NVIDIA DRIVE Sim

Lidar Validation | NVIDIA DRIVE Sim

Table of Contents

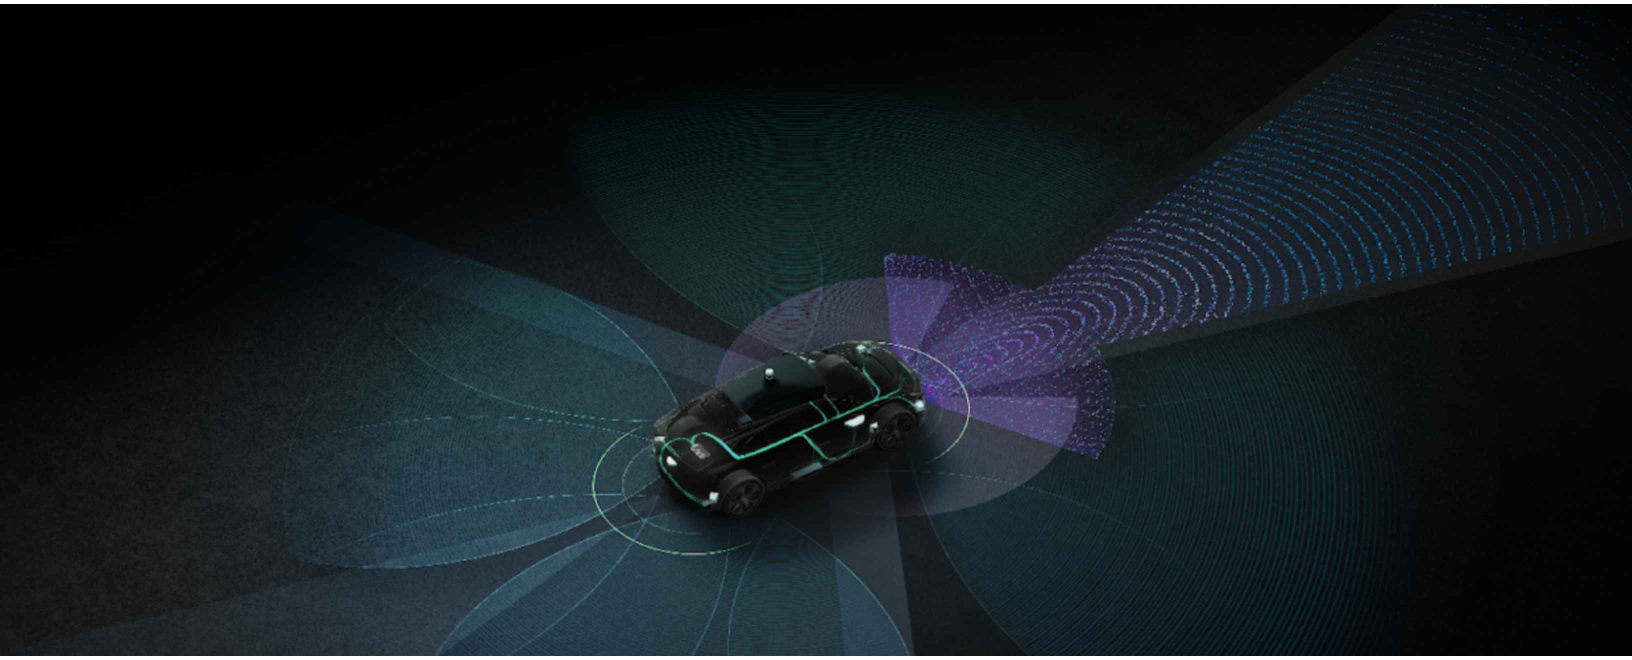

Introduction

Autonomous vehicles require an extensive amount of data to train the deep neural networks that power perception and motion control. Collecting and labeling this data is incredibly time- and cost-consuming, and in some cases, dangerous to encounter in the real world.

Highly accurate, physically based simulation addresses this problem. NVIDIA DRIVE Sim is an end-to-end simulation platform, architected from the ground up to run large-scale, physically accurate multi-sensor simulations. It allows users to generate synthetic data to train AV perception and validate motion control in a closed-loop simulation with high fidelity and accurate sensor data.

Self-driving cars use redundant and diverse sensors to collect data as well as perceive their surroundings. The sensors can be categorized into passive— cameras—and active, which include lidar, radar, and ultrasonic (USS).

This paper presents the process used to validate DRIVE Sim's active sensor models' accuracy and precision. There are multiple ways to approach sensor validation such as comparing DNNs trained on real-world data and DNNs trained on synthetic data. Also, we can validate a sensor's accuracy by comparing the synthetic data itself to the sensors' specification and real-world experiments. Each component of the model is individually validated and performs comparisons with the final output of the model.

We will focus on the latter approach as it compares the sensors' final output and takes into account the effect of all of the model's various components. As a primer to the validation topic in this post, you can read our previous analysis with camera sensor validation.

Active sensor pipeline

An active sensor sends energy into the environment and measures the backscatter reflected to the sensor to create a representation of the scene. It can either use an electromagnetic wave (lidar, radar) or an acoustic wave (USS) to generate the sensor output. The non-visual sensor pipeline is based on ray tracing and uses an iterative process as shown in Figure 1.

Rays that trace the path of light represent real-world waves and are fired into the scene in batches. For each lidar, radar, or USS ray that hits objects in the 3D scene, secondary rays are created from reflections and transmissions based on the material properties of the objects. The materials framework defines these properties using a wavelength-dependent configuration.

This process continues in an iterative manner, producing effects such as multi-bounce and time-stamp effects, including rolling shutter and Doppler effect. After the process reaches a user-defined cap on the number of batches, the RTX sensor model consolidates those returns and creates a sensor output. Post-processing nodes then use this sensor output to introduce additional noise effects and convert the data to specific protocol matching the AV stack.

