Towards Characterizing the Behavior of LiDARs in Snowy Conditions

Sébastien Michaud[1], Jean-François Lalonde[2], and Philippe Giguère[1]

[1]Computer Science and Software Engineering, Laval University
[2]Electrical and Computer Engineering, Laval University

***Abstract*— Autonomous driving vehicles must be able to handle difficult weather conditions in order to gain acceptance. For example, challenging situations such as falling snow could significantly affect the performance of vision or LiDAR-based perception systems. In this paper, we are interested in characterizing the behavior of LiDARs in snowy conditions, as there seems to be little information publicly available. In particular, we present a characterization of the behavior of 4 commonly-used LiDARs (Velodyne HDL-32E, SICK LMS151, SICK LMS200 and Hokuyo UTM-30LX-EW) during the falling snow condition. Data was collected from the 4 sensors simultaneously during 6 snowfalls. Statistical analysis of these datasets indicates that these sensors can be modeled in a probabilistic manner, allowing the use of a Bayesian framework to improve robustness. Moreover, we were able to observe the temporal evolution of the impact of the falling snow during these snowstorms, and characterize the sensitivity of each device. Finally, we concluded that the falling snow had little impact beyond a range of 10 m.**

I. Introduction

The robustness of autonomous vehicles has increased prodigiously in the recent years. While long-range autonomous driving on the highway has been around for decades already [1], advances in mapping, 3D data processing and computer vision have enabled cars to drive autonomously for thousands of miles in unconstrained, city environments [2]. While this surely is an impressive feat, one quickly notes that most of these miles have been logged in California weather, which provides optimal operating conditions for sensors such as LiDARs. In order for these systems to gain acceptance worldwide, it is crucial that they could be operated in more challenging weather conditions, such as rain, fog and snow.

As we strive to make autonomous vehicles more adaptable to varying weather conditions, it is important to understand how sensors will behave in such conditions. Of particular interest, snowy conditions may cause challenging situations for sensors such as LiDARs. Indeed, the laser beams emitted may illuminate the snowflakes themselves, thus providing echoes that do not correspond to real obstacles. Consider fig. 1 for example. The same scene appears drastically different depending on whether it was captured on a clear or snowy day. While programmable lighting may help circumvent this problem [3], current LiDARs may fail under such circumstances.

In this paper, our main contribution is to provide a characterization of the behavior of four well-known LiDARs in snowy conditions. Through an extensive empirical study performed on a novel dataset captured under varying degrees of snowfall, we evaluate how much these LiDARs are sensitive—or not—to falling snow. We show that recent advances in sensor design have increased their robustness even to significant snowfall.

Corresponding author: philippe.giguere@ift.ulaval.ca

Fig. 1. Driving in bad weather. While autonomous vehicles have attained a great level of performance in nice weather (left), bad weather can cause significant challenges due to limited visibility (right). In this paper, we characterize the behavior in snowy conditions for oft-used sensors in autonomous cars: LiDARs. *Photo credit*: Nicole Duchesne (left), Gaetan Chevalier (right).

A. Related work

It is well-known that snow poses significant challenges to sensors mounted on-board outdoor mobile robots or other autonomous vehicles. For example, in their Antarctica exploration project, Moorehead et al. indicate that "in heavy [snow] storms, [...] the laser could not be used" [4]. Similarly, Yamauchi et al. relate that "LiDAR and stereo vision provide greater accuracy and resolution in clear weather but has difficulty with precipitation and obscurants" [5]. Common approaches for dealing with this problem include filtering 3-D data [4], or video [6], but this is often not enough to completely remove artifacts.

It is therefore important to characterize how sensors behave in such conditions. To this end, Sumi et al. [7] build a specifically designed simulated snow chamber, with white polystyrene beads flown with large fans to simulate snow. In our case, we use real world conditions to acquire a novel dataset of more than 6 days of snowfall.

