3**
- [16] MMDetection3D Contributors. MMDetection3D: Open-MMLab next-generation platform for general 3D object detection. <https://github.com/open-mmlab/mmdetection3d>, 2020. **3**
- [17] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. **4, 12**
- [18] Mingyu Ding, Yuqi Huo, Hongwei Yi, Zhe Wang, Jianping Shi, Zhiwu Lu, and Ping Luo. Learning depth-guided convolutions for monocular 3d object detection. In *CVPR*, 2020. **2**
- [19] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2009. **3**
- [20] Divyansh Garg, Yan Wang, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger, and Wei-Lun Chao. Wasserstein distances for stereo disparity estimation. 2020. **2**
- [21] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012. **1, 2, 3, 4, 12**
- [22] Jiaqi Gu, Bojian Wu, Lubin Fan, Jianqiang Huang, Shen Cao, Zhiyu Xiang, and Xian-Sheng Hua. Homography loss for monocular 3d object detection. In *CVPR*, 2022. **2**
- [23] Xiaoyang Guo, Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Liga-stereo: Learning lidar geometry aware representations for stereo-based 3d detector. In *ICCV*, 2021. **2**
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. **3**
- [25] Derek Hoiem, Yodsawalai Chodpathumwan, and Qieyun Dai. Diagnosing error in object detectors. In *ECCV*, 2012. **3**
- [26] Yu Hong, Hang Dai, and Yong Ding. Cross-modality knowledge distillation network for monocular 3d object detection. In *ECCV*, 2022. **2**
- [27] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021. **1, 2, 3, 5, 8**
- [28] Jason Ku, Alex D Pon, and Steven L Waslander. Monocular 3d object detection leveraging accurate proposals and shape reconstruction. In *CVPR*, 2019. **2**
- [29] Abhinav Kumar, Garrick Brazil, Enrique Corona, Armin Parchami, and Xiaoming Liu. Deviant: Depth equivariant network for monocular 3d object detection. In *ECCV*, 2022. **1, 2, 5, 6, 7, 13**
- [30] Abhinav Kumar, Garrick Brazil, and Xiaoming Liu. Groomed-nms: Grouped mathematically differentiable nms for monocular 3d object detection. In *CVPR*, 2021. **2**
- [31] Buyu Li, Wanli Ouyang, Lu Sheng, Xingyu Zeng, and Xiaogang Wang. Gs3d: An efficient 3d object detection framework for autonomous driving. In *CVPR*, 2019. **2**
- [32] Peiliang Li, Xiaozhi Chen, and Shaojie Shen. Stereo r-cnn based 3d object detection for autonomous driving. In *CVPR*, 2019. **2**
- [33] Peixuan Li, Huaici Zhao, Pengfei Liu, and Feidao Cao. Rtm3d: Real-time monocular 3d detection from object keypoints for autonomous driving. In *ECCV*, 2020. **2**
- [34] Yinhao Li, Han Bao, Zheng Ge, Jinrong Yang, Jianjian Sun, and Zeming Li. Bevstereo: Enhancing depth estimation in multi-view 3d object detection with dynamic temporal stereo. In *AAAI*, 2023. **2**

- [35] Yingyan Li, Yuntao Chen, Jiawei He, and Zhaoxiang Zhang. Densely constrained depth estimator for monocular 3d object detection. In *ECCV*, 2022. 2
- [36] Yin hao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *AAAI*, 2023. 2
- [37] Zhuoling Li, Zhan Qu, Yang Zhou, Jianzhuang Liu, Haoqian Wang, and Lihui Jiang. Diversity matters: Fully exploiting depth clues for reliable monocular 3d object detection. In *CVPR*, 2022. 2
- [38] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers. In *ECCV*, 2022. 2
- [39] Qing Lian, Peiliang Li, and Xiaozhi Chen. Monojsq: Joint semantic and geometric cost volume for monocular 3d object detection. In *CVPR*, 2022. 2
- [40] Qing Lian, Yanbo Xu, Weilong Yao, Yingcong Chen, and Tong Zhang. Semi-supervised monocular 3d object detection by multi-view consistency. In *ECCV*, 2022. 2
- [41] Qing Lian, Botao Ye, Ruijia Xu, Weilong Yao, and Tong Zhang. Exploring geometric consistency for monocular 3d object detection. In *CVPR*, 2022. 2
- [42] Ce Liu, Shuhang Gu, Luc Van Gool, and Radu Timofte. Deep line encoding for monocular 3d object detection and depth prediction. In *BMVC*, 2021. 2
