2023 IEEE International Conference on Robotics and Automation (ICRA 2023)
May 29 - June 2, 2023. London, UK

SmartRainNet: Uncertainty Estimation For Laser Measurement in Rain

Chen Zhang[1], Zefan Huang[1], Beatrix Xue Lin Tung[1], Marcelo H. Ang Jr.[1] and Daniela Rus[2]

Abstract— **Adverse weather has raised a big challenge for autonomous vehicles. Unreliable measurements due to sensor degradation could seriously affect the performance of autonomous driving tasks, such as perception and localization. In this work, we study sensor degradation in rainy weather and present a novel method that evaluates the uncertainty for each laser measurement from a 3D LiDAR. With uncertainty estimation, downstream tasks that rely on LiDAR input (e.g., perception or localization) can increase their reliability by adjusting their reliance on laser measurements with varying fidelity. Alternatively, uncertainty estimation can be used for sensor performance evaluation. Our proposed method, SmartRainNet, uses an attention-based Mixture Density Network to model the dependence between neighboring laser measurements and then calculate the probability density for each laser measurement as an uncertainty score. We evaluate SmartRainNet on synthetic and naturalistic sensor degradation datasets and provide qualitative and quantitative results to demonstrate the effectiveness of our method in evaluating uncertainty. Finally, we demonstrate three practical applications of uncertainty estimation to address autonomous driving challenges in rainy weather.**

I. Introduction

Autonomous driving in adverse weather has been widely recognized as a challenging problem. One major concern is sensor degradation which causes sensor measurements to be inaccurate and unreliable. The uncertainties raised in sensor measurements could subsequently deteriorate the performance of downstream tasks such as perception or localization when they directly consume those corrupted measurements. Most state-of-art algorithms have shown promising results when trained and tested with clean data. But these algorithms are likely to fail in real-world environments when adverse weather substantially deteriorates the quality of input data.

Among the most popular techniques to address this problem is data restoration. Two approaches are generally used: recover corrupted data or remove corrupted data. Despite their popularity, these two approaches have limitations. First, recovery-based methods (e.g., generative models [1], [2]) may generate artifacts that do not originally belong to the scene. Second, removal-based methods (e.g., [3], [4]) usually require modeling the noise generation processes. But in practice, the complex interactions between the sensor, surrounding environments, and adverse weather make it almost infeasible to build such a model correctly.

[1] National University of Singapore, Singapore zcreation.singapore@gmail.com; huangzefan@u.nus.edu; e0052785@u.nus.edu; mpeangh@nus.edu.sg
[2]Massachusetts Institute of Technology, Cambridge, MA, USA rus@csail.mit.edu

Another popular approach to addressing sensor degradation is to implicitly consider the degradation in the model. For example, most perception algorithms perform training and test both on data collected from adverse weather and then directly aim to improve evaluation metric scores (e.g., average precision for object detection in rainy weather [5]). Nevertheless, this approach has a few limitations. First, this level of task (compared to sensor data restoration) generally requires a large amount of labeled data to obtain satisfying results. However, collecting data from adverse weather is more difficult than from clear weather because adverse weather events are usually hard to predict and capture. Second, this approach generally cannot output degradation information of their performance. Therefore, it is hard for autonomous driving systems to know the current driving capability in adverse weather, let alone adjust driving strategies based on varying weather conditions.

Fig. 1 shows laser measurement degradation under different rainy conditions. In this work, we propose a novel method to deal with sensor degradation: estimating the uncertainty for each laser measurement. The uncertainty estimation results can be input into downstream tasks as auxiliary information to represent the confidence of each measurement. Further, the uncertainty information can be used for sensor performance evaluation. Specifically, our proposed method, SmartRainNet, learns the spatial correlation between laser measurements through an attention mechanism, then models the probability distribution of each laser measurement through a Mixture Density Network. Finally, the density evaluated at each measurement is defined as the uncertainty score associated with each laser measurement.

Our SmartRainNet has been evaluated on a naturalistic degradation dataset and a synthetic degradation dataset. Experimental results show that our model can correctly find laser measurements with high uncertainties. Our method explores a new but under-explored research field: sensor degradation evaluation. In the last section, we show three relevant applications in rainy weather.