Finally, we also mention the work of Servomaa et al. [8], who use LiDARs (and other sensors) to characterize snow storms for monitoring and measurement applications. In our case, we characterize the behavior of the sensors themselves for robotics applications.

II. Data acquisition

In this section, we first report on the relevant characteristics of the four sensors used in our dataset. We then describe the physical configuration of our test setup, then outline the weather conditions for each of the six collected snowfalls. Finally, we describe how the information from the LiDARs was preprocessed before analysis.

A. Sensors

Data acquisition was performed with the following four LiDARs: the SICK LMS200, SICK LMS151, Hokuyo UTM-30LX-EW, and the Velodyne HDL-32E. Relevant sensor information is provided in table I, but the reader is referred to the manufacturers documentation for additional information[1].

The first element that provides a qualitative overview of the sensor performance is the maximum acquisition distance. This value depends on several factors, such as lighting conditions and target remission. This value is provided directly for the HDL-32E and UTM-30LX-EW, but based on a target remission greater than $75\,\%$ for the LMS200 and LMS151. Another element to consider is the shape and area covered by the beam, which influences the probability of hitting a snowflake as well as the proportion of area it covers. A final significant element which changes from one sensor to the other is the number of echoes returned. The Hokuyo sensor can return up to three echoes, which means that it could locate two snowflakes before the beam reaches the ground. Regarding the LMS151, two echoes are evaluated by the hardware, but only one is returned. Finally, note that all LiDARs use class 1 laser with a wavelength of $905\,\mathrm{nm}$.

B. Setup configuration

Data acquisition was conducted at Pouliot Hall of Laval University, where sensors were placed close to the inner wall of a window facing N50°E. As shown in fig. 2, a wooden structure held the sensors side by side at approximately $14\,\mathrm{m}$ above the ground. The main scanning plane (i.e. XY plane in the sensor reference frame) formed a $30°$ angle with respect to the building wall, so as to increase the maximum distance as much as possible without having the laser beams hitting trees or a pedestrian walkway present near the building. In addition, an RGB camera was placed alongside the LiDARs to provide visual information about the scene. In this configuration, a slight opening of the window allowed to keep the instruments inside while scanning outside. To avoid direct interference between sensors, corrugated plastic layers were placed between them. Fig. 3 shows the scene as observed by the RGB camera placed with the sensors.

C. Dataset description

Data acquisition started February 12 and ended on March 2. A total of 10 episodes were collected for a total of more than 50 hours of data. Recordings were made using the Robot Operating System (ROS) [13], which provides standardized data types as well as time synchronization. Data was acquired at different times of day and in a wide variety of conditions, covering a wide range of snowflakes size, falling rate and wind speed. Table II provides an overview of our data[2]. Of these, six are used in the current study, as highlighted in this table. The dataset is publicly available upon request.

[1]Available here: Velodyne [9], Hokuyo [10], LMS151 [11], LMS200 [12]

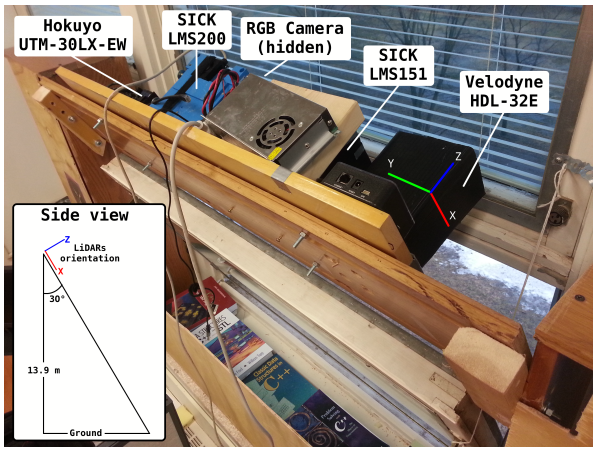

Fig. 2. The experimental setup. The 3D axis represent the orientation of the sensors and the bottom left panel represent the 2D geometry as seen from the right side of the picture.

Fig. 3. View from the RGB camera.

