

Towards Defensive Autonomous Driving: Collecting and Probing Driving Demonstrations of Mixed Qualities

Jeongwoo Oh^{1*}, Gunmin Lee^{1*}, Jeongeun Park², Wooseok Oh¹, Jaeseok Heo¹, Hojun Chung¹,
Do Hyung Kim³, Byungkyu Park³, Chang-Gun Lee³, Sungjoon Choi², and Songhwai Oh¹

Abstract—Designing or learning an autonomous driving policy is undoubtedly a challenging task as the policy has to maintain its safety in all corner cases. In order to secure safety in autonomous driving, the ability to detect hazardous situations, which can be seen as an out-of-distribution (OOD) detection problem, becomes crucial. However, conventional datasets often only contain expert driving demonstrations, although some non-expert or uncommon driving behavior data are needed to implement a safety guaranteed autonomous driving platform. To this end, we present a dataset called the *R3 Driving Dataset*, composed of driving data with different qualities. The dataset categorizes abnormal driving behaviors into eight categories and 369 different detailed situations. The situations include dangerous lane changes and near-collision situations. To further enlighten how these abnormal driving behaviors can be detected, we utilize different uncertainty estimation and anomaly detection methods for the proposed dataset. From the results of the proposed experiment, it can be inferred that by using both uncertainty estimation and anomaly detection, most of the abnormal cases in the proposed dataset can be discriminated. <https://rllab-snu.github.io/projects/R3-Driving-Dataset/doc.html>

I. INTRODUCTION

It has been a challenging task to design a controller for autonomous driving. Conventional approaches to design a controller are based on rule-based methods [1]–[4], reinforcement learning (RL) [5]–[8], and imitation learning (IL) [9]–[12]. However, rule-based methods are not proficient at covering all exception cases due to the diversity of exceptions in the open road. In the case of RL-based methods, a stable and elaborate reward function is required to learn a controller. For an IL-based method, the algorithm exploits and tries to learn to behave similarly to the provided data. Since the IL-based controller follows the provided driving data, its performance can be improved further if provided with more good driving data. In addition, an IL method has an advantage over an RL method as it does not require a reward




Fig. 1: **Data Collection Platform.** The proposed Hyundai Ioniq platform is equipped with three Velodyne LiDARs (marked in blue), a Mobileye sensor (marked in red), and a combined GPS/IMU inertial navigation system (marked in green). The dotted lines indicate the occluded sensors.

function. These characteristics allow an IL-based method to show high performance.

One of the most critical shortcomings of IL is that the performance of the policy trained via IL cannot overcome the performance of the given driving data. In order to maximize the performance of a policy, the expert data should be close to optimal. However, in autonomous driving, selecting an optimal policy is a challenging problem, as there are no standard reward functions, which we can evaluate each data in the real world. Instead of proposing a standard reward function, recent works [13], [14] have proposed to use the non-expert data, with which we proposed to be called abnormal data, to solve this problem. The abnormal data include dangerous situations in driving. In addition, as autonomous driving is a safety-critical problem, both expert and abnormal data are needed to design the controller to prepare for abnormal driving cases. Without abnormal driving demonstrations, the autonomous vehicle may become vulnerable to abnormal and unseen situations.

In contrast to the importance of such datasets, to the best of our knowledge, the existing datasets do not cover any extreme cases, covering only the expert and safe situations. To compensate for this issue, we propose a novel dataset called the *R3 Driving Dataset*, which covers both safe driving and abnormal driving cases.

As the expert demonstrations are sampled state-action pairs from the expert driving policy, the expert demonstrations can be seen as data points sampled from an expert driving distribution. On the other hand, the abnormal demonstrations are sampled from an abnormal driving policy, a policy distinct from the expert driving policy. In terms of distribution, expert and abnormal demonstrations are in an in-distribution (ID) and out-of-distribution (OOD) relationship. Therefore, the process can be seen as anomaly detection

*Equal contribution

¹J. Oh, G. Lee, W. Oh, J. Heo, H. Chung and S. Oh are with the Department of Electrical and Computer Engineering and ASRI, Seoul National University, Seoul 08826, Korea (e-mail: jeongwoo.oh@rllab.snu.ac.kr, gunmin.lee@rllab.snu.ac.kr, wooseok.oh@rllab.snu.ac.kr, jaeseok.heo@rllab.snu.ac.kr, hojun.chung@rllab.snu.ac.kr, songhwai@snu.ac.kr).

