# Label-Free Model Failure Detection for Lidar-based Point Cloud Segmentation

Daniel Bogdoll<sup>†‡\*</sup>, Finn Sartoris<sup>†‡\*</sup>, Vincent Geppert<sup>†‡\*</sup>, Svetlana Pavlitska<sup>†‡</sup>, and J. Marius Zöllner<sup>†‡</sup> <sup>†</sup>FZI Research Center for Information Technology, Germany bogdoll@fzi.de

<sup>‡</sup>Karlsruhe Institute of Technology, Germany

Abstract-Autonomous vehicles drive millions of miles on the road each year. Under such circumstances, deployed machine learning models are prone to failure both in seemingly normal situations and in the presence of outliers. However, in the training phase, they are only evaluated on small validation and test sets, which are unable to reveal model failures due to their limited scenario coverage. While it is difficult and expensive to acquire large and representative labeled datasets for evaluation, large-scale unlabeled datasets are typically available. In this work, we introduce label-free model failure detection for lidar-based point cloud segmentation, taking advantage of the abundance of unlabelled data available. We leverage different data characteristics by training a supervised and self-supervised stream for the same task to detect failure modes. We perform a large-scale qualitative analysis and present LidarCODA, the first publicly available dataset with labeled anomalies in real-world lidar data, for an extensive quantitative analysis.

#### I. INTRODUCTION

In machine learning, utilizing a fully labeled dataset for training, validation, and testing is typical. In autonomous driving, 70 - 85 % of the data is usually reserved for training, leaving only 15 - 30 % for both validation and testing [1]–[3]. These small evaluation datasets stand in stark contrast to the millions of miles driven on public roads during deployment [4]. As a result, many failure modes of machine learning models, be it in seemingly normal situations or due to corner cases, are not captured in the evaluation sets. As large-scale unlabeled fleet data is generally available due to the much-reduced cost [5]–[7], there is an untapped potential to use this data for the evaluation of machine learning models.

In order to detect failure modes, there are many active research areas. Active learning [8] is concerned with continuously enriching training data by querying samples from a set of unlabeled data points. Discrepancies between different sensor systems can also be used to query samples [9]. In error estimation, many approaches try to utilize unlabeled test sets to evaluate models [10]. Label refinement compares given labels, e.g., by an auto-labeling process, with new proposals [11]. All of these methods have in common that they utilize or compare two or more different results for the same task. However, we are unaware of approaches that take advantage of different training paradigms to detect model failures. In this work, we introduce the concept of complementary learning to leverage different data characteristics of the training dataset. We train



Fig. 1. **Model Failure Detection.** The top left point cloud shows a *supervised* and the top right a *self-supervised* motion segmentation. The supervised model falsely classifies the pedestrian in the front left as static. Our approach exposes this model failure, as highlighted in red in the bottom image.

two models, one supervised and the other self-supervised, to detect model failures. Our main contributions are:

- Label-free evaluation for lidar-based point cloud segmentation based on *complementary learning* with a supervised and a self-supervised model
- *LidarCODA*, the first publicly available real-world anomaly dataset with labeled lidar data, to quantitatively evaluate the sensitivity of the approach to outliers

All code to run our experiments and generate the Lidar-CODA dataset is available on GitHub<sup>1</sup>.

#### II. RELATED WORK

The concept of comparing the outputs of two or more neural networks was already introduced in 1994 by Cohn et al., where they queried samples for active learning based on the disagreement between neural networks [12]. Since then, the variability in model predictions has been widely used to detect anomalies or errors. Ensemble diversity is especially well studied, as it was shown to lead to better performance [13], robustness [14], uncertainty quantification [15], and detection of outliers or distribution shifts [16]–[18]. While no uniform metric for ensemble diversity exists, measures like disagreement of models, double fault measure, or output correlation are widely used [19]. Ensemble diversity can be implicitly enhanced

<sup>1</sup>github.com/daniel-bogdoll/model\_contradictions

<sup>&</sup>lt;sup>\*</sup>These authors contributed equally



Fig. 2. **Overview.** Given point clouds, we derive semantic motion labels in a supervised (sv) fashion and perform ground segmentation and derive predictive motion labels in a self-supervised (ssv) fashion. Subsequently, we perform point-wise discrepancy detection and classify potential model failures.

via random initialization [15], noise injection or dropout, or explicitly via bagging, boosting, or stacking. Compared to ensembles, mixtures of experts [20] enforce higher model specialization and thus more component diversity, leading to better detection of out-of-distribution data [21].

The approaches mentioned above involve a combination of several neural networks with similar or identical architectures. Active learning is another research field interested in the detection of model failures. Here, uncertainty derived from ensembles is resource-intensive and thus only rarely used as part of a querying strategy [22]–[24]. Similarly, disagreements in a query-by-committee setting can be used to select samples [25]. In autonomous driving, also contradicting detections from sensors can be used as triggers, e.g., when radar and camera detections do not match [9]. Discrepancies between teacher and student models, typically known from knowledge distillation, can also be utilized [26]. As test sets are often small and not representative, directly estimating the accuracy of a model with only unlabeled data is of high interest [10], [27]-[29]. Here, we often see simple classification tasks or approaches that estimate an overall error that cannot be applied to individual samples. In some cases, generated pseudo-labels are utilized for further training steps [30], [31].