D. Pre-selection of laser data

For each sensor, we selected a combination of angles and laser rings (for the Velodyne) or angles (for the others) that had a clear view of the snow-covered ground surface. Details for each sensor are provided in table III. The range of the ground in our scans was between $x = 15\,\mathrm{m}$ to $x = 22\,\mathrm{m}$, depending on the angle. To simplify the analysis, we considered as a snowflake echo any measurement which had a range reading of $x < 14.5\,\mathrm{m}$. As will be shown later in sec. IV, this approximation is valid as the vast majority of those events happened for $x < 10\,\mathrm{m}$.

[2]Wind speed, daily precipitation and temperature were measured at Québec City Jean Lesage International Airport, located at a distance of $9\,\mathrm{km}$ from Laval University. Data is available here [14].

Sensor	Maximum distance	Spot area (at 30 meters)	Spot shape	Echoes
SICK LMS200	28 m	165 cm^2	Circle	1
SICK LMS151	50 m	22 cm^2	Circle	2
Hokuyo UTM-30LX-EW	30 m	196 cm^2	Ellipse	3
Velodyne HDL-32E	70 m	51 cm^2	Rectangle	1

TABLE I

OVERVIEW OF CHARACTERISTICS SPECIFIC TO EACH LIDAR.

Beginning time	Duration (HH:MM)	Snowflakes size	Falling rate	Wind speed range ($km\,h^{-1}$)	Daily precipitation (cm)	Temperature (°C)
Feb 12, 9:47 am	09:21	Small	Variable	[2–13]	1.4	-14.1
Feb 14, 10:12 pm	04:12	Small	Very low	[5–13]	0.2	-21.4
Feb 19, 8:38 am	10:02	Big/small	High	[3–28]	4.5	-10.9
Mar 2, 1:06 pm	01:27	Big/small	Variable	[22–36]	1.6	-9.1
Mar 3, 10:33 pm	02:17	Big	Medium	[7–9]	5.4	-13.3
Mar 4, 11:45 am	04:12	Big/medium	Low/none	[20–30]	2.0	-4.3
Mar 17, 10:08 am	06:08	Big/medium	Low/none	[1–31]	2.0	-5.8
Mar 21, 6:44 pm	07:42	Medium/big	High	[5–33]	8.6	-5.1
Mar 30, 1:06 pm	04:45	Medium/big	High	[4–8]	8.5	-3.0
Apr 2, 1:56 pm	01:51	Small/rain	High	[2–10]	1.2	-8.4

TABLE II

OVERVIEW OF OUR SNOW DATASET. DATES IN BOLD CORRESPOND TO THE SIX DAYS USED IN THE PRESENT STUDY.

Sensor	Acquisition frequency	Selected beams/angles	Selected rings	Window size
LMS200	9.375 Hz	55–115	N/A	106 s
LMS151	25 Hz	310–220	N/A	40 s
Hokuyo	20 Hz	440–590	N/A	100 s
Velodyne	10 Hz	-0.05–0.25 rad	17–31	40 s

TABLE III

DETAILS OF MEASUREMENT SELECTION FOR THE ANALYSIS. THE WINDOW SIZE IS THE TEMPORAL WINDOW USED TO CALCULATE STATISTICS DURING THE TEMPORAL EVOLUTION OF A SNOWFALL.

III. TEMPORAL ANALYSIS

In this section, we analyze the temporal behavior of the four sensors for the duration of six complete snowstorms. In particular, we are interested in seeing how the fraction of echoes in snowflakes evolves over time, for all four sensors. First, we will discuss the highly dynamical nature of snowstorms. This will be exemplified by how consecutive scans can have significant quantitative and spatial differences in the distributions of the snowflakes echoes, which justify the use of averaging windows for our analysis. We will then present the actual temporal evolution of these statistics in the form of graphs for all four sensors, and finally briefly discuss the results for each sensor.