II. Related Works

Research about sensor degradation can be traced to laser sensing studies, which mainly concentrate on understanding laser transmission behavior through atmospheric gases and particles. For example, in [6], an energy loss equation was developed for laser beams propagating through the atmosphere. This work established the theoretical foundation for a series of subsequent studies, such as predicting rain influence [7] or simulating rain effects on LiDAR sensors [8], [9]. More recently, sensor degradation has been raised

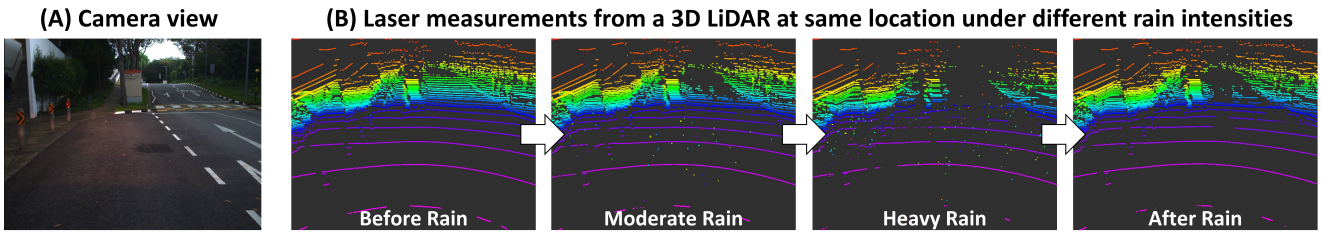

Fig. 1. **Laser measurements degradation.** (A) Surrounding environment. (B) Laser measurement changes from the start to the end of the rain. As shown in (B), laser measurement degrades severely during the rain. For example, some laser measurements on the road are absent, and others become noise. These noisy and incomplete measurements, blended with normal measurements, could negatively affect downstream tasks (e.g., perception and localization). To address this problem, we introduce a novel task that estimates the uncertainty of each laser measurement.

as a concern for autonomous driving. For example, authors [10], [11] conducted a series of qualitative analyses to discuss weather's adverse effects on autonomous systems. Quantitative analyses are also available, but most of them are conducted in controlled and simulated environments using fixed targets for basic statistical analysis, such as laser measurement changes in intensity, range, or the number of measurements returned [9], [12], [13].

The first work to quantify LiDAR degradation in real-world driving environments is proposed in [14], which formulates degradation evaluation as an anomaly detection problem and estimates the degradation for each LiDAR scan. In their work, an anomaly detection model is applied to learn a hypersphere in the feature space that encompasses the majority of normal LiDAR scans (collected in clear weather) while excluding anomalous LiDAR scans (collected in rainy weather). The distance from a test sample to this hypersphere center is then defined as the degradation score. This work studies degradation in each laser measurement. The essential idea is as follows; instead of formulating an anomaly detection problem, which requires selecting anomaly samples based on subjective criteria, our method constructs a probability distribution of normal laser measurements (i.e., collected in clear weather) and then uses the probability density function to distinguish anomaly measurements. More design details are given in the next section.

III. Methodology

This work aims to find unreliable laser measurements in the environment. The fundamental idea is to build a probability distribution that captures the behavior of reliable measurements in clear weather and use that distribution to discover unreliable measurements, such as those affected by rain degradation or unexpected measurements. A typical property of laser measurement is its multimodal property. LiDAR emits light pulses to measure the distance to the object. During the propagation, the laser light can encounter multiple obstacles along the path, and each part of reflected light could be potentially picked up as a valid measurement. So in this work, we build a conditional mixture model that considers the spatial correlation between laser measurements and the multimodal property of laser measurements. The probability density is then defined as the uncertainty for each laser measurement. The pipeline of our approach is displayed in Fig. 2.

A. Preprocessing

To use convolution neural networks (CNNs) to process point clouds, a common practice is to convert those LiDAR point clouds into images. However, our work adopts this practice mainly for another purpose: to record both valid measurements (observed measurements) and absent measurements. A typical 3D rotating LiDAR produces a fixed number of laser measurements during each scan. For example, in our experiment setup, we use a 32-channel 3D LiDAR with 0.2° horizontal resolution. This 3D LiDAR fires 360/0.2 times horizontally, and each firing produces 32 laser measurements vertically. However, some laser measurements may not be valid if the emitted laser pulses cannot return or if returned signals are too weak to be recognized by the LiDAR. In rainy weather, this phenomenon becomes more prevalent as the interactions between rain and laser pulses could severely distort laser signals to cause more losses of laser measurements. In some circumstances, it is possible that laser measurements on an obstacle front are completely absent. Therefore, to reduce such risks, knowing which laser measurements do not have valid returns is necessary. An image is a suitable representation to represent both valid and absent laser measurements. For each LiDAR scan, we create an image with a height equal to the number of LiDAR channels and a width equal to the number of horizontal firings. Each pixel location (i, j) is related to a laser measurement $\boldsymbol{s}_{i,j}$ from channel i and horizontal firing j, and each measurement contains five values: $\{d, r, x, y, z\}$ where d is the distance to the object, r is the intensity of the reflected beam from the object, and x, y, z are 3D Cartesian coordinates of the object. If the laser measurement is absent, all values are set to zeros.