Disagreements can also be used for detecting erroneous labels. Ground truth labels in large vision datasets are often error-prone when auto-labeling processes based on large models are employed [32]. Detecting label errors with disagreements can be done by predicting a novel or refined label, and uncertainties can be generated by predicting multiple such labels [11], [33], [34]. This way, also noisy labels introduced by human errors can be detected [35].

Robustness during deployment is often achieved with sensor fusion, which, quite differently, purposefully aims to complement the weaknesses of one sensor with the strengths of another. Thus, disagreements are both typical and wanted, with the aim of resolving them [36]. However, also data from a single sensor can be split into multiple streams to increase robustness. For example, object detection can be improved by combining appearance and geometry [37] or temporality and geometry [38], [39]. In performance monitoring [40], [41], but also in outlier or anomaly detection [42], typically, a primary model performing a regular task is accompanied by a learned or model-based module that provides some sort of uncertainty for the results of the regular task. **Research Gap.** Many of the analyzed works utilizing disagreements deal with toy problems and only analyze classification tasks, which are not sufficient to truly understand the shortcomings of a model that is designed for the complex task of autonomous driving. Many works analyze model outputs of the same architecture, leveraging differences during training. However, this way, the same data characteristics are being used during training. Disagreement-based approaches in designing triggers for active learning [9] and increasing robustness during deployment [37] are most similar to our approach, but these industry demonstrations are not accompanied by scientific works and are thus hard to evaluate. Finally, to the best of our knowledge, no work exists that utilizes different training paradigms to detect model failures through disagreements.

#### III. METHOD

To assess model performance without labeled validation or test sets, we perform complementary training for the same task in order to detect model failures and classify challenging scenarios. The ability to detect model failures is based on the intuition that different training paradigms leverage different data characteristics from the same training data set. We demonstrate this approach with the motion segmentation of lidar point clouds for autonomous driving. As shown in Fig. 2, we first derive motion labels in a supervised and selfsupervised fashion. Here, the first paradigm leverages human knowledge through labels, given only context from static scenes. On the other hand, the second paradigm leverages temporal information inherent in the data. Typically, these paradigms are combined either in a pre-training context [43] or with a combined loss during learning [44]. Based on a pointwise comparison, we detect discrepancies and cluster them for better interpretation. Finally, an oracle examines and classifies the model failures to better understand challenging situations.

#### A. Supervised Semantic Motion Labels

We derive semantic motion labels with a supervised **semantic segmentation** model [45] to determine whether a point belongs to a static or dynamic class. Some classes do not provide clear information about the motion state of the points, e.g., points assigned to the class *cyclist* at a traffic light may be static in the case of a red light and dynamic in the case of a green light. By also performing supervised **motion segmentation** [46], we further subdivide classes into *semantic motion labels*, as shown in Fig. 3.

Manuscript 411 submitted to 2025 36th IEEE Intelligent Vehicles Symposium (IV). Received February 2, 2025.



Fig. 3. **Supervised Semantic Motion Labels.** The left semantic segmentation allows no distinction between the parked car at the bottom left and the moving car at the top right. The middle image shows a motion segmentation, where the parked car was classified as static, and the moving car as dynamic. Finally, the right image shows the fused semantic motion labels.

#### B. Self-Supervised Predictive Motion Labels

In order to predict motion labels for a given point cloud, we first filter out the ground [47] to focus on objects in the scene, a common pre-processing step of scene flow models [48]–[52]. Based on self-supervised flow prediction [51] of the remaining points, we aim to derive motion labels, indicating whether a point is static or dynamic. The model takes consecutive point clouds as input and predicts the future motion for each lidar point in the form of a 3D displacement vector. The scene flow model does not distinguish between the point's own motion and the observer's ego-motion and represents the overall motion of a point between two consecutive frames. In order to derive relative displacements, we need to correct for the ego-motion. This can be done by leveraging or learning odometry information [53]. After predicting the future point cloud  $X_{t+1} = X_t + f_t$ , we apply the learned rigid body transformation  $T_{t+1 \rightarrow t}$  of an odometry model, transforming the predicted point cloud back into the coordinate system of  $X_t$ . This gives the future point cloud  $X_{t+1}$ , which contains only the predicted relative motion without the ego motion.

As a result, static objects line up closely with the original data of  $X_t$ , and only dynamic objects show a predicted displacement, as shown in Fig. 4 a).

**Two-Stage Clustering.** An analysis of the velocity values of the flow predictions showed that separating static from dynamic classes is infeasible in a point-wise fashion, as a strong overlap exists. What we found, however, is a significant difference when considering instance-wise normalized standard deviations, as shown in Fig. 5. While we performed this analysis with ground-truth labels, in unlabeled data, the necessity arises to form instance clusters.