²J. Park, and S. Choi are with the Department of Artificial Intelligence, Korea University, Seoul 02841, Korea (e-mail: baro0906@korea.ac.kr, sungjoon-choi@korea.ac.kr).

³D. Kim, B. Park, and C. Lee are with the Department of Computer Science and Engineering, Seoul National University, Seoul 08826, Korea (email: totheparadise119@gmail.com, earthquake16@gmail.com, cglee@snu.ac.kr).

This work was partly supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-01190, [SW Star Lab] Robot Learning: Efficient, Safe, and Socially-Acceptable Machine Learning, 75%) and the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-01309, Development of AI Technology for Guidance of a Mobile Robot to its Goal with Uncertain Maps in Indoor/Outdoor Environments, 25%).
Corresponding author: Songhwai Oh

when discriminating expert and abnormal demonstrations.

As conventional works do not exist, the method of evaluating such datasets is not well defined. Therefore, we also propose a combined method of uncertainty metric and OOD detection for the evaluation metric of this dataset. The results show that the epistemic uncertainty, calculated using the mixture density network (MDN) [15], best detects failing lane-keeping and unstable driving situations, which are closely related to the decision of the ego driver. On the other hand, the reconstruction loss, calculated using the variational autoencoder (VAE), best detects near-collision, dangerous overtaking, and dangerous lane-changing situations, mainly influenced by the surrounding vehicles.

II. RELATED WORK

Several approaches to provide an autonomous driving dataset [16]–[25] have been made. Geiger et al. [16] provide the researchers with the most commonly used dataset. The dataset includes data collected via a camera, LiDAR, and GPS and has annotations of bounding boxes, semantic labels, and lane marking. Caesar et al. [17] provide the same categories of data to [16], but have behavioral label annotations. Unlike [16] and [17], Geyer et al. [19] provide additional data of other vehicles. Sun et al. [18] provide the researchers with all the information in [16], [17], [19]. However, the above datasets only provide expert demonstrations, while this paper provides a new dataset including abnormal driving situations.

A number of methods [26]–[30] are applied to discriminate OOD samples from in-distribution samples. An et al. [26] proposed a method using the reconstruction loss of VAE as a metric for OOD detection. Hendrycks et al. [27] proposed an evaluation method and baselines for the OOD detection. The baselines are called the area under the receiver operating characteristic curve (AUROC) and the area under the precision-recall curve (AUPR). Liang et al. [28] proposed a method using temperature scaling and an input preprocessing method to discriminate unknown samples without modifying the existing structure of the neural network model. Lee et al. [29] proposed to modify the loss function by adding the GAN loss and confidence loss terms to the cross-entropy loss in order to discriminate an OOD sample during the training process. Islam et al. [30] proposed a method to discriminate OOD samples using uncertainty values.

Yannis et al. [31] provided the statistics about vehicle accident cases in Europe. The author categorized the traffic accident cases into 44 different categories. The vehicle accident types include pedestrians and vehicle collisions in both crossroads and straight roads, collisions between vehicles on straight roads and crossroads, and single-vehicle accidents in straight, curved, and crossroads. The proposed dataset in this paper follows the categorizations in [31].

III. DATA COLLECTION METHOD

A. Object Detection

For object detection, LiDAR sensors are used, which obtain point cloud information of surroundings. Our object detection algorithm consists of two clustering steps and one gradient descent method. The proposed detection algorithm is based on the method proposed by Yan et al. [32].