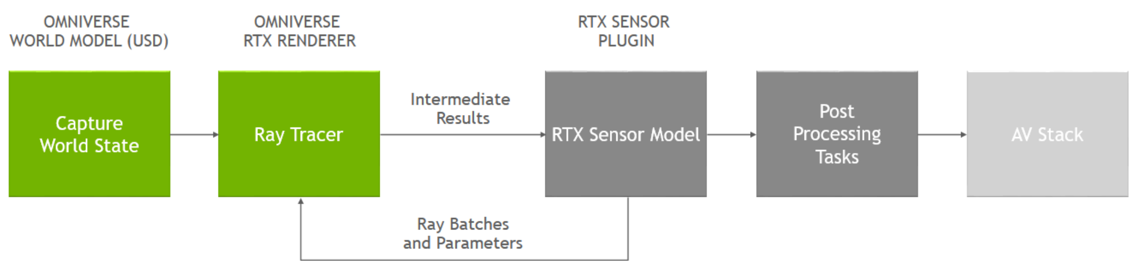

Figure 1: The active sensors pipeline

Lidar validation process

In this paper, we perform two types of comparisons to validate the lidar simulation model. The first compares the results from DRIVE Sim to the sensor's specification. The second method compares simulation results to real-world data.

For real-world data collection, two different scenarios define the extent of the data capture and are explained in detail later in this paper. Then, we create a digital twin environment in simulation, collecting the same data for detailed analysis, thereby validating the lidar model.

Validation is an extensive and comprehensive process that covers all aspects of a sensor and the data collection environment. This report covers just two scenarios, but is quantitatively conducted for validating many aspects of the lidar simulation model.

Location requirements

We outlined three major requirements for the experiment location. First, we needed a track measuring more than 100 meters in length to match the scenario execution plan.

The experiments' environment also needed to be sterilized as much as possible, requiring a location that was empty and spacious to minimize noise and unwanted reflections. By choosing a spacious and empty location, the reconstruction task in DRIVE Sim for a digital twin was greatly simplified.

We chose the Transportation Research Center (TRC) in California for the data collection environment. TRC is an advanced mobility testing site, featuring different tracks to test autonomous vehicle performance and safety, including a 1-mile intersection, 2 miles of oval track, and a 17- and 22 -acre vehicle dynamics areas. We used the intersection track and the 17 -acre vehicle dynamics area, shown in Figure 2.

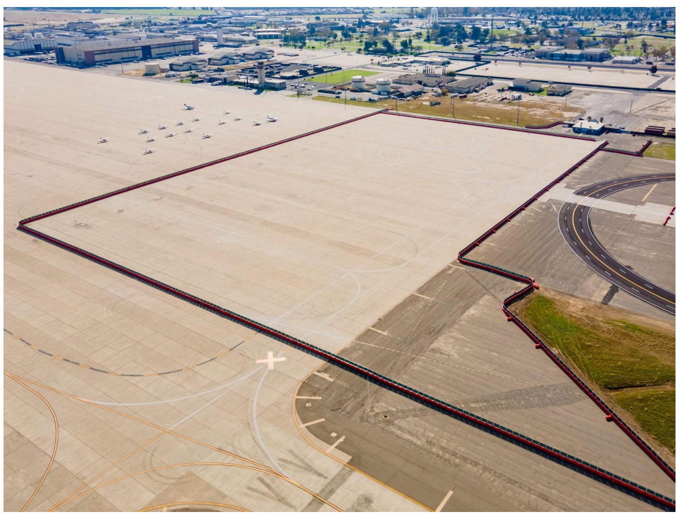

Figure 2: Aerial image of lidar test site at TRC

Vehicle setup

The lidar sensor used in the validation process is part of the NVIDIA DRIVE Hyperion platform and is available for AV development. The sensors were mounted on a development vehicle as shown in Figures 3a-c, with the 360-degree rotating lidar sensor mounted on top of the car (LD1).

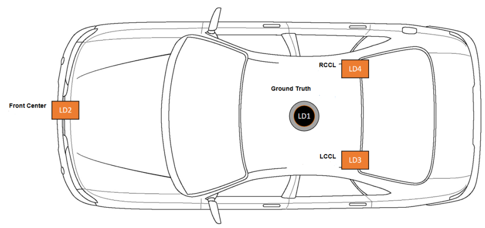

Figure 3a: Side view of vehicle

Figure 3b: Front view of vehicle

Figure 3c: Top view of lidars' sensors setup

Lidar validation

The first lidar scenario involved a firing pattern comparison between DRIVE Sim and the lidar's specification, including azimuth, elevation, and firing time deviation for each emitter of the sensor.

Elevation is the angle between the sensor XY plane and Z axis, and azimuth is the angle on the XY plane, as shown in Figure 3d.

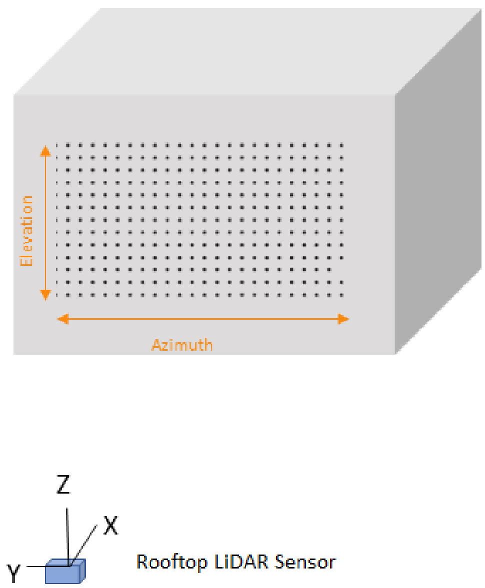

Figure 3d: Azimuth and elevation of lidar scan

Additionally, we can add configurable noise to each emitter's angle (or direction) to apply randomness, simulating imperfections in the real-world. Figures 4 through 8 detail the standard deviation for each application of the noise metric. The computed mean and standard deviation for azimuth, elevation, and firing times are identified on each plot.