A. Extraction of temporal statistics

Snowstorms are highly dynamic processes, with large variation in snowfall rates over their durations. Moreover, the snow physical characteristics (size, shape or reflectance) might vary significantly during a storm, affected by ambient conditions such as humidity level and temperature. Also, wind gusts might pull snow back up in the air or drive it sideways, affecting its effective fall rate. Consequently, one expects during a snowstorm to see significant short, medium and long term variations in the fraction of LiDAR echoes corresponding to the falling snow.

Computing and reporting the temporal statistics for every scan would put too much emphasis on the very short-term statistics. Indeed, the inter-scan variation in the fraction of snowflake echoes can be significant. To better illustrate this point, we have overlaid four consecutive scans in the same plot for the LMS200 and for the first echo returned by the multi-echo Hokuyo sensor in fig. 4, for an intense snowing episode from the 02-19 dataset (see tab. II). In these figures, we can see strong variations in the fraction of snowflake echoes and their spatial distribution. One can readily see the fluctuation in these fractions as reported in the brackets of the legend in fig. 4.

To smooth out these fluctuations, statistics are extracted from a number of consecutive scans contained in a time window of around 1 minute (detailed values in tab. III). Fig. 5 shows this smoothed fraction of snowflake echoes compared to all returned laser measurements as a function of time, for the six snowiest days of our dataset. To allow for better visualization, only the LMS200 and the Hokuyo's first echo are plotted at their actual scale (1x): Others have been scaled up (from 30x to 200x), with their corresponding scaling factors reported in the legend. As will be shown below, some sensors were much more sensitive than others.

B. Detailed analysis, per sensor

1) SICK Sensors LMS200 and LMS151: Our first conclusion based on fig. 5 is that the most sensitive device was the older LMS200, first introduced in the mid-2000s. For the most intense snowstorms (fig. 5. b) 02-19, d) 03-17, e) 03-21 and f) 03-30), it peaked at around 15% of measurements triggered by the falling snowflakes, for averaging windows of 106 s. As an older-generation device, it probably uses less sophisticated algorithms and sensing, and was not directly

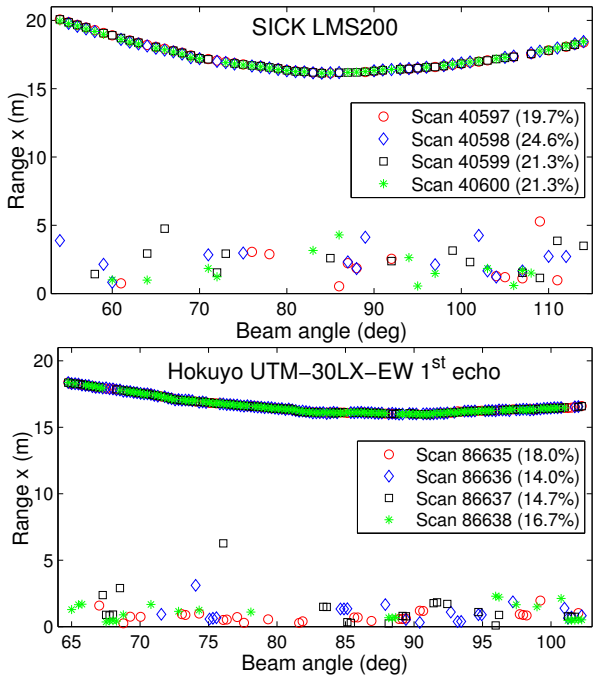

Fig. 4. Four overlaid consecutive scans for the LMS200 sensor (top), and the first echo scans for the Hokuyo sensor (bottom), taken from the 02-19 dataset. Each symbol corresponds to a particular scan. The curved line at the top corresponds to the snow surface on the ground. One can see the rapid variation of the snowflake echoes between scans, and how they are mostly limited to a range $x < 5\,\mathrm{m}$. The percentages (in brackets) are the proportion of those echoes in the snowflakes.

targeted for harsh outdoor environments. Indeed, its technical description [12] indicates that "raindrops and snow-flakes are cut out using pixel-oriented evaluation", but this seems only applicable to obstacle detection (field computation), not the actual measurements. No further details are given. On the other hand, the more recent SICK LMS151 exhibits much less sensitivity to snowflakes: The reduction factor for the fraction of snowflakes echoes is in the order of 200-300, granting this device a much higher immunity to snowstorms. Indeed, the highest peak was around 0.1 % of echoes in snowflakes during the 02-19 dataset. This seems to support the claim, obtained in the documentation from the manufacturer, that this model is targeted for "all weather conditions" [15].