Before the clustering, the obtained point cloud is pre-processed to exclude the ground data, which is redundant for the clustering algorithm. The first clustering algorithm is the Euclidean clustering, which is used to cluster points

in the point cloud. After that, a convex hull is built for each cluster. The second Euclidean clustering is performed among the convex hulls to solve multiple clusters from the same object in the blind spot of LiDARs or when the objects are far from the LiDAR. Similar to the first Euclidean clustering, the second clustering algorithm considers convex hulls within a certain distance to be the same cluster, and the distances between any two convex hulls are calculated using the Gilbert–Johnson–Keerthi distance algorithm [33] to reduce time complexity. After grouping one cluster per object, the object parameters are optimized by the gradient descent method using the following loss function:

$$L = \alpha_w(w - w_g)^2 + \alpha_l(l - l_g)^2 + \frac{\alpha_p}{N} \sum_{i=1}^N d_{\text{box}}(p_i)^2 + \alpha_{\text{cvh}}(d_{\text{cvh}}(O_{\text{ego}}) - d_{\text{box}}(O_{\text{ego}}))^2, \quad (1)$$

where w and l are the width and length of the bounding box, w_g and l_g are the desired width and length of a vehicle, $d_{\text{cvh}}(p)$ and $d_{\text{box}}(p)$ are the distance to the convex hull and the bounding box from point p , p_i is the i th point of point cloud, which is a convex hull, N is the number of valid vertices on the convex hull, O_{ego} is the origin of the ego vehicle, and α_p , α_w , α_{cvh} and α_l are hyperparameters. We treat the point on the convex hull as valid only when it is not occluded from the ego vehicle. The overall complexity is $O(N_{\text{all}} + N_r^2)$, where N_{all} is the number of points in one frame, and N_r is the number of redundant points after removing the ground.

B. Object Tracking

With only one time frame of the point cloud, the dynamic information of vehicles cannot be estimated. However, dynamic information such as linear velocity, linear acceleration, and angular velocity of surrounding vehicles in the driving environment is essential for identifying dangerous situations. Therefore, an object tracking algorithm using the extended Kalman filter (EKF) [34] is implemented to estimate dynamic information.

The observed object list is given at each timestep to the object tracking algorithm, and the current object tracking list is updated. First, the next frame states of objects in the current object tracking list are predicted. Next, we apply the cost function proposed by Chui et al. [35] in order to decide the matching order of object pairs between the current and observed object lists.

After the cost calculation, similar to [35], objects are matched greedily until a certain threshold for the calculated cost is met. The objects are matched from the lowest to highest cost. The objects that are matched in the current object list go through the update step of EKF to estimate the next state. For the objects that are not matched in the current object list, the state information from the prediction is used to estimate the next state. If the object is not in the current object list for five consecutive frames, it is removed from the current object list. The objects that are not matched in the observed object list are thought of as new objects and inserted into the current object list.

C. Prediction Uncertainty in a Mixture Density Network

One of the most popular methods to estimate the uncertainty is the mixture density network (MDN) [15]. A MDN is composed of multiple Gaussian models, whose parameters

are outputs of a neural network. The MDN with the output y and input x can be written as:

$$p(y|x) = \sum_{j=1}^K \pi_j(x) \mathcal{N}(y; \mu_j(x), \Sigma_j(x)), \quad (2)$$

where K is the number of mixtures, $\pi_j(x)$, $\mu_j(x)$, $\Sigma_j(x)$ are the j th mixture weight function, mean function, and variance function of a Gaussian mixture model (GMM).

The total uncertainty of an MDN is the variance of the MDN network. The variance can be calculated by the sum of aleatoric and epistemic uncertainties [36]. The total uncertainty for input x can be calculated as:

$$V(y|x) = \sum_{j=1}^K \pi_j(x) \Sigma_j(x) + \sum_{j=1}^K \pi_j(x) \left\| \mu_j(x) - \sum_{k=1}^K \pi_k(x) \mu_k(x) \right\|^2, \quad (3)$$

where π_k and μ_k is the k th mixture weight function and mean function of the GMM. The other notations are same as (2).

Three types of different uncertainties are used to evaluate the dataset. The first uncertainty is the epistemic uncertainty. The epistemic uncertainty is the second term of the right-hand side of (3). This uncertainty shows the difference between the predicted value of the j th Gaussian and the other Gaussians. Therefore, this can be seen as calculating how each predicted value is different from others.

The second uncertainty is the aleatoric uncertainty, which is the first term of the right-hand side of (3). This uncertainty is the weighted sum of the model variance, which can be seen as capturing the noise in the given dataset.