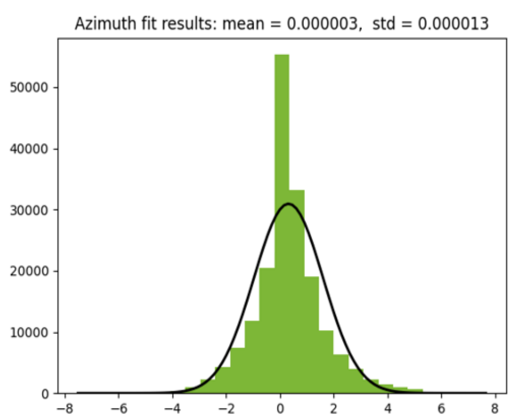

Figure 4: Azimuth Deviation (without noise)

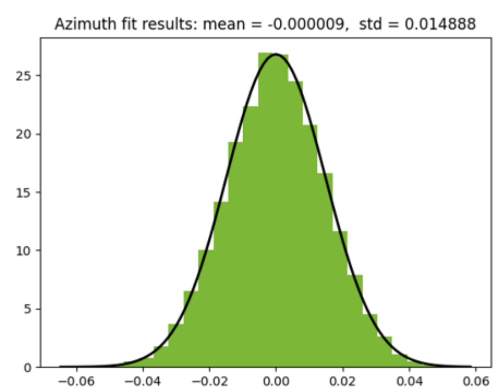

Figure 5: Azimuth Deviation (with noise)

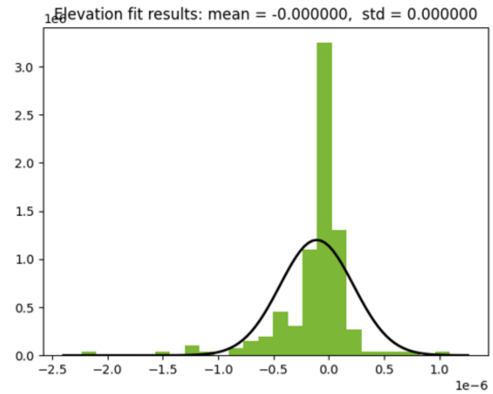

Figure 6: Elevation deviation (without noise)

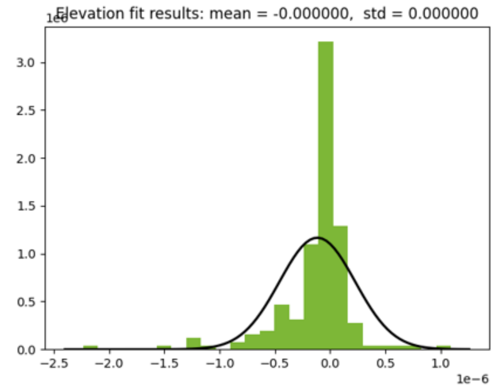

Figure 7: Elevation deviation (with noise)

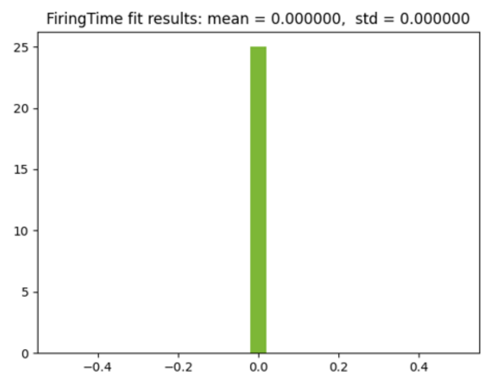

Figure 8: Firing time deviation

As shown in the results, the mean deviation is 0 for elevation, azimuth, and firing time, and the standard deviation is no larger than 0.015 in azimuth. Figures 6 and 7 show a standard deviation of 0 for the elevation noise, thereby producing a deterministic distribution.

These results show the lidar model only differs from the idealized firing pattern up to the configured noise. The model matches the specification well, providing enough confidence to proceed to the next step in the validation process.

Understanding lidar point clouds

Before detailing the results from the simulated and real data collections, it is important to demonstrate some of the basic elements of a lidar point cloud.

Figure 9 shows a visualization of a real world point cloud from this scenario. The important detail to highlight here is the intensity dropoff as a function of range, as well as the azimuth and elevation distribution. Visually observing the result here will help digest the trends in the plots below, as this is observed in the data for both simulation and real world.

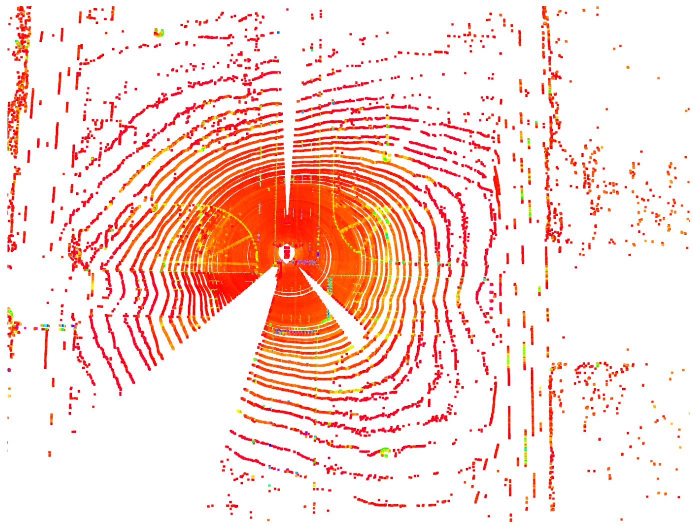

Figure 9: Real-world point cloud showing both intensity and point density dropoff as a function of distance from the sensor

The intensity color pattern in Figure 9 is a gist rainbow color scheme, where red represents weaker returns and pink represents stronger returns. This color scheme is used throughout the lidar blog sections 0, 1, and 2 to encode the intensity returns for visible consumption. Figure 10 shows the color spectrum as it is mapped to intensity values from [0,1] going from left to right.