As shown in Fig. 4 b), we utilize DBSCAN [54] to cluster the point clouds spatially. A cluster is classified as *potentially dynamic* if the normalized standard deviation of the cluster's velocity is below 0.12, as identified through a grid search. To further reduce false positives, potentially dynamic points are clustered based on their flow vectors. Points with a similar flow



Fig. 4. Self-Supervised Predictive Motion Labels. The first image shows the original point cloud in green and the point cloud transformed by the scene flow model and compensated by the ego-motion in red. The second and third images show the result of the spatial and flow-based clustering, respectively. The fourth image shows the final predictive motion labels, with dynamic points in red and static points in green.

are clustered, causing fewer static clusters to be incorrectly classified as dynamic. Next, spatially separated points are grouped together based on their flow, as shown in Fig. 4 c). Here, the points on the left and right edges belonging to static objects now form a cluster, which changes the distribution of flow vectors within this cluster, causing it to be classified as static. Finally, the newly found clusters are classified as *dynamic* if the speed of the cluster is above 4 km/h, based on typical velocity profiles of pedestrians, as visible in Fig. 4 d).

#### C. Discrepancy Detection and Failure Classification

After obtaining motion labels from both the *supervised* and the *self-supervised* stream, contradictions between the labels are detected, as shown in Fig. 2. Only the lidar points per frame for which both streams predicted a label are considered. Given a semantic and a predictive motion label for each lidar point, there exist four categories: Points which both models deem static (green •); points which both models deem dynamic (blue •); points where the *supervised* stream predicts a static point and the *self-supervised* a dynamic one (red •), and points where the *supervised* a static one (yellow •). Finally, we cluster instances with contradicting labels so an oracle can classify both complete scenes and single instances.

#### D. Implementation Details

For all models shown in Fig. 2, we utilized publicly available models and model architectures to demonstrate the modularity of our approach. We trained the supervised semantic segmentation model SalsaNext [45] and the supervised motion segmentation model of Chen et al. [46] on the KITTI-360 dataset [55], [56], as it is a large dataset that contains semantic labels, motion labels, and odometry data. The training was performed on an NVIDIA RTX A6000. Hyperparameters were taken from the original papers [45], [46]. For the remaining three models, we used available pre-trained model weights. For ground segmentation, we deployed GndNet [47]. For selfsupervised scene flow estimation, we used FlowStep3D [51]. For the self-supervised odometry model, we deployed De-LORA [53]. More details can be found in [57].



Fig. 5. **Self-Supervised Label Generation**. The left graph shows that the magnitude of point-wise flow vectors is insufficient to distinguish between dynamic and static points. The boxplot on the right shows that the normalized standard deviation per instance is significantly lower for dynamic instances.

#### IV. EVALUATION

Model failures can occur in seemingly normal situations [58]–[61], but models are especially prone to failure in the presence of anomalies or outliers [58], [62]–[67]. In Section IV-A, we first analyze our approach given regular data. Here, no ground truth for especially challenging situations is available. Thus, we perform a qualitative evaluation by manually analyzing our method on over 20,000 frames.

In Section IV-B, the second part of the evaluation, we focus on anomalies. Here, ground truth is available. As CODA [68], the largest real-world dataset with anomalies, only provides labels in camera space, we first introduce *LidarCODA*, an extension that provides labels in lidar space. Subsequently, we perform a quantitative evaluation to better understand the sensitivity of our method towards outliers.

Fig. 6 provides an overview of the discrepancy detection, the second-to-last step of our approach, as shown in Fig. 2. For **regular scenarios**, the majority of points are predicted as static by both models and around 5 % of the points show model contradictions. For **outlier scenarios**, many more disagreements take place. However, a large variety among the subsets of *LidarCODA* can be observed. In the following, we first qualitatively examine scenarios in which the models disagree. Subsequently, we quantitatively examine the correlation between disagreements and anomlies in the environment.

#### A. Regular Scenarios

Regular scenarios represent the majority of miles driven during deployment, and it is important to understand situations in which models fail. However, the evaluation under regular scenarios is challenging, as no ground truth is available. Thus, we manually examined our approach on over 20,000 frames.

**Evaluation Data.** For training, KITTI-360 and several sequences of the KITTI Odometry were used. To minimize perceptual failures due to a domain shift, we perform the qualitative evaluation on the remaining 20,350 frames of the KITTI Odometry sequences 11-21. The datasets are closely related, as both were captured in Karlsruhe, Germany with a Velodyne HDL-64E lidar.



Fig. 6. **Discrepancy Detection.** Given mostly normal data in the KITTI Odometry dataset, the two streams agree on 95 % of the data (green, blue). For outlier scenarios, as provided by LidarCODA, many more discrepancies (red, yellow) are detected.

**Evaluation.** As shown in Fig. 2, the final stage of our approach is the classification of model failure modes. Our qualitative evaluation is performed by a human oracle. For visual inspection, we utilize lidar points mapped onto the corresponding RGB image for an improved scene understanding, as shown in Fig. 7 with the color scheme introduced in Section III-C. In most cases, both streams were correctly consistent. In the following, we qualitatively present representative examples of detected model failures and highlight those that occurred frequently, suggesting general model flaws.