2) Hokuyo UTM-30LX-EW: For this sensor, we resorted to a slightly different approach for comparison, as the device has been designed to return multiple echoes. We thus extracted statistics for the two most relevant cases: the first and last echoes. Statistics for the first echo indicate how sensitive the device is, if one wishes to detect the presence or absence of falling snow. This information could be used, for example, to adapt the driving strategy of an autonomous vehicle or inform vision algorithms of the presence of particles in the air. Using the last echo increases the probability that obstacles, such as another vehicle or the snow-covered ground, will be detected. This information would be used for localization and navigation purposes.

For the first echo, we observed that the device behaved similarly to the LMS200. Indeed, the Hokuyo first echo (blue line) closely tracks the LMS200 curves (red dashed line) almost everywhere in fig. 5, with a few exceptions. When using the last echo, the sensor behaves like the LMS151, not surprisingly as this sensor performs a 2-echo analysis and filtering. The last echo of the Hokuyo tends to reject the falling snow, but not as well as the LMS151, as it peaked at around 0.5 % in some episodes. Nevertheless, this difference might not be sufficient to impact algorithms relying on laser data. Note that tab. IV shows similar correlations between these three sensors for averages taken over the complete 02-19 dataset.

LMS200	Hokuyo first echo	Hokuyo last echo	LMS151	Velodyne HDL-32E
2.67%	3.55%	0.0113%	0.00178%	0.0100%

TABLE IV

OVERALL AVERAGE SNOWFLAKE ECHOES FOR THE COMPLETE 02-19 DATASET, PER SENSOR. THESE AVERAGES ARE SIGNIFICANTLY LOWER THAN THE INSTANTANEOUS VALUES DISPLAYED IN FIG. 5, AS SNOW WAS NOT FALLING AT ALL TIMES DURING THAT PERIOD.

3) Velodyne HDL-32E: For all purposes, the behavior of the Velodyne was similar to the last echo of the Hokuyo sensor. This is seen both in the temporal behavior in fig. 5 and in the average value displayed in tab. IV.

IV. Distribution of Snowflake Echoes as a Function of Range

In the previous section, we showed how the expected fraction of snowflake echoes varied temporally during snowstorms. In some sense, it provided for a *temporal* modeling of the interaction between a snowstorm and a given LiDAR. In this section, we evacuate the temporal aspect and instead focus on how the range x affects the probability for a snowflake to trigger a measurement. To this end, we will use histograms to estimate a probability density function of those events, and show that for the weather conditions and the sensors we tested, there seems to be an upper bound on the range x beyond which falling snowflakes no longer trigger a measurement: in other words, snowflakes become invisible to the sensor past a certain range.

A. Modeling the impact of range on snowflake detection

When modeling a range sensor, one has to obtain the probability distribution of certain events (e.g. snowflakes) as a function of this range. Over the years, many researchers have proposed probabilistic models for sensors, notably [16]. In the previous section, we have estimated the probability for a given sensor S that a snowflake would generate an echo $E_{\text{snowflake}}$ given the weather condition W, or $P_S(E_{\text{snowflake}}|W)$. In this section, we take a closer look at which range x such events would be generated, that is $P_S(E_{\text{snowflake}}|x, W)$. Having such a formulation would allow for a more statistically-sound treatment of the information, such as within a Bayesian probabilistic framework. To this effect, we use histograms as approximations for the previous