B. Spatial Correlation Learning

An attention mechanism is then used to capture the correlation between neighboring laser measurements that share similar characteristics. Attention mechanisms were firstly introduced in [15] and [16] in conjunction with a recurrent neural network (RNN)-based architecture to create attention between tokens in the source sentence and target sentence. The attention mechanism later adopted in [17] is

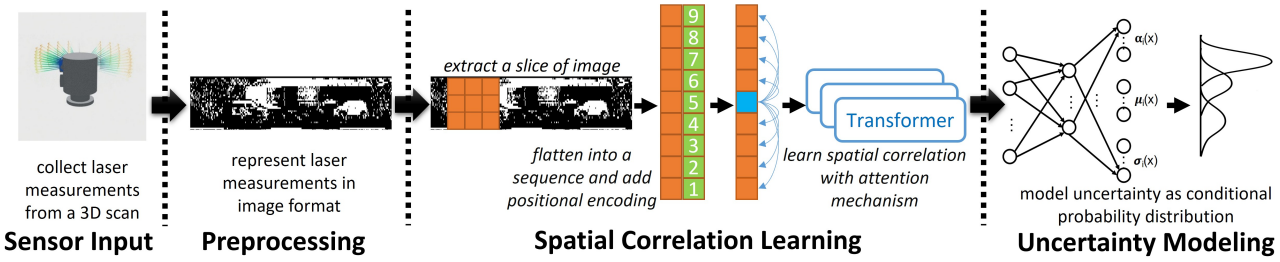

Fig. 2. **SmartRainNet Architecture.** First, 3D LiDAR scans are converted into 2D LiDAR images. Then the attention mechanism is applied to learn the spatial correlation between laser measurements extracted from a slice of the image. After that, a Mixture Density Network is applied to transform this spatial correlation into a conditional probabilities distribution. Finally, the conditional probability density is defined as the uncertainty associated with each laser measurement.

completely free from complex recurrent or convolutional neural networks. The new architecture, called Transformer, revolutionized the use of attention with great success not only in the fields of natural language processing [18], [19], [20] but also in computer vision [21], [22], [23], [24]. Basically, the Transformer architecture captures the correlation by mapping between a query and a set of key-value pairs that embed the similarity or compatibility information with respect to the query, similar to convolution operations in CNNs. Although CNN and Transformer are theoretically equivalent [25], only the Transformer is appropriate for the task discussed in this work. CNNs can only calculate correlation for pixels centered at the kernel. When multi-layer CNN kernels with a large receptive field are applied, pixels near the image edge will not be considered (unless padding was used at the cost of introducing many uninformative pixels). On the contrary, Transformer can easily attend to every pixel to the neighboring pixels using the query mechanism. In this work, we calculate the attention value for each laser measurement $\boldsymbol{s}_{i,j}$ considering a set of neighboring locations $\mathcal{P}(i,j)$ as:

$$\sum_{a,b\in\mathcal{P}(i,j)} \text{softmax}_{a,b}(\boldsymbol{q}_{i,j}^{\top}\boldsymbol{k}_{a,b} + \boldsymbol{q}_{i,j}^{\top}\boldsymbol{r}_{a-i,b-j})\boldsymbol{v}_{a,b} \quad (1)$$

where queries $\boldsymbol{q}_{i,j} = W_Q\boldsymbol{s}_{i,j}$, keys $\boldsymbol{k}_{a,b} = W_K\boldsymbol{s}_{a,b}$, values $\boldsymbol{v}_{a,b} = W_V\boldsymbol{s}_{a,b}$ and W_Q, W_K, W_V are weight matrices to be learned. $\boldsymbol{r}_{a-i,b-j}$ is the relative positional encoding that embeds spatial information between the calculating location (i,j) and attending location (a,b).

C. Uncertainty Modeling

This step integrates the spatial correlation into a probabilistic framework. First, let us consider the characteristics of laser measurements in clear weather. A laser emits a light pulse and measures the time it takes to reflect from an object. If the laser is partially reflected by multiple objects along its propagation path, it is possible to receive multiple return signals. Fig. 3 displays a real-world example. In this example, we display the histograms of accumulated laser measurements from four image regions: the edge of signage, the surface of signage, the ground plane, and the bushes. It can be seen that Histogram 2 and 3 in Fig. 3 B only have a single peak but Histogram 1 and 4 have two peaks. Such a difference is caused by the complex interactions between the lasers and the environment. A laser emits a light beam whose diameter will diverge with increasing distance. When the laser pulse strikes a reflective object, it creates a laser spot of a certain radius. Depending on the environment, the measurement returned may be different: if the laser is incident on the wall or floor larger than the spot, it will most likely produce a single consistent measurement because most laser energy will be reflected (e.g., at Location 2 and 3). But if the laser spot is incident on an edge, the laser pulse will be partially reflected from the first measuring object, and its remaining part will continue the propagation. As a result, the laser receiver may receive multiple echoes, and each echo potentially becomes a valid measurement. Measurements from Location 1 and 4 are examples of this.