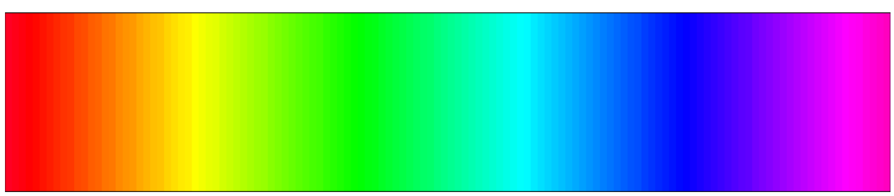

Figure 10: Gist rainbow color pattern for intensity mapping. Intensities range from 0 to 1 map to the colors here from left to right

Metrics and validation details

In this study, each test scenario and range permutation is set to record point cloud data from the real world setup and the simulated digital twin.

We compute three metrics from these datasets: the hit ratio, intensity variation, and range variation. The hit ratio is the number of points returned with respect to a given region of interest on the target panel divided by the theoretical number of expected points for the same region of interest based on the firing pattern of the sensor. Intensity is the raw strength of the returned point, and range is the distance measured from the sensor to the hit point location.

The raw data is then used to compute metric statistics that are central to the comparison and validation of simulation vs. real world. This computation includes means and standard deviations (STD) for all three metrics. The presentation of the raw range and intensity focuses on STD (variations) to depict the precision between simulation and real world.

With STD, the mean intensities and ranges are offset for a number of reasons not related to the accuracy in the simulation predictions. We do not know the exact range to the target, thereby there is an offset between simulation and real data that is not representative of the accuracy of the validation.

Additionally, the rotating lidar introduces firing pattern uncertainties that do not have an exact counterpart in simulated firing pattern uncertainties. Finally, the real world intensity data is dependent upon outside environmental influences and target panel construction details that cannot be entirely modeled in simulation. These influences include macroscopic deformations, not microfacet variations, in the target panel (e.g. creases), fewer than-expected target reflectivities, per sensor intrinsic settings and calibration differences, as well as environmental conditions (e.g. solar glint noise).

There are noticeable differences in measurements that are largely independent of the simulation model and are acceptable differences that come with real world data collection. These factors apply to all the observable measurement errors throughout this study and discussions below.

As a result, the STD is a more interesting metric to observe than the individual mean measurements. It is the reason behind the raw point cloud data analysis emphasizing the intensity and range variations centered at a zero mean.

It is also important to note that the simulation model was held constant throughout all the scenarios and test cases to maintain a constant baseline for comparisons. There were no changes in the simulation model (e.g. lidar intrinsic parameters), just the movement of the target panel to the prescribed ranges and performing measurements.

However, the simulation model is agile enough for “tuning” to augment the output, such that simulation can have accurate comparisons with any set of measurements. That was not the objective here, so the simulation model was held constant throughout all the validation tests.

The results are organized according to the lidar scenario and target panel reflectivity. The first lidar scenario simply has a 5% target reflective panel for each of the prescribed ranges (approximately 7.5m, 12.5m, 22.5m, 52m, and 102m) where a minimum of 200 point clouds are collected for each range. The raw data statistics are presented at only two discrete ranges for conciseness and detail the hit ratio metrics along with the intensity and range STDs.

Finally, summary statistics are presented for all ranges on hit ratio and intensity and discussed to quantify the correlation between simulation and real world data results. The same procedure is then applied to lidar scenario 2, which adds a transparent plexiglass plate in front of the 5%

reflective target panel. The same set of tests and results are performed on a 84% reflective target panel for lidar scenarios 1 and 2. Those details are covered in the appendix.

Lidar scenario 0 - comparable evaluation

Validating simulation data against real world data is always a challenge. Real world environments include many unforeseen factors that produce observable differences when compared to simulation results.

Though we can minimize these differences with careful measurement setup and techniques, it is also advantageous to include multiple real world data collections to provide greater insight in the validation process.

Lidar scenario 0 is a previously conducted experiment, where the results from that data are used to provide additional comparison to what scenarios 1 and 2 perform. It is a simple evaluation scene of real world data with target panels of known reflectivities at an approximate range of 12.5m.

The lidar used here is the same type of rotary sensor that is used in lidar scenarios 1 and 2. Figure 11 depicts the point cloud of this scenario capture. The results from the target panels are displayed in the lidar scenario 1 section, so that all results can be observed in a single cohesive diagram.

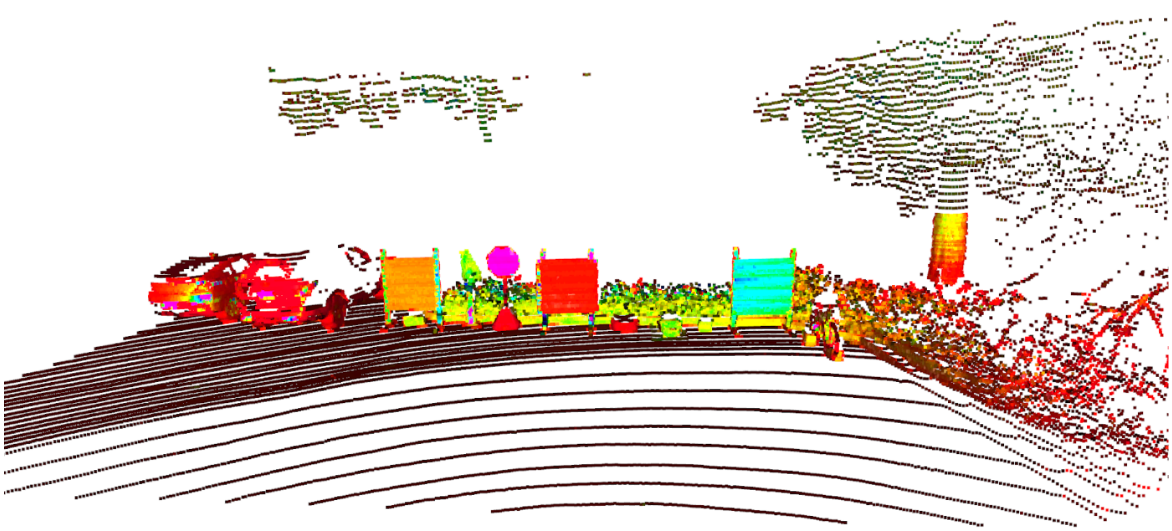

Figure 11: 12% Target panel (left), 5% target panel (middle), 84% target panel (right)

Lidar scenario 1 - single opaque target panel results

The next scenario simulates real world experiments as a function of range. Figure 12 depicts the range-based measurements discussed in the metrics and validation details section.