- [43] Xianpeng Liu, Nan Xue, and Tianfu Wu. Learning auxiliary monocular contexts helps monocular 3d object detection. In *AAAI*, 2022. 3
- [44] Zechen Liu, Zizhang Wu, and Roland Tóth. Smoke: Single-stage monocular 3d object detection via keypoint estimation. In *CVPR Workshops*, 2020. 2, 3, 4, 14
- [45] Zongdai Liu, Dingfu Zhou, Feixiang Lu, Jin Fang, and Liangjun Zhang. Autosshape: Real-time shape-aware monocular 3d object detection. In *ICCV*, 2021. 2
- [46] Yan Lu, Xinzhu Ma, Lei Yang, Tianzhu Zhang, Yating Liu, Qi Chu, Junjie Yan, and Wanli Ouyang. Geometry uncertainty projection network for monocular 3d object detection. In *ICCV*, 2019. 1, 2, 3, 4, 5, 6, 7, 14
- [47] Xinzhu Ma, Shinan Liu, Zhiyi Xia, Hongwen Zhang, Xingyu Zeng, and Wanli Ouyang. Rethinking pseudo-lidar representation. In *ECCV*, 2020. 2
- [48] Xinzhu Ma, Yuan Meng, Yinmin Zhang, Lei Bai, Jun Hou, Shuai Yi, and Wanli Ouyang. An empirical study of pseudo-labeling for image-based 3d object detection. *arXiv preprint arXiv:2208.07137*, 2022. 2, 3, 6, 13
- [49] Xinzhu Ma, Wanli Ouyang, Andrea Simonelli, and Elisa Ricci. 3d object detection from images for autonomous driving: a survey. *arXiv preprint arXiv:2202.02980*, 2022. 1, 2
- [50] Xinzhu Ma, Zhihui Wang, Haojie Li, Pengbo Zhang, Wanli Ouyang, and Xin Fan. Accurate monocular 3d object detection via color-embedded 3d reconstruction for autonomous driving. In *ICCV*, 2019. 2
- [51] Xinzhu Ma, Yinmin Zhang, Dan Xu, Dongzhan Zhou, Shuai Yi, Haojie Li, and Wanli Ouyang. Delving into localization errors for monocular 3d object detection. In *CVPR*, 2021. 1, 2, 3, 7, 14
- [52] Fabian Manhardt, Wadim Kehl, and Adrien Gaidon. Roi-10d: Monocular lifting of 2d detection to 6d pose and metric shape. In *CVPR*, 2019. 2
- [53] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *ICCV*, 2021. 2, 6, 7
- [54] Liang Peng, Fei Liu, Zhengxu Yu, Senbo Yan, Dan Deng, Zheng Yang, Haifeng Liu, and Deng Cai. Lidar point cloud guided monocular 3d object detection. In *ECCV*, 2022. 2, 3, 6
- [55] Liang Peng, Xiaopei Wu, Zheng Yang, Haifeng Liu, and Deng Cai. Did-m3d: Decoupling instance depth for monocular 3d object detection. In *ECCV*, 2022. 1, 4, 6
- [56] Wanli Peng, Hao Pan, He Liu, and Yi Sun. Ida-3d: Instance-depth-aware 3d object detection from stereo vision for autonomous driving. In *CVPR*, 2020. 2
- [57] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, 2020. 2
- [58] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *ICCV*, 2019. 12, 13
- [59] Rui Qian, Divyansh Garg, Yan Wang, Yurong You, Serge Be-longie, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger, and Wei-Lun Chao. End-to-end pseudo-lidar for image-based 3d object detection. In *CVPR*, 2020. 2
- [60] Zequn Qin and Xi Li. Monoground: Detecting monocular 3d objects from the ground. In *CVPR*, 2022. 2
- [61] Zengyi Qin, Jinglu Wang, and Yan Lu. Monogmet: A geometric reasoning network for monocular 3d object localization. In *AAAI*, 2019. 2
- [62] Zengyi Qin, Jinglu Wang, and Yan Lu. Triangulation learning network: from monocular to stereo 3d object detection. In *CVPR*, 2019. 2
- [63] Cody Reading, Ali Harakeh, Julia Chae, and Steven L Waslander. Categorical depth distribution network for monocular 3d object detection. In *CVPR*, 2021. 2, 3, 13
- [64] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 2015. 2
- [65] Thomas Roddick, Alex Kendall, and Roberto Cipolla. Orthographic feature transform for monocular 3d object detection. In *BMVC*, 2019. 2
- [66] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In *WACV*, 2022. 3
- [67] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *CVPR*, 2020. 6, 7, 13
- [68] Xuepeng Shi, Zhixiang Chen, and Tae-Kyun Kim. Distance-normalized unified representation for monocular 3d object detection. In *ECCV*, 2020. 2