First, we discuss model failures of the supervised stream, which is the model under test in most cases. Here, the two streams contradict each other. We show representative examples in Fig. 7. Scene 1 shows a turning car and two moving bicyclists, where one bicyclist is wrongly labeled as static by the supervised stream. Scene 2 contains two walking pedestrians that are wrongly classified as static by the supervised stream. Scene 3 shows a parking car misclassified as dynamic by the supervised stream. Scenes 4 and 5 show a car moving slowly and a car moving backward, respectively. These cases demonstrate effectively that our approach enables the detection of regular but challenging scenarios that lead to model failures. Such model failures remain undetected in small evaluation datasets. We find various weak points in each stream, characterized by repeated occurrence. Specifically, the supervised model under test has weaknesses in distinguishing between dynamic and static objects in specific situations, e.g., at red lights or when a car is parked directly in front of the ego vehicle. Examples of such situations are given in scenes 6 and 7 of Fig. 7.

Next, we analyze scenarios where self-supervised model failures occur, detected by correct predictions of the supervised stream. Fig. 8 shows representative scenes. Scene 1 contains two distant pedestrians walking, wrongly classified as static by the self-supervised stream. In scene 2, a parked car is misclassified as dynamic by the self-supervised stream. Scenes 3, 4, and 5 show walking pedestrians or moving cars incorrectly classified as dynamic. These cases demonstrate that our approach enables the detection of challenging temporal scenarios. The self-supervised stream classifies an above-average number of objects as dynamic when the ego-vehicle turns or goes over speed bumps. An example is shown in Fig. 8, where in scene 6, the vehicle turns, and in scene



Fig. 7. **Supervised Model Failures**. These exemplary images show model failures of the *supervised* stream, which can be detected due to contradicting outputs of the *self-supervised* model.

8, it drives over a speed bump. Another weak point is fast oncoming vehicles on highways, often classified as static, as seen in scene 7. Finally, a common weakness is small clusters on the side, which are incorrectly classified as dynamic, as in scene 9, where a window is classified as dynamic.

In rare cases, both models are incorrectly consistent, i.e., both streams agree, but the label is incorrect in both cases. We show examples in Fig. 9. Here, the left scene shows two walking pedestrians that are incorrectly classified as static, and the right scene shows a parked car that is classified as dynamic.

#### **B.** Outlier Scenarios

Scenarios with anomalies or outliers are known to lead to model failures [69]. So far, the evaluation of lidar-based anomaly detection models has been challenging. Evaluation datasets are either unavailable [70], [71] or utilize known classes but exclude them from training data [72]. Thus, we first introduce *LidarCODA*, a real-world anomaly dataset with labeled anomalies in lidar space. Subsequently, we quantitatively evaluate how sensitive our approach is to anomalies in the environment.

**Evaluation Data.** For the evaluation, we utilized data from the CODA dataset [68], the only publicly available real-world dataset that provides lidar data and includes anomalies [73]. The CODA dataset provides anomaly labels for objects based on three existing data sets: KITTI [74], ONCE [5], and nuScenes [1]. CODA defines an anomaly as an object that "blocks or is about to block a potential path of the self-driving vehicle" [68] and/or "does not belong to any of the common classes of autonomous driving benchmarks" [68]. While this risk-aware definition is not always in line with the methodology of our approach, where objects that block the path in front of the ego vehicle are not necessarily hard to segment or predict, we chose it because CODA is currently the most relevant real-world anomaly dataset [73].



Fig. 8. Self-Supervised Model Failures. These exemplary images show model failures of the *self-supervised* model, which can be detected due to contradicting outputs of the *supervised* stream.

For the CODA-KITTI split, the authors of CODA manually reviewed all *misc* labels available in the ground truth and relabeled some as anomalies according to a labeling policy. This enables us to examine our approach with only a small domain gap quantitatively. For CODA-nuScenes, the authors similarly adopted available annotations in a manual process. Finally, for CODA-ONCE, they deployed an automated anomaly detection approach, making this subset the most relevant. CODA includes 1,500 scenes with a total of 5,937 anomaly instances. Among those, 4,746 belong to the superclass *traffic\_facility*, followed by 929 *vehicle* and 197 *obstruction* instances. With 396, most *vehicle* instances can be found in CODA-KITTI.

The CODA dataset provides anomaly labels only in the form of 2D bounding boxes in image space. However, point-wise labels in 3D lidar space are necessary to utilize CODA for our approach. Therefore, we present *LidarCODA*, a dataset based on the CODA dataset [68] for evaluation. Based on a frustumbased filter, subsequent clustering, and manual inspection, we transfer the original 2D labels from image space into refined, point-wise 3D labels that go beyond the coarse characteristic of the provided bounding boxes. More details can be found in [75]. *LidarCODA* is the first real-world anomaly dataset for *object-level* anomalies [62] with annotated lidar data, as shown in Fig. 10. Here, the different lidar systems utilized also become clearly visible. Due to the sparse point cloud



Fig. 9. **Simultaneous Model Failures.** Examples where both streams produce model failures. Both cases are misclassified and are, therefore, consistent.

TABLE I							
LIDARCODA SUBSETS. EVALUATION OF OUR APPROACH ON							
LIDARCODA AND ITS THREE SUBSETS.							