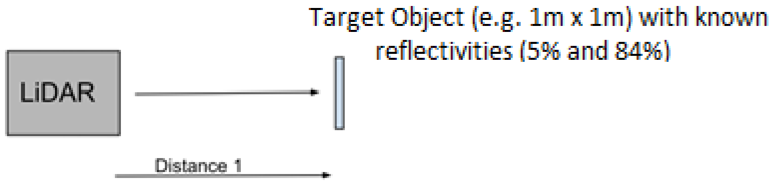

Figure 12: Lidar reflective panels experiment

An example of this scenario is shown in the following images. Figure 13 displays an image from the actual experiment using the front DRIVE Hyperion camera to capture the upright 5% reflective panel, at a distance of 22.5m for the ego vehicle.

We then create the same scenario in DRIVE Sim as a digital twin environment, shown In Figure 14. Though the surrounding environment is not exactly as in the real world data, the details for the target panel match between the real world and simulation.

Figure 13: Reflective panel with 5% reflectiveness at 20m away from ego vehicle

Figure 14: Reconstruction of the scene in DRIVE Sim

We repeated the same procedure for all of the reflective panels, at every range (10 experiments in total). For each case, comparisons were performed on the percentage of hits on target, range variation, and intensity variation. The 5% reflective target is presented here and the 84% reflective target is presented in the appendix, given the commonality of results and simplicity.

5% Reflectivity target panel validation

This section details the analysis and validation of simulation versus real world data for the 5% reflective target panel.

Real world data

The figures of 15a and b show the point clouds captured from the real world recorded data for lidar scenario 1 and 0. Only the portion of the data spanning the 5% reflective target panel is extracted and used for the comparison to sim data.

Lidar Scenario 0 (real) Lidar Scenario 1 (real)

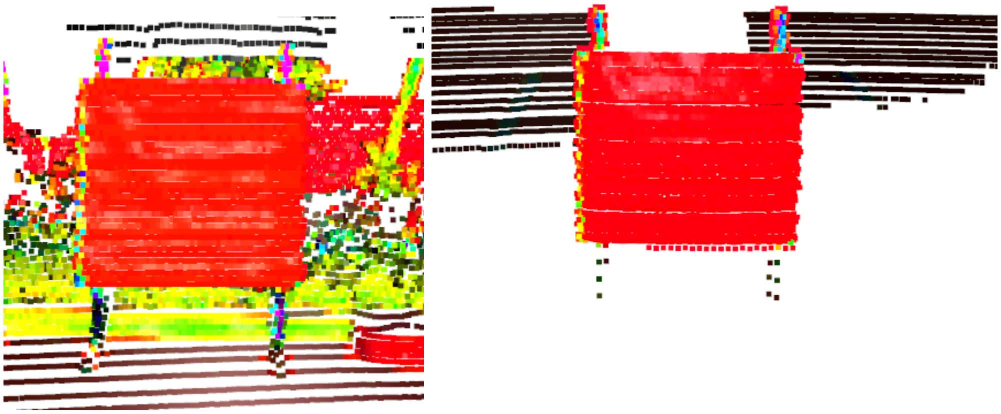

Figure 15(a) Figure 15(b)
Real world point cloud images from recorded data of 5% reflective target

Simulated data

The point cloud in Figure 16 is captured from the simulated recorded data for lidar scenario 1. Just as in the real world data collection, we extract only the portion of the data spanning the 5% reflective target panel and use it for the comparison to real world data.

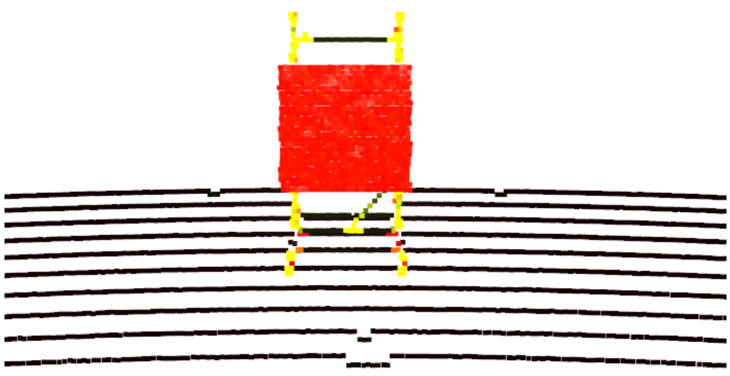

Figure 16: Simulation replica point cloud of real world scene of 5% reflective target

Figures 17a-c and 18a-c detail the raw real world and simulation data for lidar scenario 1, covering hit ratios, intensity and object distance variations at two discrete ranges upon a 5% reflective target panel.