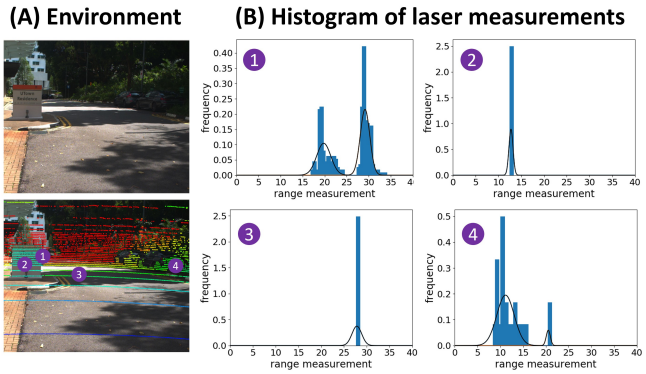

Fig. 3. **Multi-modal laser measurements in clear weather.** We project LiDAR scans onto the camera image and accumulate 20 seconds of laser measurements from four locations: the edge of signage, the front surface of signage, the ground, and the bushes. The histogram of laser measurements at each location is plotted in (B). Laser measurements can be multi-valued due to edge effects (e.g., Location 1) or multi-path effects (e.g., Location 4). This justifies using a mixture model to model the multi-mode data.

To deal with multi-modal measurements, a natural choice is to choose a mixture model. Here we use a mixture of Gaussian distributions in this work. Also, each laser measurement is spatially related to its nearby measurements. Therefore, we build a conditional probability density function of laser measurement $\boldsymbol{s}_{i,j}$, conditioned on neighboring measurements $P_{i,j} = \{\boldsymbol{s}_{a,b}|a,b \in \mathcal{P}_{i,j}\}$. This all together gives a Mixture Density Network [26]:

$$p(s_{i,j}|P_{i,j}) = \sum_{k=1}^{K} \pi_k \mathcal{N}(s_{i,j}|\mu_k, \sigma_k) \quad (2)$$

where K is the number of Gaussian components, π_k is mixing coefficient for k_{th} Gaussian component that satisfy $0 \leqslant \pi_k \leqslant 1$ and $\sum_{k=1}^{K} \pi_k = 1$. μ_k and σ_k represent the center of k_{th} and standard deviation of the k_{th} Gaussian kernel.

The conditional density function represents the relative likelihood of a laser measurement to be expected from inferring its neighboring measurements. Therefore, a high density means it is more likely to observe that laser measurement, thus containing low uncertainty. In contrast, a low density means that laser measurement is very likely to be spurious or affected by degradation, such as noise or absent measurement due to random interactions between raindrops and lasers. These spurious, unexpected measurements are considered to be of high uncertainty.

IV. Experiments

In this section, we evaluate SmartRainNet on a naturalistic driving dataset and a synthetic degradation dataset and present qualitative and quantitative evaluation results.

A. Dataset and Task Descriptions

Naturalistic Driving Dataset (NDD) This dataset consists of LiDAR scans recorded from a 32-channel RoboSense 3D LiDAR mounted on a full-sized data collection vehicle. This dataset is made of two parts. The first part is collected in clear weather while manually driving the vehicle along the data collection route. The second part is collected in rainy weather stationarily throughout the entire rainfall event (from the beginning to the end of the rain). Besides LiDAR scans, we also recorded the real-time local rainfall intensity using a rainfall measurement sensor (Lufft WS100).

Synthetic Degradation Dataset (SDD) This dataset is created by adding simulated degradation to laser measurements collected from clear weather. Two types of degradation are added: absent measurements to simulate rain absorption and noise measurements to simulate laser measurements reflected from raindrops. To generate them, we randomly select a proportion of normal measurements from the first part of the Naturalistic Driving Dataset (collected in clear weather) and replace them with simulated degraded measurements.

B. Experimental Setup

Our 3D LiDAR has 32 channels, a horizontal resolution of 0.2°, and a maximum range of 60 m. Each LiDAR scan is converted to a LiDAR image with the height of 32 and the width of (360/0.2). In the experiment, all valid range measurements are scaled to $(0, 1]$, and absent measurements are set to 1. The Transformer layer consists of two identical self-attention blocks, and each block has the same input and output dimension of 5. The linear layer inside the block has a dimension of 20. The last layer of the second block is connected to a linear layer with an output dimension of 4, corresponding to the means and variances for two Gaussian components. We trained the model using LiDAR scans collected in clear weather with Negative Log-Likelihood loss and Adam optimization [27].