The third uncertainty is the π -entropy, which denotes the entropy of mixture weights. This term is similar to epistemic uncertainty, but differs from the sense that it does not consider any disagreement among means. The equation for calculating the π -entropy is shown as:

$$\pi\text{-entropy} = - \sum_{j=1}^K \pi_j(x) \log(\pi_j(x)), \quad (4)$$

D. Using Reconstruction Loss as a Measurement of Uncertainty

A variational autoencoder (VAE) [37] is composed of an encoder model and a decoder model. An encoder model can be seen as a neural network that reduces the dimension of input data. During this dimension reduction process, the encoder extracts features that are crucial to the given data. A decoder model can be seen as a neural network that reconstructs the encoded features back to the original dimension. The encoder-decoder model is trained at once by comparing the original and reconstructed data. The model learns to reconstruct the original data with the least amount of data loss. For implementation, we used L2-loss as the loss function of VAE.

An et al. [26] proposed that during the encoding process, the encoder selects essential features, as the loss of data is inevitable during the dimension deduction process. However, when the encoder is introduced to OOD data, the encoder selects different features to extract. When these features




Fig. 2: **Sensor Configuration.** The illustration of the dimensions and the installation positions of the data collection platform.

are reconstructed, the reconstruction result is distinct from the input data, causing the reconstruction loss to elevate. Therefore, the reconstruction loss can be used as a method for anomaly detection.

E. Overview of the Collected Data

We have safely collected expert demonstrations in three different locations, a highway, urban road, and Future Mobility Technology Center (FMTC) in Siheung-si, Korea. In case of abnormal demonstrations, the data are collected only in FMTC due to the safety concerns. The abnormal demonstrations are collected based on the driving accident cases listed in [31]. The vehicle trajectories are set beforehand for safety during the data collection, and the control tower has been alert throughout the experiment. The number of expert demonstrations collected is 42,400, and the number of abnormal demonstrations collected is 20,355. The exact figures are shown in Table I and II.

IV. EXPERIMENT SETUP

A. Hardware Setting

As shown in Figure 2, various sensors are deployed on the tested vehicle (Hyundai Ioniq EV). Three LiDAR modules (Velodyne VLP-16 Hi-Res), two in the front and one in the back, are dedicated to collect surrounding environment information. These modules are used to collect the point cloud, which is exploited to gain object information around the ego vehicle. The GNSS sensor (Inertial Lab INS D) is placed on the vehicle's base link (in the middle of the trunk). This sensor receives the satellited-based location information to localize the tested vehicle. Also, it generates IMU data of the vehicle. The GNSS sensor collects the longitude and latitude, the angle from magnetic north, velocity, acceleration, and angular velocity of the ego vehicle. Lastly, the vision sensor (Mobileye ELD) is located on the vehicle's front windshield. This sensor is used to collect the lane data. Based on the above data collection platform setup, demonstrations have been collected.

B. Data Collection Method

We collect the proposed dataset in three different locations: highway, urban road, and Future Mobility Technology Center (FMTC). Also, note that the data collection rate is 10 Hz.

1) *Expert Dataset:* Expert data comprises three significant segments; a highway, city road, and test road in FMTC. The highway data, which are marked with the red line, are collected on the highway in Gwanak (Seoul), Siheung, and Incheon. The urban road data are collected around the FMTC in Siheung, marked with the blue line. Finally, the test road data are collected on the FMTC test site, in which the experiments are artificially designed. The photo of the

	location	# of exp	# of frames	Total
expert	FMTC	5	18,400	42,400
	urban	4	12,000	
	highway	4	12,000	

TABLE I: **Expert Demonstrations.** The number of episodes and frames collected in each location. # of exp indicates the number of experiments done in each location. # of frames indicates the number of frames collected in each location. Total indicates the total number of frames collected.



Fig. 3: **FMTC Snapshot.** The snapshot shows the autonomous driving test site in FMTC, Siheung. The track is divided into straight road, crossroad, roundabout, toll gate, overpass, and parking lot. In this paper, only the straight road, crossroad, and roundabout are used.

FMTC track is shown in Figure 3. The number of frames collected in each category is shown in Table I.

2) *Abnormal Dataset:* As a safety system is needed during the abnormal data collection, abnormal data are collected in the FMTC. In addition, we would like to emphasize the fact that no one is injured while collecting the abnormal driving demonstrations. The abnormal driving demonstrations are based on common accident data set (CADaS) [31], which provides common accident situations in Europe. To secure safety during the data collection, only less risky situations are executed.