22.5m (real)

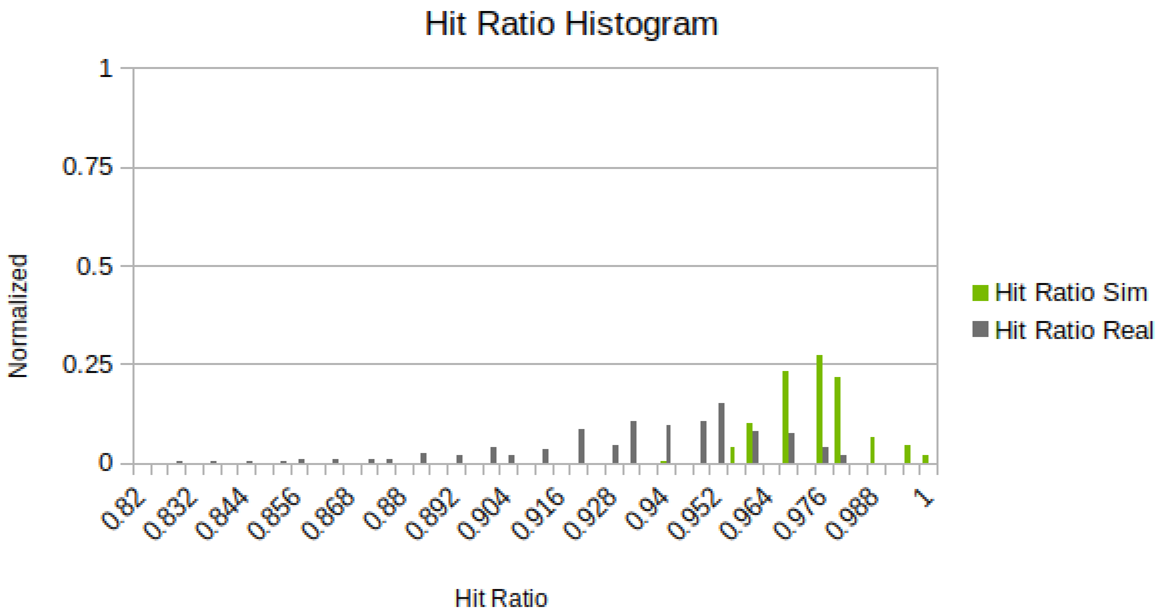

Figure 17(a): Hit ratio of simulated and real data at 22.5m range
Real mean - 0.92, Real STD - 0.06
Simulation mean - 0.992, Simulation STD - 0.01

The hit ratio data for this range, including all point clouds, correlates to the idealized number of returns for a given region of interest. The majority of the point clouds exhibit a greater than 90% ratio.

The real data exhibits greater deviation due to the uncertainties introduced by the rotating lidar and variations associated when the sensor captured the point cloud for the region of interest. Deviations are further exacerbated by unknowns associated with real sensor calibration elements.

The simulation environment does simulate firing uncertainties for a spinning lidar, but we can expect some differences without knowing the full extent of the real lidar intrinsics.

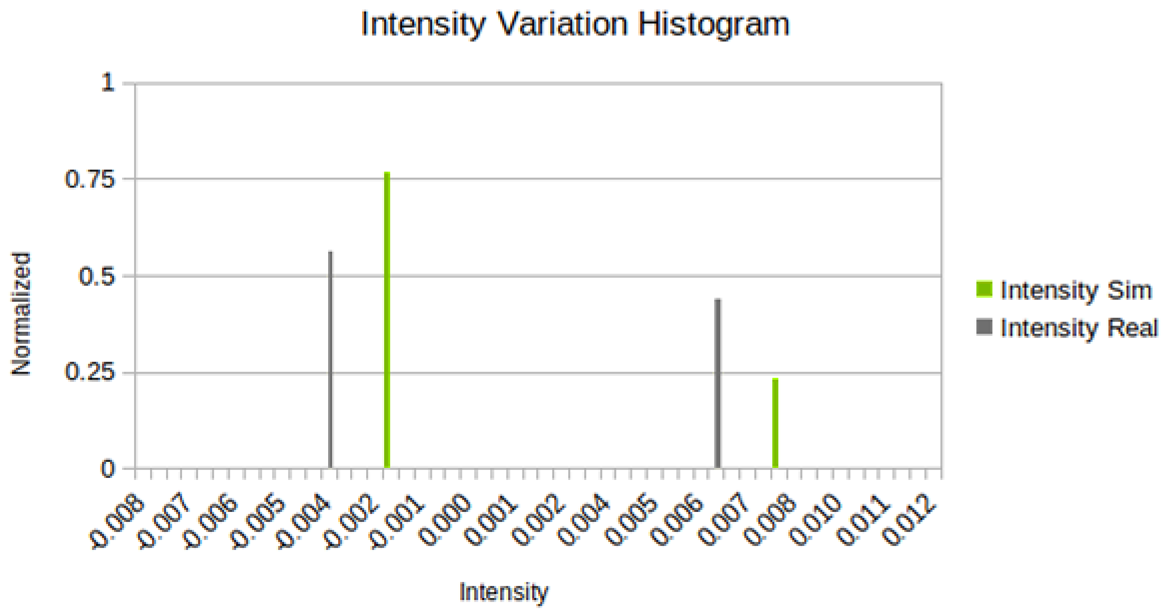

Figure 17(b): Intensity variation of simulated and real data at 22.5m range.
Real STD - 0.005; Simulation STD - 0.004

For the raw intensity variation data in Figure 17b, both simulated and real data show similar variations in the returned intensities relative to their respective means. Most of the measurements fall into two bins, based on the constant reflective panel at normal incidence.

Despite the small number of unique bins, the standard deviation is small, showing convergence of the measurements with respect to the mean.

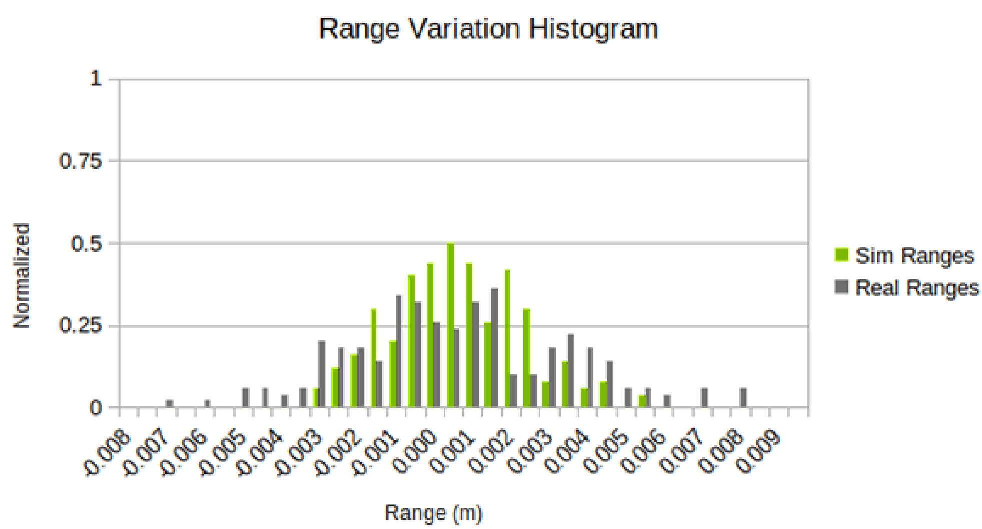

Figure 17(c): Range variation of simulated and real data at 22.5m range.
Real STD - 0.003; Simulation STD - 0.002

Similar to the intensity variations, the range variation agrees with simulation in relation to their respective mean ranges.