C. Qualitative Results

1) Uncertainty Estimation in Clear and Rainy Weather: This section analyzes uncertainty estimation in four scenarios: clear weather, noise measurements due to raindrop reflection, absent measurements due to rainwater absorption, and absent measurements on dynamic objects. We select four LiDAR scans in each scenario and estimate the uncertainty for each laser measurement. Fig. 4 presents the results in each scenario with LiDAR range images on the left and uncertainty estimation on the right.

Fig. 4 B1 shows that the uncertainty estimation is stable in clear weather. All four scans have similar uncertainty distribution: high uncertainty scores are assigned to areas that are most likely to generate multi-echoes, such as the tree crowns highlighted in yellow boxes. Fig. 4 B2 highlights uncertainty estimation for noise measurements reflected from raindrops. Those sporadic measurements are caused by random raindrop reflections and therefore are assigned high uncertainty scores. Fig. 4 B3 highlights uncertainty estimation for absent measurements (black pixels) on the road in rainy weather. As the rainfall increases from scan S1 to scan S4, many laser measurements on the road are absent due to absorption by accumulated water on the road. Correspondingly, those areas are assigned with high uncertainty scores. Fig. 4 B4 shows another type of absent measurement on a car: laser measurements travel through transparent windows that refrain back reflections or on dark paints that absorb the majority of laser light. In the range images, it can be seen that most measurements on the car cannot return (black pixels). As a result, this car is almost invisible in the point cloud as point clouds can only save those measurements returned to the LiDAR. But in our approach, these absent measurements on the car are designated with high uncertainty scores that distinguish the car from the background scenes.

2) Uncertainty Estimation Under Different Rain Intensities: Laser degradation could also change with rainfall intensities. For example, heavy rainfall creates large amounts of tiny raindrops in the air, increasing the chance of laser beams interacting with raindrops to produce distorted measurements. Also, rainwater remaining on an objects' surface can reduce their reflectance and reduce valid measurements from these objects.

To test whether these changes can be reflected in uncertainty estimation, we collect a sequence of LiDAR scans at two locations stationarily (our data collection vehicle is not moving) and select one scan each before the rain, during moderate rain, during heavy rain, and after the rain. Fig. 5 shows the histogram of uncertainty scores for laser measurements in these scans. When the rainfall changes from moderate to heavy, the number of laser measurements with higher uncertainty scores increases. When the rainfall stops, the number of laser measurements with higher uncertainty scores decreases. This shows that the uncertainty estimation