A total of 20,455 frames of abnormal data in 369 experiments are collected during the data collection. The abnormal data are categorized in two different ways: the road type and the hazard type. The road type category includes straight road and crossroad. The hazard type category includes unstable driving, failing lane keeping, dangerous lane changing, dangerous overtaking, and near collision. Note that the hazard type categorization is not entirely distinguishable from each other since the driving situations cannot be strictly divided into separate categories.

3) *Data Structure:* The collected dataset includes the essential data for the autonomous driving dataset. One ego vehicle timestep data file includes the latitude (x), longitude (y), the velocity (v) and the angular velocity (ω) of the ego vehicle, the angle from the magnetic north and the ego vehicle (θ), the heading (a_x) and lateral (a_y) direction acceleration of the ego vehicle, the lane deviation, and the discretized decision of the ego vehicle.

The state of each surrounding object includes the position of the object in Cartesian coordinate (x , y , θ), the velocity (v), and the angular velocity (ω) calculated in the ego vehicle frame. The data also includes the acceleration magnitude of the object heading direction (a_x), the length and width of the object (l , w), a unique id for the object (id), and the number of frames that appears near the ego vehicle (age). v , a_x , ω , and age are extracted using object tracking, and x , y , θ , l , and w are extracted using object detection.

Finally, the bin file includes the point cloud data of each frame. For detailed information, note the GitHub repository in the following link: <https://github.com/rllab-snu/R3-Driving-Dataset>.

C. Anomaly Detection

To provide a decent baseline for discriminating expert and abnormal demonstrations, we evaluated the proposed collected dataset by using anomaly detection algorithms. As explained in Section I, the expert dataset and the abnormal dataset can be seen as an ID and OOD relationship. Therefore, using anomaly detection in discrimination of expert and abnormal demonstrations is valid.

As the state dimension varies due to the varying number of surrounding vehicles, the maximum number of the other vehicle is set to five.

Two different experiments are executed to show the performance difference between single-frame state-action pairs and multi-frame sequential state-action pairs for training data. The input data for single-frame state-action pairs for both MDN and VAE are the state data. The output data of MDN are the action data. In the case of using the multi-frame sequential state-action pairs (five-frame pairs are used in experiments), the input data for both MDN and VAE is the five state data and the first four action data. The output data for MDN are the last (fifth) action data.

V. EXPERIMENT RESULTS

The comparative visual results of epistemic uncertainty, aleatoric uncertainty, π -entropy, and reconstruction loss using one frame and five consecutive frames are shown in Figure 4. The numerical results of the experiments are shown in Table III. For the evaluation methods, the AUROC and the AUPR score are used.

From the results shown in Table III, it can be inferred that the aleatoric uncertainty cannot discriminate the abnormal situations. The AUROC and the AUPR scores of aleatoric uncertainty are less than 0.65, which illustrates that aleatoric uncertainty is not a suitable anomaly detection method. In addition, as the number of frames increases, all the scores drop below 0.5, which is worse than random selection.

In the case of the epistemic uncertainty, it shows the best results among the algorithms at the discriminating failing lane-keeping cases collected in both crossroad and straight road when using a single frame, as shown in Table III. The discrimination ability is reinforced when using five consecutive frames, excelling at the unstable driving cases and failing lane-keeping cases, which are the best results among the entire experiments.

In the case of the reconstruction loss, it shows the best results among the algorithms at the discriminating near-collision cases collected in both crossroad and straight road, and the dangerous lane changing and overtaking cases in a straight road when using a single frame, as shown in Table III. These results are the best results among the entire experiments in the above four cases. Unlike epistemic uncertainty, the results of the reconstruction loss deteriorate when the number of frames increases to five.