- [69] Andrea Simonelli, Samuel Rota Bulò, Lorenzo Porzi, Peter Kotschieder, and Elisa Ricci. Are we missing confidence in pseudo-lidar methods for monocular 3d object detection? In *ICCV*, 2021. 2, 7
- [70] Andrea Simonelli, Samuel Rota Bulò, Lorenzo Porzi, Manuel López-Antequera, and Peter Kotschieder. Disen-

- tangling monocular 3d object detection. In *ICCV*, 2019. 2, 12
- [71] Andrea Simonelli, Samuel Rota Bulo, Lorenzo Porzi, Elisa Ricci, and Peter Kotschieder. Towards generalization across depth for monocular 3d object detection. In *ECCV*, 2020. 2
- [72] Vishwanath A Sindagi, Yin Zhou, and Oncel Tuzel. Mvx-net: Multimodal voxelnet for 3d object detection. In *ICRA*, 2019. 13
- [73] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. *arXiv:1708.07120*, 2019. 4
- [74] Jiaming Sun, Linghao Chen, Yiming Xie, Siyu Zhang, Qin-hong Jiang, Xiaowei Zhou, and Hujun Bao. Disp r-cnn: Stereo 3d object detection via shape prior guided instance disparity estimation. In *CVPR*, 2020. 2
- [75] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *ICCV*, 2020. 1, 14
- [76] OpenPCDet Development Team. Openpcdet: An open-source toolbox for 3d object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet>, 2020. 3
- [77] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, 2019. 2
- [78] Li Wang, Liang Du, Xiaoqing Ye, Yanwei Fu, Guodong Guo, Xiangyang Xue, Jianfeng Feng, and Li Zhang. Depth-conditioned dynamic message propagation for monocular 3d object detection. In *CVPR*, 2021. 2
- [79] Tai Wang, Jiangmiao Pang, and Dahua Lin. Monocular 3d object detection with depth from motion. In *ECCV*, 2022. 2
- [80] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective. In *CoRL*, 2022. 2, 3, 5, 7
- [81] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *ICCV Workshop*, 2021. 2, 3, 4, 5, 7, 8
- [82] Xinlong Wang, Wei Yin, Tao Kong, Yuning Jiang, Lei Li, and Chunhua Shen. Task-aware monocular depth estimation for 3d object detection. In *AAAI*, 2020. 2
- [83] 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 *CVPR*, 2019. 1, 2
- [84] Yue Wang, Vitor Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *CoRL*, 2021. 2
- [85] Xinshuo Weng and Kris Kitani. Monocular 3d object detection with pseudo-lidar point cloud. In *ICCV Workshops*, 2019. 2
- [86] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 1
- [87] Bin Xu and Zhenzhong Chen. Multi-level fusion based 3d object detection from monocular images. In *CVPR*, 2018. 2
- [88] Zhenbo Xu, Wei Zhang, Xiaoqing Ye, Xiao Tan, Wei Yang, Shilei Wen, Errui Ding, Ajin Meng, and Liusheng Huang. Zoomnet: Part-aware adaptive zooming neural network for 3d object detection. In *AAAI*, 2020. 2
- [89] 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. In *ICLR*, 2020. 2
- [90] Fisher Yu, Dequan Wang, Evan Shelhamer, and Trevor Darrell. Deep layer aggregation. In *CVPR*, 2018. 3
- [91] Renrui Zhang, Han Qiu, Tai Wang, Xuanzhuo Xu, Ziyu Guo, Yu Qiao, Peng Gao, and Hongsheng Li. Monodetr: Depth-aware transformer for monocular 3d object detection. *arXiv preprint arXiv:2203.13310*, 2022. 3
- [92] Yunpeng Zhang, Jiwen Lu, and Jie Zhou. Objects are different: Flexible monocular 3d object detection. In *CVPR*, 2021. 14
- [93] Yinmin Zhang, Xinzhu Ma, Shuai Yi, Jun Hou, Zhihui Wang, Wanli Ouyang, and Dan Xu. Learning geometry-guided depth via projective modeling for monocular 3d object detection. *arXiv preprint arXiv:2107.13931*, 2021. 2
- [94] Yunpeng Zhang, Wenzhao Zheng, Zheng Zhu, Guan Huang, Dalong Du, Jie Zhou, and Jiwen Lu. Dimension embeddings for monocular 3d object detection. In *CVPR*, 2022. 2
- [95] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 1, 2, 3
- [96] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets v2: More deformable, better results. In *CVPR*, 2019. 3, 14