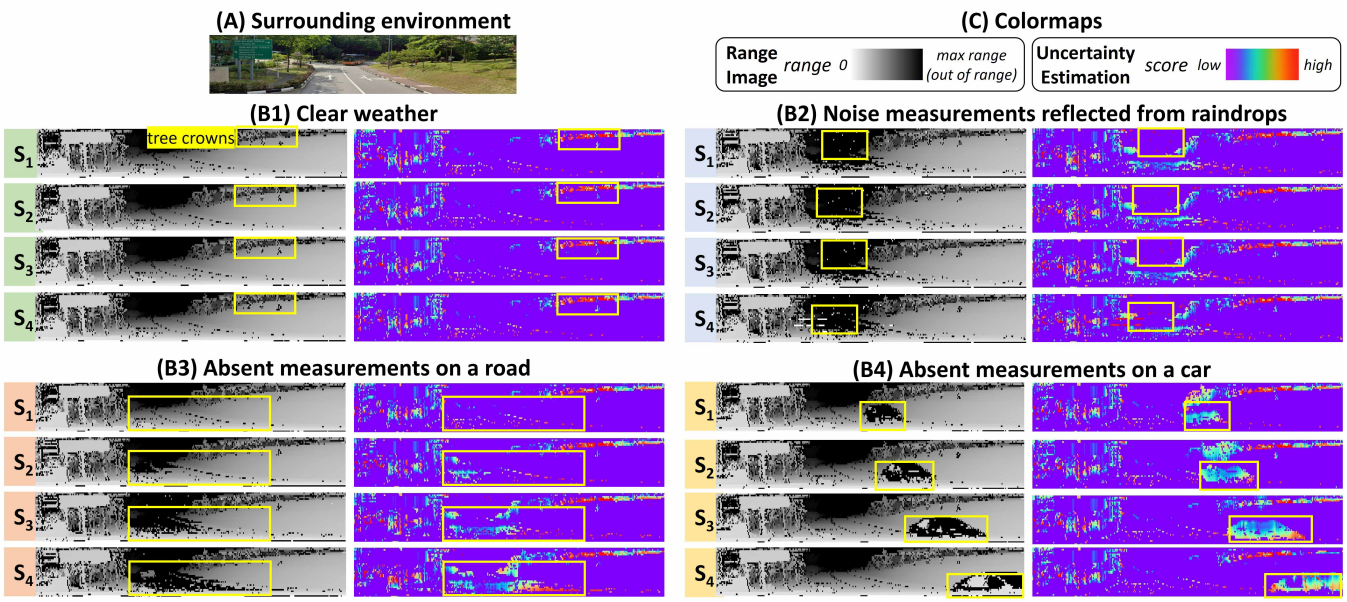

Fig. 4. **Visualization of uncertainty estimation at location (A) in four scenarios (B1-B4).** In each scenario, we select four LiDAR scans (S1-S4), display LiDAR range images on the left and uncertainty estimation results on the right, and highlight the areas discussed in yellow boxes. The colormaps for the LIDAR range image and uncertainty estimation are displayed in (C). In (B1), uncertainty estimation is consistent in clear weather: most areas have low uncertainty scores, and high uncertainty scores are assigned to areas that are unlikely to have consistent measurements, such as tree crowns. In (B2), some noise measurements are generated due to the reflection from raindrops. Those sporadic measurements are assigned with high certainty score that easily distinguishes them from the background. In (B3), many laser measurements on the road are absent as the rainfall intensity increases from S1 to S4. Those absent measurements are generated mostly due to rainwater absorption and are therefore assigned high uncertainty scores. (B4) shows an important application of uncertainty estimation. From the LiDAR range images, we can see that most laser measurements on the car are absent (black pixels). This could be either due to the dark paint that absorbs laser light or the specular surfaces, such as windows, that divert reflected lasers away from the LiDAR receiver. As most measurements are missing, only a small fraction of the car can be seen in the point cloud, increasing the difficulty of car detection. However, these unusual measurements are identified and assigned with higher uncertainty scores that distinguish themselves from the background, and such an area should be regarded with extra precaution for autonomous vehicles.

correctly reflects the degradation changes exerted by the rainfall.

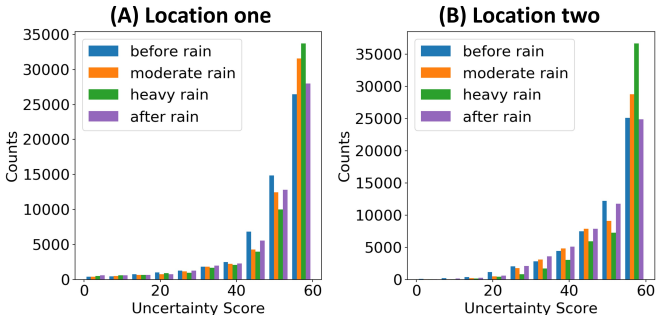

Fig. 5. **Visualization of uncertainty estimation change during the rainfall.** We collect a sequence of LiDAR scans stationarily at two different locations and select one scan each before the rain, during moderate rain, during heavy rain, and after the rain. The uncertainty score distributions of laser measurements from those scans are displayed in (A) and (B). The number of laser measurements with higher uncertainty scores increases when the rain intensity increases and then decreases after the rain. This means that the uncertainty estimation correctly identifies more unreliable measurements due to heavier degradation by rain.

D. Quantitative Results

In this experiment, we create a binary classification task with two classes: low-uncertainty class and high-uncertainty class, and evaluate our model as a binary classifier. First, we create a binary label for each laser measurement in Natural Driving Dataset (NDD) and Synthetic Degradation Dataset (SDD) following these steps: For the NDD, we select a few scans before the rain and convert them all into LiDAR images. At each pixel location, we accumulate laser measurements from different scans and group them into clusters. After that, for each laser measurement from rainy weather, we attach it with a low-uncertainty label if its measurement belongs to any of these clusters at that location or otherwise the high-uncertainty label. For the SDD, we select a few scans collected in clear weather, randomly select a percentage of laser measurements in these scans, modify their measurements (either remove their measurement values or add noise), and attach them with high-uncertainty labels. For the remaining measurements that are not selected, we assign them with the low-uncertainty labels. Next, we run our model to predict the uncertainty score for each laser measurement and set the decision threshold to k. The evaluation metric $acc@k$ represents the fraction of correct predictions evaluated at k.

Experimental data were collected from three different locations for each dataset. Results are listed in Table I. For the NDDs, the highest average accuracy is 0.83, and the lowest average accuracy is 0.75. For the SDDs, the highest accuracy is 0.71, and the lowest accuracy is 0.61. The accuracies for SDDs are relatively lower than those in NDDs,