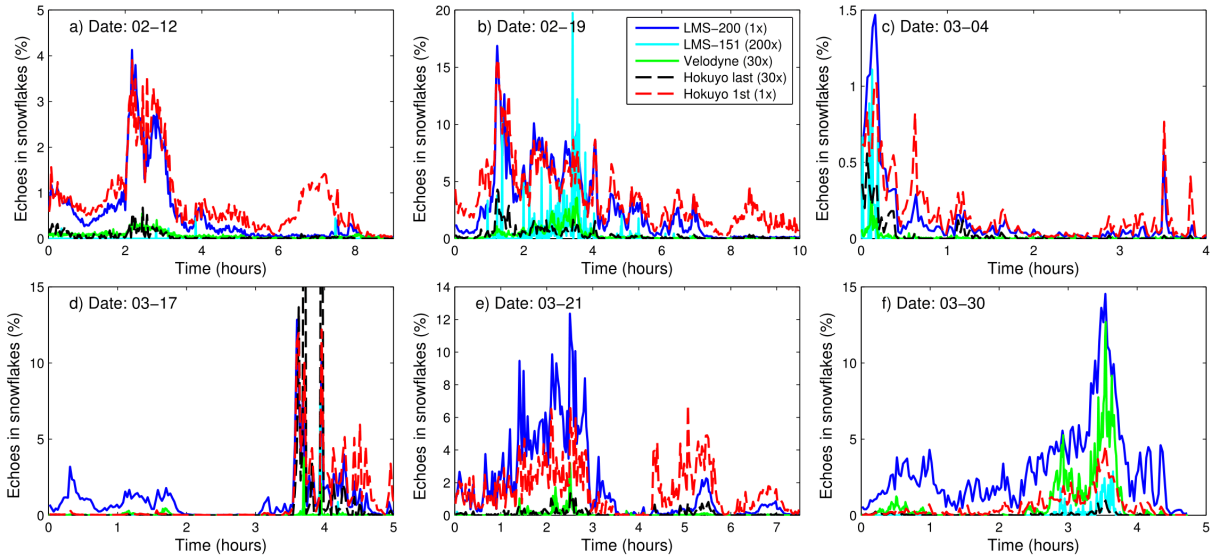

Fig. 5. Temporal evolution of the percentage of echoes coming from the falling snow (range x <5 m) during the 6 most intense episodes, for all 4 sensors. The data is smoothed by taking statistics for small time windows. Except for the LMS200 and Hokuyo first echo, all other sensors statistics have been scaled up (factor in bracket of legend b) for ease of visual comparison. Time is in hour, starting from the beginning of the data capture sequence.

distribution. In fig. 7, we have plotted these histograms for each of the four sensors. For ease of comparison, they have all been normalized by their total area in the interval $0 < x < 14\,\mathrm{m}$, as the total count varies widely between the sensors. The numbers in brackets in the legend indicate the fraction of echoes generated by snowflakes compared to the total number of data points, for a given dataset.

The general shape of these histograms is close to a log-normal distribution, with the exception of the LMS200 for a number of dates (02-12 through 03-17), which seems to follow a sum of two log-normal distributions. We attribute this log-normal shape to the interaction between two different phenomena, illustrated in a cartoon-type model in fig. 6. At short ranges $x < 3\,\mathrm{m}$, the building acts as a shield and decreases the probability of having a snowflake in the path of the laser. We recognize that this phenomenon would be most likely absent on an autonomous vehicle, thereby increasing the probability of having echoes in snowflakes at close range. However, we believe that this difference is not problematic, as close obstacles would be easily detected from *i)* the overwhelming number of LiDAR echoes on this obstacle *ii)* other sensing modalities such as vision or radar. Furthermore, if the LiDAR is to be mounted on a rooftop, one can safely ignore echoes in the first 2 m, either in software or directly through the sensor itself (via its configuration). The other phenomenon, illustrated as the red dashed line in fig. 6, is the probability of optical detection of a snowflake by the sensor as a function of the range x. We argue that this shape is due to the rapidly decreasing light intensity of the echoes in snowflakes, as a function of x. Combining these two phenomenon yields a log-normal shaped curve (black line in fig. 6). Overall, this seems to indicate that a simple probabilistic model $P_S(E_{\text{snowflake}}|x, W)$ can be derived for these sensors.