An interesting part of the result is that the reconstruction loss and epistemic uncertainty excel in different cases. In the case of MDN, the network excels at discriminating unstable driving and failing lane-keeping cases. When the epistemic uncertainty is high, the action of the ego vehicle has high variance. Hence, it can be inferred that the ego vehicle is under unstable driving or failing lane-keeping. On the other hand, the reconstruction error excels at discriminating near-collision, dangerous lane changing, and overtaking cases. The reconstruction error increases when the surroundings

	road type	hazard type	# of exp	# of frames	sub total	Total
abnormal	straight road	unstable driving	45	6,360	13,300	20,355
		failing lane keeping	49	6,650		
		dangerous lane changing	50	1,785		
		dangerous overtaking	30	925		
		near collision	96	5,620		
	crossroad	unstable driving	28	1,380	7,055	
		failing lane keeping	8	530		
		near collision	180	5,145		

TABLE II: **Abnormal Demonstrations.** The number of episodes and frames collected in each road and hazard type. # of exp indicates the number of experiments done about each hazard type. # of frames indicates the number of frames collected in each category. Subtotal and total each indicate the total number of frames collected in each road type and the total number of collected frames.




Fig. 4: **Anomaly Detection Results.** The results of using the uncertainties and reconstruction loss using a single frame and five consecutive frames for anomaly detection. Epis, alea, pi_entropy, and recon each indicate the epistemic uncertainty, the aleatoric uncertainty, the π -entropy, and the reconstruction loss. The lane_keeping indicates the failing lane-keeping case. The overtaking indicates the dangerous overtaking case. The lane_changing indicates the dangerous lane changing case.

	unstable driving (straight)	failing lane keeping (straight)	near-collision (straight)	unstable driving (cross)	failing lane keeping (straight)	dangerous lane changing (straight)	dangerous overtaking (straight)	near-collision (cross)
Epistemic	0.769/0.686 0.843/0.818	0.824/0.758 0.801/0.788	0.533/0.406 0.647/0.456	0.799/0.270 0.901/0.681	0.842/0.166 0.860/0.340	0.649/0.234 0.746/0.283	0.541/0.099 0.783/0.194	0.607/0.419 0.788/0.646
Aleatoric	0.624/0.476 0.402/0.414	0.577/0.446 0.369/0.374	0.401/0.323 0.536/0.416	0.610/0.152 0.294/0.094	0.602/0.065 0.205/0.032	0.379/0.135 0.400/0.124	0.410/0.085 0.422/0.078	0.376/0.291 0.448/0.338
π -entropy	0.804/0.666 0.329/0.334	0.820/0.692 0.357/0.355	0.655/0.499 0.593/0.433	0.857/0.455 0.265/0.088	0.807/0.156 0.228/0.033	0.699/0.274 0.530/0.159	0.685/0.147 0.516/0.096	0.642/0.532 0.370/0.276
Recon loss	0.363/0.332 0.548/0.399	0.402/0.357 0.536/0.402	0.778/0.554 0.718/0.467	0.227/0.084 0.593/0.140	0.243/0.035 0.492/0.049	0.827/0.349 0.774/0.255	0.850/0.248 0.804/0.169	0.835/0.622 0.779/0.479

TABLE III: **Anomaly Detection Results.** The numerical result of each experiment. The first line of each row indicates the AUROC (left) and AUPR (right) scores calculated in the experiments using a single frame. The second line of each row indicates the AUROC (left) and AUPR (right) scores calculated in the experiments using five consecutive frames. The bold figure indicates the best AUROC and AUPR scores for each situation.

of the ego vehicle are different from the data used to train the network. Therefore, it can be inferred that near-collision, dangerous lane changing, and overtaking are mainly affected by the surrounding of the ego vehicle. In conclusion, the epistemic uncertainty is useful at detecting abnormal behaviors of the ego vehicles, while the reconstruction loss is useful for detecting abnormal surroundings.

When using five consecutive frames, the performance of the epistemic uncertainty increases, as the information from the past is also given to the network. This increased information allows the epistemic uncertainty to consider the sequential results of the action during prediction. This allows the epistemic uncertainty to better discriminate abnormal cases. On the other hand, the performance of the recon-

struction loss decreases when using five consecutive frames. As the VAE input has changed into consecutive frames, the results of some of the data that has been viewed abnormal are changed. This change of view leads the reconstruction error to lessen, causing deterioration in AUROC and AUPR scores.

From the experimental results, it is desirable to use the epistemic uncertainty with multiple consecutive frames and the reconstruction loss with a single current frame for anomaly detection, as they cover all the abnormal cases. From the results, it can be inferred that unstable driving and failing lane-keeping cases can be regarded as problems with the ego-driver, as the epistemic uncertainty better captures these cases. On the other hand, near-collision, dangerous

overtaking, and lane changing cases can be regarded as problems related to surrounding drivers, as the reconstruction loss better captures these cases.