Dataset	#Frames	mIoU ↑	AP ↑	AR ↑	<b>F1</b> ↑
LidarCODA	1,412	8.9	13.2	26.2	17.5
LidarCODA-ONCE LidarCODA-KITTI LidarCODA-nuScenes	1,034 307 71	<u>8.9</u> <b>10.9</b> 0.4	<b>14.0</b> <u>13.3</u> 0.7	<u>27.1</u> <b>29.0</b> 1.0	<b>18.4</b> <u>18.3</u> 0.8

of nuScenes, many small or distant labeled anomalies in the image space are only covered by a few or no lidar points.



Fig. 10. LidarCODA. Annotated lidar scenes in the three data splits ONCE, KITTI, and nuScenes, from left to right. Anomalies are shown in red.

**Evaluation.** To quantitatively evaluate the correlation between model disagreements and anomalies present in the environment, we utilize the standard metrics mean Intersection over Union (mIoU), Average Precision (AP), Average Recall (AR), and F1 score, as shown in Tables I and II. For a fair evaluation, we consider all points of the lidar point cloud, even if our approach did not label individual points, e.g., because they were filtered out during pre-processing. Such cases were counted as false negatives if an anomaly was missed.

To better understand the suitability of LidarCODA due to introduced domain shifts, either due to new environments or due to new sensor setups, we perform experiments on the individual subsets, as shown in Table I. The results clearly show that our approach struggles with the nuScenes subset, which is primarily due to the large domain shift w.r.t to the sensor setup. Our approach is more sensitive towards anomalies for the subsets ONCE and KITTI. This is reflected in Fig. 6, where the CODA subsets also show much higher detection rates compared to the analysis with regular scenarios. This aligns with the much higher number of anomalies, even though the subsets reveal strongly varying behavior patterns. We generally observe a moderate correlation between model disagreements and anomalies present in the environment.

Next, we investigate the sensitivity of our approach towards the superclasses provided by CODA by evaluating points that have labels assigned by both our approach and the ground truth. As shown in Table II, our approach shows different levels of sensitivity given different types of anomalies, being most sensitive to *cyclists* and objects of the *misc* class. The *misc* class consists of objects that are "unrecognizable or difficult to categorize" [68]. These results align well with our approach, where *cyclist* instances, which are hard to predict by the self-supervised stream, and *misc* instances, which are rare and thus hard to classify by the supervised stream, lead to model disagreements, as the complementary model does not display similar weaknesses in both cases.

 
 TABLE II

 CODA Superclasses. Evaluation of our approach on seven superclasses.

Superclass	#Instances	mIoU ↑	AP ↑	AR ↑	<b>F1</b> ↑
Pedestrian	16	33.9	44.1	37.3	40.4
Cyclist	22	41.6	58.3	49.5	53.5
Vehicle	736	33.0	48.3	41.0	44.4
Animal	5	0.0	0.0	0.0	0.0
Traffic facility	3,360	28.6	39.9	33.9	36.7
Obstruction	125	20.5	34.0	22.9	27.4
Misc	15	<u>36.7</u>	60.7	37.9	<u>46.7</u>

#### V. CONCLUSION

In this work, we have presented a label-free approach to detect model failures for motion segmentation of point clouds. We leverage complementary training paradigms to detect contradicting outputs on the same task, consisting of a *supervised* stream for semantic motion labels and a *selfsupervised* stream for predictive motion labels. This way, we detect model failures far beyond the limited scope of a small evaluation dataset. We first inspect model failures in regular scenarios. By manually analyzing over 20,000 frames qualitatively, we detect model failures in seemingly normal scenarios and are able to categorize frequently occurring cases. In regular scenarios, our method disregards 95 % of the data as typical, which makes human analysis of the remaining 5 % feasible even for larger datasets.

Second, we analyzed the sensitivity of our approach towards scenarios with anomalies, as defined by the CODA dataset, and showed a moderate sensitivity. In order to quantitatively examine our approach, we first introduced *LidarCODA*, the first real-world dataset with labeled anomalies in lidar space. We demonstrated that our approach shows an increased sensitivity to often hard-to-detect and hard-to-predict bicycles as well as hard-to-categorize objects.

Our approach effectively unveils model failures far beyond those that can be detected with small evaluation datasets. This leads to an increased understanding of the model performance in large-scale deployments, leveraging abundantly available unlabeled data. Model failures detected by our approach can be utilized to collect additional training data representing both static and temporally challenging scenarios.

**Limitations.** When both streams are wrong, model failures go undetected. This behavior is known and unavoidable [18], [30] and can be mitigated by deploying multiple approaches or triggers to detect challenging scenarios [9]. In addition, our self-supervised stream depends on a clustering strategy, which can lead to faulty clusters and misclassifications.

#### ACKNOWLEDGMENT

This work results partly from the XXXXXX project supported by the XXXXXXX.