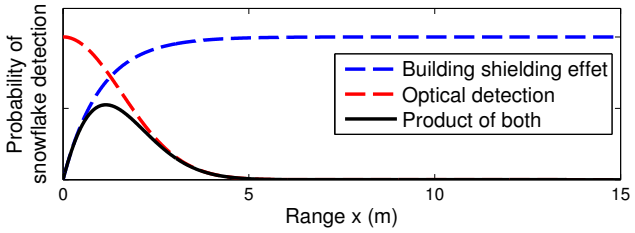

Fig. 6. Cartoon representation of the interaction between the probability of detecting a snowflake (in red) and the diminution of snowflakes due to the shielding effect of the building (in blue). The black line is the product of the two, and bear a close resemblance to the actual histograms extracted from our dataset.

B. Sensor results

As can be seen from the histograms in fig. 7, most sensors exhibit the log-normal or sum-of-log-normal distributions discussed above. We note that for certain days, the distributions are shifted to the right (greater range x). In particular, for the 03-21 and the 03-30 distributions, this shift is substantial (on the order of 1 m). We suspect that for these days, the snowflakes were significantly larger, thus allowing for a stronger optical echo and extended range of detection.

For all sensors, we can also conclude that beyond the range $x > 10\,\mathrm{m}$, snowflakes are no longer detected, i.e. they become invisible. A small notable exception would be for the Velodyne, for which snowflakes were detected all the way to $x = 14\,\mathrm{m}$, albeit at a significantly reduced rate. Again, we do not think that this would significantly impair their use in conditions similar to our test setup.

V. Discussion and Conclusion

In this paper, we explored the impact of falling snow on the usability of 4 commonly deployed LiDARs in the context of autonomous driving vehicles. To this end, we collected

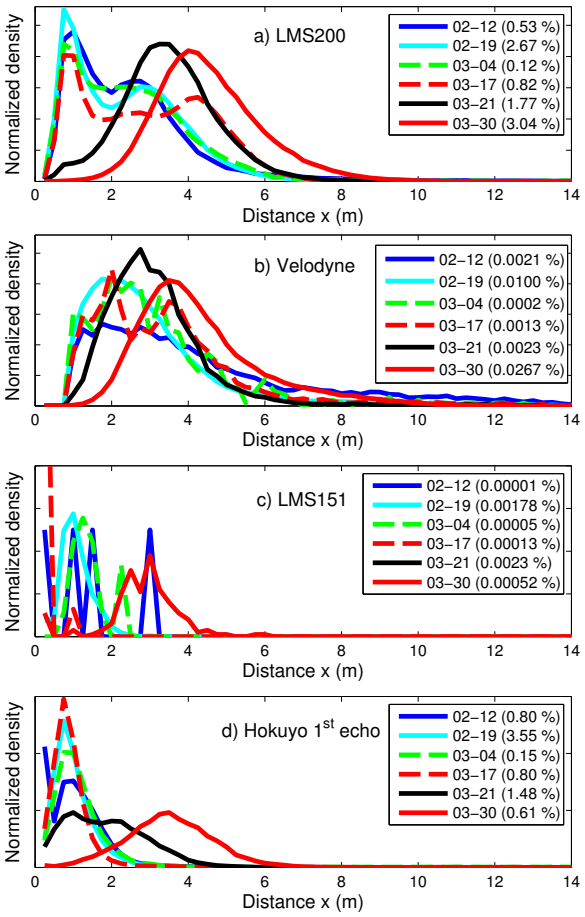

Fig. 7. Histograms of echoes in falling snow during important snowfall days, as a function of distance x reported by the sensor. Each histogram has been normalized by its area, for ease of comparison. The numbers in brackets are the fraction of data points in the complete dataset that correspond to snowflake echoes. Note that for the 03-21 dataset, the LMS151 was not working properly: thus no data is included for that day.