VI. CONCLUSION

This paper presents a novel dataset called the *R3 Driving Dataset*, which contains normal and abnormal driving demonstrations. The number of normal and abnormal demonstrations is 42,400 and 20,355, respectively. The data collection platform is also proposed, along with baseline methods for anomaly detection of the proposed dataset. The MDN and VAE are used for anomaly detection, and the evaluation metrics of the algorithm are the AUROC and AUPR. From the results, we can infer that near-collision, dangerous lane changing, and overtaking cases can be detected effectively using the reconstruction loss, while unstable driving and failing lane-keeping cases can be detected more effectively using the epistemic uncertainty.

REFERENCES

- [1] W. Xiao, N. Mehdipour, A. Collin, A. Y. Bin-Nun, E. Frazzoli, R. D. Tebbens, and C. Belta, "Rule-based optimal control for autonomous driving," in *Proc. of the ACM/IEEE 12th International Conference on Cyber-Physical Systems*, 2021.
- [2] M. Parseh, F. Asplund, M. Nybacka, L. Svensson, and M. Törngren, "Pre-crash vehicle control and manoeuvre planning: A step towards minimizing collision severity for highly automated vehicles," in *Proc. of the IEEE International Conference of Vehicular Electronics and Safety, Cairo, Egypt*, September 4-6, 2019.
- [3] L. Claussmann, M. Revilloud, D. Gruyer, and S. Glaser, "A review of motion planning for highway autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 5, pp. 1826–1848, 2020.
- [4] M. Campbell, M. Egerstedt, J. P. How, and R. M. Murray, "Autonomous driving in urban environments: Approaches, lessons and challenges," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4649–4672, 2010.
- [5] L. Claussmann, M. Revilloud, S. Glaser, and D. Gruyer, "A study on ai-based approaches for high-level decision making in highway autonomous driving," in *Proc. of the IEEE International Conference on Systems, Man, and Cybernetics, Banff, AB, Canada*, October 5-8, 2017.
- [6] A. E. Sallab, M. Abdou, E. Perot, and S. K. Yogamani, "Deep reinforcement learning framework for autonomous driving," *CoRR*, vol. abs/1704.02532, 2017.
- [7] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani, and P. Pérez, "Deep reinforcement learning for autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–18, 2021.
- [8] J. Chen, B. Yuan, and M. Tomizuka, "Model-free deep reinforcement learning for urban autonomous driving," in *Proc. of the IEEE Intelligent Transportation Systems Conference, Auckland, New Zealand*, October, 2019.
- [9] Y. Pan, C. Cheng, K. Saigol, K. Lee, X. Yan, E. A. Theodorou, and B. Boots, "Imitation learning for agile autonomous driving," *International Journal of Robotics Research*, vol. 39, no. 2-3, 2020.
- [10] J. Chen, B. Yuan, and M. Tomizuka, "Deep imitation learning for autonomous driving in generic urban scenarios with enhanced safety," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Macau, SAR, China*, November 3-8, 2019.
- [11] F. Codevilla, M. Müller, A. M. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *Proc. of the IEEE International Conference on Robotics and Automation, Brisbane, Australia*, May 21-25, 2018.
- [12] Z. Xu, Y. Sun, and M. Liu, "iCurb: Imitation learning-based detection of road curbs using aerial images for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1097–1104, 2021.
- [13] G. Lee, D. Kim, W. Oh, K. Lee, and S. Oh, "Mixgail: Autonomous driving using demonstrations with mixed qualities," in *Proc. of the IEEE International Conference on Intelligent Robots and Systems, Las Vegas, NV, USA*, October 24, 2020 - January 24, 2021.
- [14] S. Choi, K. Lee, and S. Oh, "Robust learning from demonstrations with mixed qualities using leveraged Gaussian processes," *IEEE Transactions on Robotics*, vol. 35, no. 3, pp. 564–576, 2019.
- [15] C. Bishop, "Mixture density networks," Aston University, Birmingham, U.K., Tech. Rep., February, 1994.