TABLE I
QUANTITATIVE EVALUATION RESULTS

Dataset	Degradation	Duration (min)	Evaluation Results				
			$acc@0.1$	$acc@0.2$	$acc@0.5$	$acc@0.8$	$acc@avg$[3]
Natural Driving Dataset-1	natural	19.2 / 32.3 / 5.0[1]	0.80	0.81	0.84	0.85	0.83
Natural Driving Dataset-2	natural	32.9 / 18.2 / 2.0	0.85	0.84	0.83	0.83	0.84
Natural Driving Dataset-3	natural	4.0 / 65.6 / 27.3	0.80	0.77	0.73	0.71	0.75
Synthetic Degradation Dataset-1	synthetic	2.0[2]	0.81 / 0.85 / 0.80[4]	0.75 / 0.81 / 0.73	0.68 / 0.72 / 0.62	0.58 / 0.61 / 0.57	0.71
Synthetic Degradation Dataset-2	synthetic	2.0	0.69 / 0.84 / 0.76	0.70 / 0.81 / 0.70	0.65 / 0.70 / 0.59	0.57 / 0.60 / 0.54	0.68
Synthetic Degradation Dataset-3	synthetic	2.0	0.69 / 0.76 / 0.77	0.62 / 0.70 / 0.71	0.50 / 0.59 / 0.61	0.40 / 0.49 / 0.54	0.61

[1] 1: duration for light rain / moderate rain / heavy rain degraded measurements to each scan [2] 2: duration for clear weather [3] 3: average across k = 0.1, 0.2, 0.5, 0.8 [4] 4: added 10%, 20% and 30%

but such a difference is reasonable. Our model learns the distribution of reliable measurements from real-world data. The simulated degradation added to SDDs is very likely to have a distribution different from the natural degradation.

E. Applications

Uncertainty estimation can be combined with downstream tasks to assist autonomous driving systems in tackling adverse weather challenges. For example, Fig. 6 demonstrates three relevant applications. The first application is sensor degradation evaluation. Sensor degradation evaluation is necessary for navigation in adverse weather because sensor reliability could decrease with the weather and negatively affect the system's safety. Also, sensor degradation information can help autonomous driving systems adjust driving strategies to respond to different levels of degradation. In Fig. 5, we have seen that the histogram skews towards the right when rainfall intensity increases. Therefore, we calculate the skewness of the histogram as the amount of degradation contained in each LiDAR scan. The degradation evaluation is shown in Fig. 6 (A). The degradation score increases when the rainfall starts and decreases slowly when the degradation effects exerted by the rainfall gradually diminish.

The second application is to filter out unreliable LiDAR measurements caused by rain, which could help localization algorithms extract reliable environmental features for pose estimation. Fig. 6 (B) shows the filtering result. After the filtering, most structural features, e.g., the outline of building walls, are still preserved while the majority of unreliable measurements, e.g., reflections from raindrops, are discarded. This could potentially improve the stability of the localization system.

The third application is to enhance perception performance. Currently, most perception algorithms struggle to detect dark-colored or overly-reflective vehicles because the laser beam projected on the car could be absorbed heavily or diverted away from the receiver, causing very few laser beams to be reflected back to the sensor. As a result, the measurements on the car will be very sparse, thereby increasing the missed detection rate. Fig. 6 (C) shows such a case where only a small fraction of the car can be seen in the LiDAR point cloud. Our uncertainty estimation method successfully identifies that area as unusual and assigns large uncertainty scores to distinguish it from background scenes. This uncertainty estimation could potentially compensate for the limitation of existing perception algorithms.

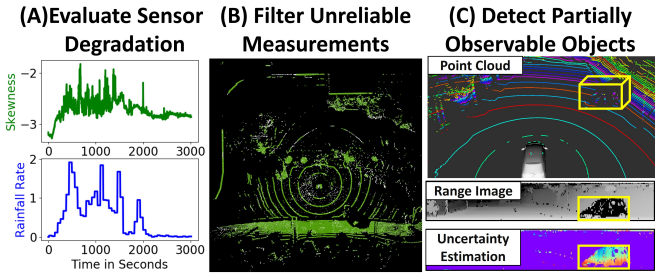

Fig. 6. **Three applications for uncertainty estimation.** (A) Evaluate LiDAR sensor degradation. Uncertainty estimation results can be converted into degradation scores representing the amount of degradation in each LiDAR scan. The navigation system can use the degradation scores to respond intelligently to different degradation levels. (B) Filter out unreliable measurements to increase the stability of feature-based localization systems. Environmental features (e.g., the shape of walls) can still be clearly seen in the filtered point cloud (green). (C) Assist perception systems in detecting partially-observed objects. In the LiDAR point cloud (top) and range image (middle), only a small fraction of measurements on the car is visible, making car detection difficult. However, in the uncertainty estimation result (bottom), the absent measurements on the car are assigned with high uncertainty scores that distinguish the car from the background. This could help reduce the missed detection rate.