The same results are presented in Figures 18a-c, but now at a range of approximately 101.5m.

101.5m range

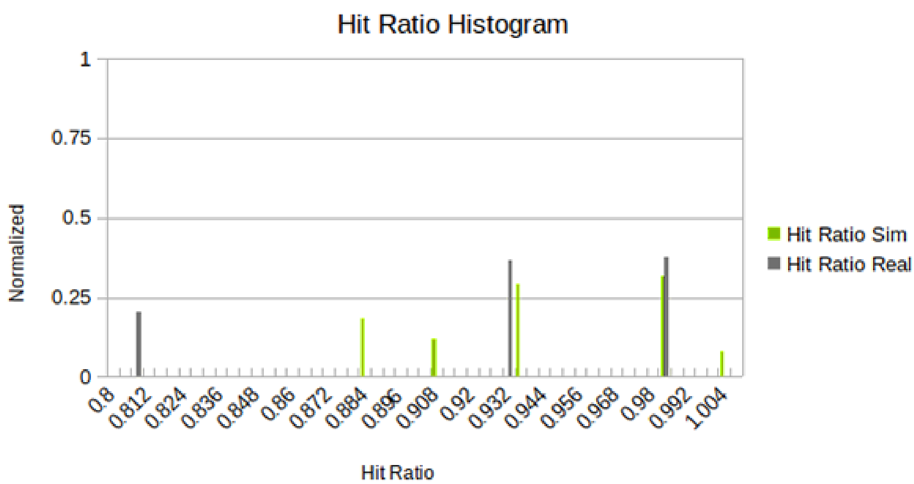

Figure 18(a): Hit ratio of simulated and real data at 101.5m range.
Real mean - 0.933, Real STD - 0.22
Simulation mean - 0.95, Simulation STD - 0.087

Larger distances produce greater standard deviations in the observed hit ratios. This is an expected trend given the smaller point density on the target region of interest and the firing pattern uncertainty. Mean ratios are still greater than 93%.

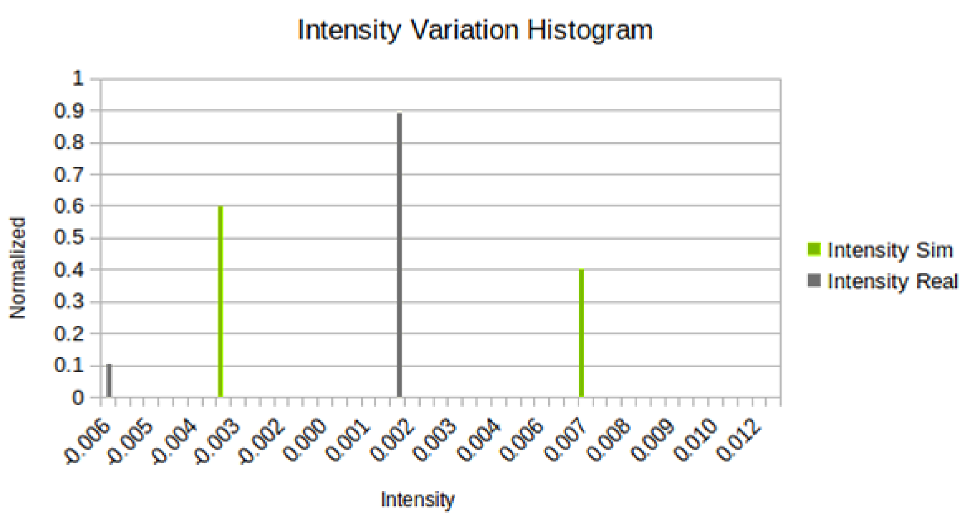

Figure 18(b) - Intensity variation of simulated and real data at 101.5m range.
Real STD - 0.003; Simulation STD - 0.005

The standard deviation of the intensities remains small relative to their respective means.

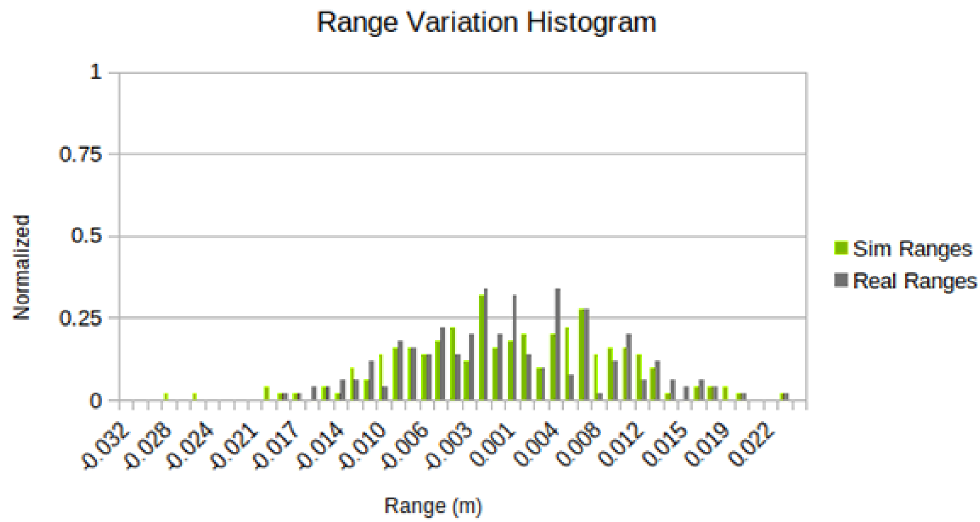

Figure 18(c): Range variation of simulated and real data at 101.5m range.
Real STD - 0.008; Simulation STD - 0.009

Similar to the data in the intensity variation data, the range data shows very similar variations between simulated and real data relative to their respective means.