#### REFERENCES

- [1] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [2] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov, "Scalability in perception for autonomous driving: Waymo open dataset," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [3] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell, "Bdd100k: A diverse driving dataset for heterogeneous multitask learning," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [4] R. Bishop, "60 million miles and counting: Robotaxis shift into high gear," https://www.forbes.com/sites/richardbishop1/2024/07/27/60million-miles-and-counting-robotaxis-shift-into-high-gear/, 2024, accessed: 2025-01-29.
- [5] J. Mao, M. Niu, C. Jiang, H. Liang, J. Chen, X. Liang, Y. Li, C. Ye, W. Zhang, Z. Li, J. Yu, H. Xu, and C. Xu, "One million scenes for autonomous driving: Once dataset," in *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [6] J. Han, X. Liang, H. Xu, K. Chen, L. Hong, J. Mao, C. Ye, W. Zhang, Z. Li, X. Liang, and C. Xu, "Soda10m: A large-scale 2d self/semisupervised object detection dataset for autonomous driving," in *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [7] Y. Li, Z. Li, N. Chen, M. Gong, Z. Lyu, Z. Wang, P. Jiang, and C. Feng, "Multiagent multitraversal multimodal self-driving: Open mars dataset," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [8] P. Kumar and A. Gupta, "Active Learning Query Strategies for Classification, Regression, and Clustering: A Survey," *Journal of Computer Science and Technology*, vol. 35, no. 4, 2020.
- [9] A. Karpathy, "Tesla keynote: Cvpr 2021 workshop on autonomous driving," https://www.youtube.com/watch?v=g6bOwQdCJrc, 2021, accessed: 2024-06-13.
- [10] W. Deng and L. Zheng, "Are Labels Always Necessary for Classifier Accuracy Evaluation?" in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [11] M. Schubert, T. Riedlinger, K. Kahl, D. Kröll, S. Schoenen, S. Šegvić, and M. Rottmann, "Identifying Label Errors in Object Detection Datasets by Loss Inspection," in Winter Conference on Applications of Computer Vision (WACV), 2024.
- [12] D. Cohn, L. Atlas, and R. Ladner, "Improving generalization with active learning," *Machine Learning*, vol. 15, no. 2, 1994.
- [13] T. G. Dietterich, "Ensemble methods in machine learning," in *International workshop on multiple classifier systems*. Springer, 2000.
- [14] T. Pang, K. Xu, C. Du, N. Chen, and J. Zhu, "Improving adversarial robustness via promoting ensemble diversity," in *International Conference* on Machine Learning (ICML), 2019.
- [15] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Conference* on Neural Information Processing Systems (NeurIPS), 2017.
- [16] H. A. Mehrtens, C. González, and A. Mukhopadhyay, "Improving robustness and calibration in ensembles with diversity regularization," in *German Conference on Pattern Recognition (GCPR)*, 2022.
- [17] M. Pagliardini, M. Jaggi, F. Fleuret, and S. P. Karimireddy, "Agree to disagree: Diversity through disagreement for better transferability," in *International Conference on Learning Representations (ICLR)*, 2023.
- [18] L. Wang, J. Wang, Y. Zheng, S. Jain, C.-C. M. Yeh, Z. Zhuang, J. Ebrahimi, and W. Zhang, "Learning from Disagreement for Event Detection," in *International Conference on Big Data*, 2022.
- [19] L. I. Kuncheva and C. J. Whitaker, "Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy," *Machine learning*, vol. 51, 2003.
- [20] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton, "Adaptive mixtures of local experts," *Neural computation*, vol. 3, no. 1, 1991.
- [21] S. Pavlitskaya, C. Hubschneider, and M. Weber, "Evaluating mixture-ofexperts architectures for network aggregation," in *Deep Neural Networks* and Data for Automated Driving: Robustness, Uncertainty Quantification, and Insights Towards Safety, 2022.