- [16] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [17] 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 *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA*, June 13-19, 2020.
- [18] P. Sun, H. Kretschmar, 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 *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA*, June 13-19, 2020.
- [19] J. Geyer, Y. Kassahun, M. Mahmudi, X. Ricou, R. Durgesh, A. S. Chung, L. Hauswald, V. H. Pham, M. Mühlegg, S. Dorn, T. Fernandez, M. Jänicke, S. Mirashi, C. Savani, M. Sturm, O. Vorobiov, M. Oelker, S. Garreis, and P. Schuberth, "A2D2: Audi autonomous driving dataset," *CoRR*, vol. abs/2004.06320, 2020.
- [20] P. Dollár, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: A benchmark," in *Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Miami, Florida, USA*, June 20-25, 2009.
- [21] G. J. Brostow, J. Fauqueur, and R. Cipolla, "Semantic object classes in video: A high-definition ground truth database," *Pattern Recognition Letters*, vol. 30, no. 2, pp. 88–97, 2009.
- [22] A. Buyval, A. Gabdullin, R. Mustafin, and I. Shimchik, "Realtime vehicle and pedestrian tracking for didi udacity self-driving car challenge," in *Proc. of the IEEE International Conference on Robotics and Automation, 2018, Brisbane, Australia*, May 21-25, 2018.
- [23] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA*, June 27-30, 2016.
- [24] X. Huang, P. Wang, X. Cheng, D. Zhou, Q. Geng, and R. Yang, "The apolloscope open dataset for autonomous driving and its application," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 10, pp. 2702–2719, 2020.
- [25] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 Year, 1000km: The Oxford RobotCar Dataset," *The International Journal of Robotics Research (IJRR)*, vol. 36, no. 1, pp. 3–15, 2017.
- [26] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," *Special Lecture on IE*, vol. 2, no. 1, pp. 1–18, 2015.
- [27] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," in *Proc. of the 5th International Conference on Learning Representations, Toulon, France*, April 24-26, 2017.
- [28] S. Liang, Y. Li, and R. Srikant, "Enhancing the reliability of out-of-distribution image detection in neural networks," in *Proc. of the 6th International Conference on Learning Representations, Vancouver, BC, Canada*, April 30 - May 3, 2018.
- [29] K. Lee, H. Lee, K. Lee, and J. Shin, "Training confidence-calibrated classifiers for detecting out-of-distribution samples," in *Proc. of the 6th International Conference on Learning Representations, Vancouver, BC, Canada*, April 30 - May 3, 2018.
- [30] R. U. Islam, M. S. Hossain, and K. Andersson, "A novel anomaly detection algorithm for sensor data under uncertainty," *Soft Computing*, vol. 22, no. 5, pp. 1623–1639, 2018.
- [31] G. Yannis, P. Evgenikos, and A. Chaziris, "Cadas-A common road accident data framework in Europe," in *Proc. of the IRTAD conference, Seoul, Korea*, 2009.
- [32] Z. Yan, T. Duckett, and N. Bellotto, "Online learning for 3D LiDAR-based human detection: Experimental analysis of point cloud clustering and classification methods," *Autonomous Robots*, vol. 44, no. 2, pp. 147–164, 2019.
- [33] E. G. Gilbert, D. W. Johnson, and S. S. Keerthi, "A fast procedure for computing the distance between complex objects in three-dimensional space," *IEEE Journal on Robotics and Automation*, vol. 4, no. 2, pp. 193–203, 1988.
- [34] G. C. Charles K. Chui, *Kalman Filtering with Real-Time Applications*, 1st ed. Berlin, Heidelberg: Springer-Verlag, 1987.
- [35] H. Chiu, A. Prioletti, J. Li, and J. Bohg, "Probabilistic 3D multi-object tracking for autonomous driving," *CoRR*, vol. abs/2001.05673, 2020.
- [36] S. Choi, K. Lee, S. Lim, and S. Oh, "Uncertainty-aware learning from demonstration using mixture density networks with sampling-free variance modeling," in *Proc. of the IEEE International Conference on Robotics and Automation, Brisbane, Australia*, May 21-25, 2018.
- [37] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. of the 2nd International Conference on Learning Representations, Banff, AB, Canada*, April 14-16, 2014.