data during 6 snowstorms in the winter of 2015. Upon analysis, we found that the SICK LMS200 was the most sensitive LiDAR, having a peak average rate of up to 15 % of echoes coming from falling snow. Meanwhile, all 3 others never exceeded 1 %. We also presented a simple probabilistic model to take into account the effect of the range on snowflakes interference. Based on a histogram analysis, we concluded that for our experimental setup, this model can be approximated by a log-normal distribution. Most importantly, our data indicate that the impact of snowflakes on LiDAR beyond a range of $10\,\mathrm{m}$ is very limited.

A number of questions remain to explore. For example, as the LiDAR beam travels through the falling snow, its intensity will diminish. Since the maximum range of a LiDAR is heavily related to this beam intensity, we expect the maximum range to be affected during snowstorms. In our setup, we have not witnessed this issue, indicating that this effect probably happens beyond our maximum distance of $20\,\mathrm{m}$. Another aspect to be investigated is the relationship between the returned intensities and the surface type (ground or snowflakes). Also, because of the shielding effect of the building, very few snowflakes were present at close range; It might be the case that at closer range, a snowflake might be detected at more than one angle, effectively occluding small targets. Moreover, we have not investigated the impact on the measurement noise for the snowy ground surface in the presence of falling snow. Finally, it would be interesting to mount these LiDARs on a moving vehicle to investigate the impact of the vehicle velocity on the sensing behavior.

References

[1] D. Pomerleau and T. Jochem, "Rapidly adapting machine vision for automated vehicle steering," *IEEE Expert: Special Issue on Intelligent System and their Applications*, vol. 11, no. 2, pp. 19–27, April 1996.

[2] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, *et al.*, "Autonomous driving in urban environments: Boss and the urban challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.

[3] R. Tamburo, E. Nurvitadhi, A. Chugh, M. Chen, A. Rowe, T. Kanade, and S. G. Narasimhan, "Programmable automotive headlights," in *European Conference on Computer Vision*. Springer, 2014, pp. 750–765.

[4] S. Moorehead, R. Simmons, D. Apostolopoulos, and W. L. Whittaker, "Autonomous navigation field results of a planetary analog robot in antarctica," in *International Symposium on Artificial Intelligence, Robotics and Automation in Space*, June 1999.

[5] B. Yamauchi, "Fusing ultra-wideband radar and lidar for small ugv navigation in all-weather conditions," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2010, pp. 76 920O–76 920O.

[6] P. C. Barnum, S. Narasimhan, and T. Kanade, "Analysis of rain and snow in frequency space," *International Journal of Computer Vision*, vol. 86, no. 2-3, pp. 256–274, 2010.

[7] Y. Sumi, B. Kim, Y. Matsumoto, E. Horiuchi, O. Matsumoto, and K. Ohba, "Outdoor environment simulators for vision-based safety sensors; artificial sunlight lampheads and simulated-snow chamber," in *IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*, Nov 2013, pp. 125–130.

[8] H. Servomaa, K.-i. Muramoto, and T. Shiina, "Snowfall characteristics observed by weather radars, an optical lidar and a video camera," *IEICE Transactions on Information and Systems*, vol. 85, no. 8, pp. 1314–1324, 2002.

[9] Velodyne, "Velodyne HDL-32E User's Manual."

[10] Hokuyo, "Scanning Laser Range Finder UTM-30LX-EW Specification."

[11] SICK, "SICK LMS151 Datasheet."

[12] ——, "Technical Documentation LMS200/211/221/291 Laser Measurement Systems," Dec. 2006.

[13] Willow Garage. Robot operating system. [Online]. Available: http://www.ros.org/

[14] G. of Canada. Canadian weather. [Online]. Available: http://weather.gc.ca/canada_e.html

[15] SICK, "SICK LMS1xx Operating Instructions."

[16] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.