- [22] C. Schröder, A. Niekler, and M. Potthast, "Revisiting Uncertainty-based Query Strategies for Active Learning with Transformers," in *Findings* of the Association for Computational Linguistics, 2022.
- [23] A. Roitberg, Z. Al-Halah, and R. Stiefelhagen, "Informed democracy: Voting-based novelty detection for action recognition," in *British Machine Vision Conference (BMVC)*, 2018.
- [24] S. W. Yahaya, A. Lotfi, and M. Mahmud, "A consensus novelty detection ensemble approach for anomaly detection in activities of daily living," *Applied Soft Computing*, vol. 83, 2019.
- [25] H. Hino and S. Eguchi, "Active Learning by Query by Committee with Robust Divergences," *Information Geometry*, vol. 6, 2023.
- [26] G. Chen, W. Choi, X. Yu, T. Han, and M. Chandraker, "Learning Efficient Object Detection Models with Knowledge Distillation," in *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [27] R. Peng, H. Zou, H. Wang, Y. Zeng, Z. Huang, and J. Zhao, "Energybased Automated Model Evaluation," in *International Conference on Learning Representations (ICLR)*, 2024.
- [28] C. Baek, Y. Jiang, A. Raghunathan, and Z. Kolter, "Agreement-on-the-Line: Predicting the Performance of Neural Networks under Distribution Shift," in *Conference on Neural Information Processing Systems* (*NeurIPS*), 2022.
- [29] J. Chen, F. Liu, B. Avci, X. Wu, Y. Liang, and S. Jha, "Detecting Errors and Estimating Accuracy on Unlabeled Data with Self-training Ensembles," in *Conference on Neural Information Processing Systems* (*NeurIPS*), 2021.
- [30] J. Wang, L. Wang, Y. Zheng, C.-C. M. Yeh, S. Jain, and W. Zhang, "Learning-From-Disagreement: A Model Comparison and Visual Analytics Framework," *Transactions on Visualization and Computer Graphics*, vol. 29, no. 9, 2023.
- [31] Y. Yu, Z. Yang, A. Wei, Y. Ma, and J. Steinhardt, "Predicting Out-of-Distribution Error with the Projection Norm," in *International Conference on Machine Learning (ICML)*, 2022.
- [32] K. Chachula, J. Lyskawa, B. Olber, P. Fratczak, A. Popowicz, and K. Radlak, "Combating noisy labels in object detection datasets," arXiv:2211.13993, 2023.
- [33] Z. Hu, K. Gao, X. Zhang, J. Wang, H. Wang, and J. Han, "Probability Differential-Based Class Label Noise Purification for Object Detection in Aerial Images," *Geoscience and Remote Sensing Letters*, vol. 19, 2022.
- [34] A. Bär, J. Uhrig, J. P. Umesh, M. Cordts, and T. Fingscheidt, "A Novel Benchmark for Refinement of Noisy Localization Labels in Autolabeled Datasets for Object Detection," in *Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2023.
- [35] C. Northcutt, L. Jiang, and I. Chuang, "Confident Learning: Estimating Uncertainty in Dataset Labels," *Journal of Artificial Intelligence Research*, vol. 70, 2021.
- [36] K. Kowol, M. Rottmann, S. Bracke, and H. Gottschalk, "YOdar: Uncertainty-based Sensor Fusion for Vehicle Detection with Camera and Radar Sensors," in *International Conference on Agents and Artificial Intelligence (ICAART)*, 2021.
- [37] A. Shashua, "Intel newsroom: CES 2021: Under the Hood," https:// youtube.com/watch?v=B7YNj66GxRA, 2021, accessed: 2024-06-13.
- [38] S. Baur, F. Moosmann, and A. Geiger, "Liso: Lidar-only self-supervised 3d object detection," in *European Conference on Computer Vision* (ECCV), 2024.
- [39] T. Lentsch, H. Caesar, and D. M. Gavrila, "UNION: Unsupervised 3D object detection using object appearance-based pseudo-classes," in *Conference on Neural Information Processing Systems (NeurIPS)*, 2024.
- [40] W. Shao, B. Li, W. Yu, J. Xu, and H. Wang, "When Is It Likely to Fail? Performance Monitor for Black-Box Trajectory Prediction Model," *Transactions on Automation Science and Engineering*, 2024.
- [41] C. Buerkle, F. Oboril, J. Burr, and K.-U. Scholl, "Safe perception - a hierarchical monitor approach," in *International Conference on Intelligent Transportation Systems (ITSC)*, 2022.
- [42] A. Delić, M. Grcić, and S. Šegvić, "Outlier detection by ensembling uncertainty with negative objectness," in *British Machine Vision Conference (BMVC)*, 2024.
- [43] T. Chen, S. Liu, S. Chang, Y. Cheng, L. Amini, and Z. Wang, "Adversarial Robustness: From Self-Supervised Pre-Training to Fine-Tuning," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [44] Y. Cheng, W. Wang, L. Jiang, and W. Macherey, "Self-supervised and Supervised Joint Training for Resource-rich Machine Translation," in *International Conference on Machine Learning (ICML)*, 2021.