V. CONCLUSION

Autonomous driving in adverse weather has been widely recognized as a critical problem for performing autonomous tasks. This work presents a novel method to address sensor degradation problems. Our method estimates the uncertainty of each laser measurement using an attention-based Mixture Density Network to capture the spatial correlation between laser measurements and then predict unreliable measurements. Using uncertainty information can potentially improve the performance of downstream tasks, such as LiDAR-based perception and localization, or integrate with decision-making modules to create intelligent driving strategies to deal with degradation effects. In the future, more experiments in dynamic environments will be discussed.

ACKNOWLEDGMENTS

This research was partially supported by the National Research Foundation, Prime Minister's Office, Singapore, under its CREATE programme, Singapore-MIT Alliance for Research and Technology (SMART) Future Urban Mobility (FM) IRG.

REFERENCES

[1] X. Jin, Z. Chen, and W. Li, "Ai-gan: Asynchronous interactive generative adversarial network for single image rain removal," *Pattern Recognition*, vol. 100, p. 107143, 2020.

[2] R. Li, L.-F. Cheong, and R. T. Tan, "Heavy rain image restoration: Integrating physics model and conditional adversarial learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1633–1642, 2019.

[3] R. Heinzler, F. Piewak, P. Schindler, and W. Stork, "Cnn-based lidar point cloud de-noising in adverse weather," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2514–2521, 2020.

[4] N. Charron, S. Phillips, and S. L. Waslander, "De-noising of lidar point clouds corrupted by snowfall," in *2018 15th Conference on Computer and Robot Vision (CRV)*, pp. 254–261, IEEE, 2018.

[5] M. Hnewa and H. Radha, "Object detection under rainy conditions for autonomous vehicles: A review of state-of-the-art and emerging techniques," *IEEE Signal Processing Magazine*, vol. 38, no. 1, pp. 53–67, 2020.

[6] V. Zuev, "Laser-light transmission through the atmosphere," in *Laser Monitoring of the Atmosphere*, pp. 29–69, Springer, 1976.

[7] C. Goodin, D. Carruth, M. Doude, and C. Hudson, "Predicting the influence of rain on lidar in adas," *Electronics*, vol. 8, no. 1, p. 89, 2019.

[8] S. Hasirlioglu and A. Riener, "A general approach for simulating rain effects on sensor data in real and virtual environments," *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 3, pp. 426–438, 2019.

[9] S. Hasirlioglu, I. Doric, C. Lauerer, and T. Brandmeier, "Modeling and simulation of rain for the test of automotive sensor systems," in *2016 IEEE Intelligent Vehicles Symposium (IV)*, pp. 286–291, IEEE, 2016.

[10] A. S. Mohammed, A. Amamou, F. K. Ayevide, S. Kelouwani, K. Agbossou, and N. Zioui, "The perception system of intelligent ground vehicles in all weather conditions: A systematic literature review," *Sensors*, vol. 20, no. 22, p. 6532, 2020.

[11] S. Zang, M. Ding, D. Smith, P. Tyler, T. Rakotoarivelo, and M. A. Kaafar, "The impact of adverse weather conditions on autonomous vehicles: how rain, snow, fog, and hail affect the performance of a self-driving car," *IEEE vehicular technology magazine*, vol. 14, no. 2, pp. 103–111, 2019.

[12] A. Filgueira, H. González-Jorge, S. Lagüela, L. Díaz-Vilariño, and P. Arias, "Quantifying the influence of rain in lidar performance," *Measurement*, vol. 95, pp. 143–148, 2017.

[13] R. Heinzler, P. Schindler, J. Seekircher, W. Ritter, and W. Stork, "Weather influence and classification with automotive lidar sensors," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1527–1534, IEEE, 2019.

[14] C. Zhang, Z. Huang, M. H. Ang, and D. Rus, "Lidar degradation quantification for autonomous driving in rain," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3458–3464, IEEE.

[15] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.

[16] M.-T. Luong, H. Pham, and C. D. Manning, "Effective approaches to attention-based neural machine translation," *arXiv preprint arXiv:1508.04025*, 2015.

[17] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[18] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.

[19] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, "Improving language understanding by generative pre-training," 2018.

[20] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.

[21] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.

[22] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, "End-to-end object detection with transformers," in *European conference on computer vision*, pp. 213–229, Springer, 2020.

[23] M. Chen, A. Radford, R. Child, J. Wu, H. Jun, D. Luan, and I. Sutskever, "Generative pretraining from pixels," in *International Conference on Machine Learning*, pp. 1691–1703, PMLR, 2020.

[24] R. Liu, Z. Yuan, T. Liu, and Z. Xiong, "End-to-end lane shape prediction with transformers," in *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 3694–3702, 2021.

[25] J.-B. Cordonnier, A. Loukas, and M. Jaggi, "On the relationship between self-attention and convolutional layers," *arXiv preprint arXiv:1911.03584*, 2019.

[26] C. M. Bishop, "Mixture density networks," 1994.

[27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.