- [45] T. Cortinhal, G. Tzelepis, and E. E. Aksoy, "Salsanext: Fast, uncertaintyaware semantic segmentation of lidar point clouds for autonomous driving," in Advances in Visual Computing, 2020.
- [46] X. Chen, S. Li, B. Mersch, L. Wiesmann, J. Gall, J. Behley, and C. Stachniss, "Moving Object Segmentation in 3D LiDAR Data: A Learning-based Approach Exploiting Sequential Data," *Robotics and Automation Letters*, 2021.
- [47] A. Paigwar, O. Erkent, D. Sierra-Gonzalez, and C. Laugier, "GndNet: Fast Ground Plane Estimation and Point Cloud Segmentation for Autonomous Vehicles," in *International Conference on Intelligent Robots* and Systems (IROS), 2020.
- [48] W. Wu, Z. Y. Wang, Z. Li, W. Liu, and L. Fuxin, "Pointpwc-net: Cost volume on point clouds for (self-) supervised scene flow estimation," in *European Conference on Computer Vision (ECCV)*, 2020.
- [49] H. Mittal, B. Okorn, and D. Held, "Just go with the flow: Self-supervised scene flow estimation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [50] I. Tishchenko, S. Lombardi, M. R. Oswald, and M. Pollefeys, "Selfsupervised learning of non-rigid residual flow and ego-motion," in *International Conference on 3D Vision (3DV)*, 2020.
- [51] Y. Kittenplon, Y. C. Eldar, and D. Raviv, "Flowstep3d: Model unrolling for self-supervised scene flow estimation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [52] S. Baur, D. Emmerichs, F. Moosmann, P. Pinggera, B. Ommer, and A. Geiger, "SLIM: Self-Supervised LiDAR Scene Flow and Motion Segmentation," in *International Conference on Computer Vision (ICCV)*, 2021.
- [53] J. Nubert, S. Khattak, and M. Hutter, "Self-supervised learning of lidar odometry for robotic applications," in *International Conference* on Robotics and Automation (ICRA), 2021.
- [54] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," in *International Conference on Knowledge Discovery and Data Mining*, 1996.
- [55] Y. Liao, J. Xie, and A. Geiger, "KITTI-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d," *Pattern Analysis and Machine Intelligence (PAMI)*, 2022.
- [56] J. Sanchez, "recoverkitti360label," https://github.com/JulesSanchez/ recoverKITTI360label/pull/3, 2022, accessed: 2024-06-26.
- [57] F. Sartoris, "Anomaly Detection in Lidar Data by Combining Supervised and Self-Supervised Methods," Bachelor's Thesis, Karlsruhe Institute of Technology (KIT), 2022.
- [58] F. Heidecker, J. Breitenstein, R. Kevin, L. Jonas, C. Stiller, T. Fingscheidt, and B. Sick, "An Application-Driven Conceptualization of Corner Cases for Perception in Highly Automated Driving," in *Intelligent Vehicles Symposium (IV)*, 2021.
- [59] J. Pfeil, J. Wieland, T. Michalke, and A. Theissler, "On Why the System Makes the Corner Case: AI-based Holistic Anomaly Detection for Autonomous Driving," in *Intelligent Vehicles Symposium (IV)*, 2022.
- [60] J. Zhou and J. Beyerer, "Corner cases in data-driven automated driving: Definitions, properties and solutions," in *Intelligent Vehicles Symposium* (IV), 2023.
- [61] F. Heidecker, M. Bieshaar, and B. Sick, "Corner Case Definition in Machine Learning Processes for the Perception of Highly Automated Driving," *AI Perspectives & Advances*, vol. 6, 2024.
- [62] J. Breitenstein, J. A. Termohlen, D. Lipinski, and T. Fingscheidt, "Systematization of Corner Cases for Visual Perception in Automated Driving," in *Intelligent Vehicles Symposium (IV)*, 2020.
- [63] J. Breitenstein, J.-A. Termöhlen, D. Lipinski, and T. Fingscheidt, "Corner cases for visual perception in automated driving: Some guidance on detection approaches," *arXiv preprint:2102.05897*, 2021.
- [64] D. Bogdoll, J. Breitenstein, F. Heidecker, M. Bieshaar, B. Sick, T. Fingscheidt, and J. M. Zöllner, "Description of Corner Cases in Automated Driving: Goals and Challenges," in *International Conference on Computer Vision (ICCV) Workshops*, 2021.
- [65] D. Bogdoll, M. Nitsche, and J. M. Zollner, "Anomaly Detection in Autonomous Driving: A Survey," in *Conference on Computer Vision* and Pattern Recognition (CVPR) Workshops, 2022.
- [66] L. Vater, M. Sonntag, J. Hiller, P. Schaudt, and L. Eckstein, "A Systematic Approach Towards the Definition of the Terms Edge Case and Corner Case for Automated Driving," in *International Conference on Electrical, Computer, Communications and Mechatronics Engineering* (ICECCME), 2023.

- [67] H. X. Liu and S. Feng, "Curse of rarity for autonomous vehicles," *Nature Communications*, vol. 15, no. 1, 2024.
- [68] K. Li, K. Chen, H. Wang, L. Hong, C. Ye, J. Han, Y. Chen, W. Zhang, C. Xu, D.-Y. Yeung, X. Liang, Z. Li, and H. Xu, "CODA: A Real-World Road Corner Case Dataset for Object Detection in Autonomous Driving," in *European Conference on Computer Vision (ECCV)*, 2022.
- [69] D. Bogdoll, M. Nitsche, and J. M. Zöllner, "Anomaly Detection in Autonomous Driving: A Survey," in *Conference on Computer Vision* and Pattern Recognition (CVPR) Workshops, 2022.
- [70] K. Wong, S. Wang, M. Ren, M. Liang, and R. Urtasun, "Identifying Unknown Instances for Autonomous Driving," in *Conference on Robot Learning CORL*, 2020.
- [71] L. Nunes, X. Chen, R. Marcuzzi, A. Osep, L. Leal-Taixé, C. Stachniss, and J. Behley, "Unsupervised class-agnostic instance segmentation of 3d lidar data for autonomous vehicles," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, 2022.
- [72] J. Cen, P. Yun, J. Cai, M. Wang, and M. Liu, "Open-set 3d object detection," *International Conference on 3D Vision (3DV)*, 2021.
- [73] D. Bogdoll, S. Uhlemeyer, K. Kowol, and J. M. Zöllner, "Perception datasets for anomaly detection in autonomous driving: A survey," in *Intelligent Vehicles Symposium (IV)*, 2023.
- [74] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [75] V. Geppert, "Anomaly Detection with Model Contradictions for Autonomous Driving," Bachelor Thesis, Karlsruhe Institute of Technology (KIT), 2023.