Tables 1a and 1b in the data table appendix detail the summary statistics for mean and STD of the hit ratios, as well as the intensities for the real and simulated datasets. This data summary is used in the following diagrams for the validation study.

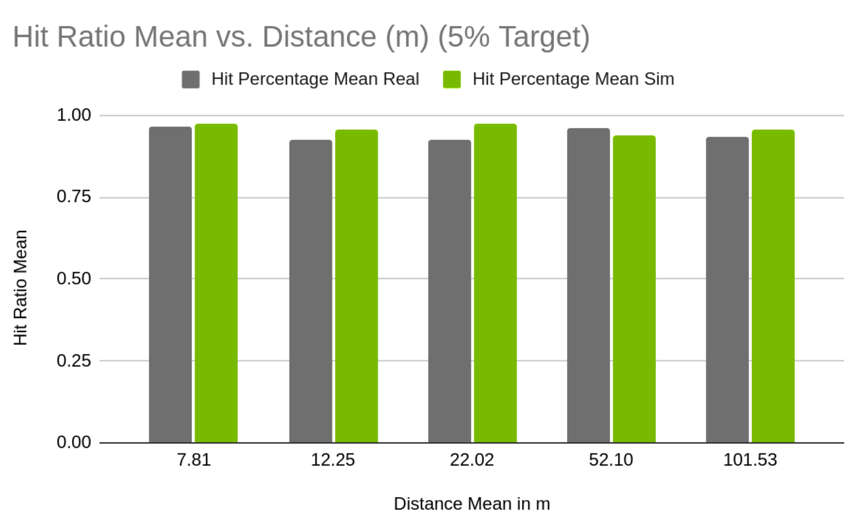

Figure 19 - 5% reflective panel hit ratios as a function of range.

The data in Figure 19 shows correspondence between the firing pattern hit point coverages for both simulated and real data. All ranges show a relative error of less than 5% comparing real and simulated data. Some larger variances in the real world data are apparent and can be attributed to the intrinsic variations in the rotary lidar.

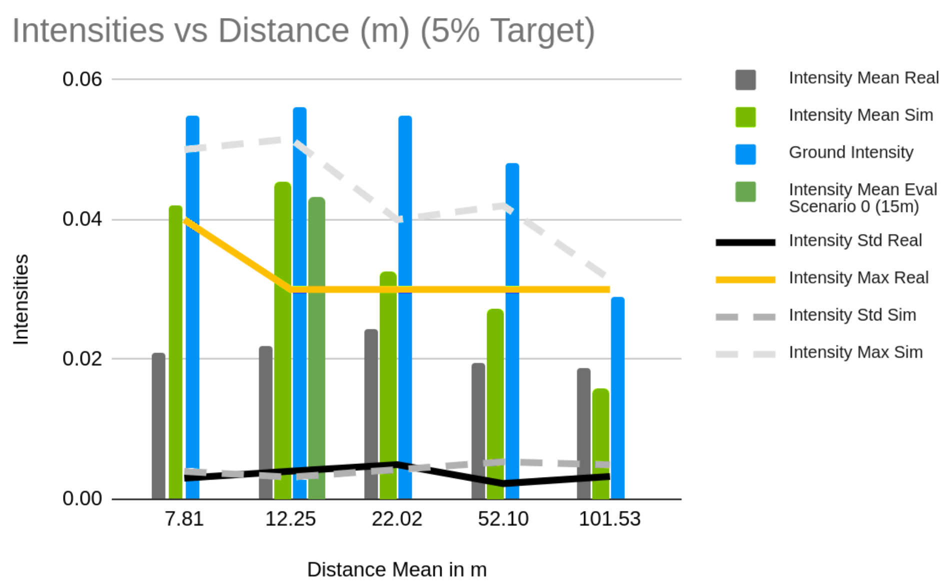

Figure 20: 5% reflective panel intensities as a function of range.

The simulation vs. real world intensity range dependence plots show intensity mean, max and STD. In addition, the mean ground intensity and intensity from the 5% reflective panel from lidar scenario 0 are all shown here for comparison.

Discussion of scenario 0 and 1 results

Overall, we observe similar intensity and range-based trends between simulated and real data. This result can only be achieved through a rigorous simulation model that incorporates all aspects of the physical phenomena from lidar source to detection.

This process includes sensor material reflectance and transmittance modeling through a bidirectional scattering distribution function (BSDF), wave propagation and radiometry, detector processing converting photons to electronic counts, gain functions, and noise variances to model the intensity predictions from the sensor.

The data shows differences in the standard deviations are less than a factor of 2 for all measurements. Despite the differences shown in mean intensities between simulated and real data, the comparative results and trends are still accurate.

The simulation results show a near constant mean intensity below the idealized 5% reflectivity, and then a monotonic downward trend as range increases. We expect this overall trend since the intensity is not entirely representative of the panel reflectivity.

Losses as a function of range will reduce intensities despite the constant reflectivity of the panel. Simulation predicts such variations—as evident in the non-zero standard deviation—but it also shows the downward trend due to aspects of wave propagation and detector sensing of the returned beams. We observe some downward trends in real data as well, but it is more apparent in sim.

We collected and analyzed the ground intensity from real world data as a function of range to provide a different view from the reflectivity target panels.

As we move further from the target, we expect intensity returns to decrease. The downward trends were more observable in the ground mean intensity, showing a range-dependent intensity falloff similar to the simulation results.

The standard deviations for both simulated and real data show similar trends and depict upward trends as a function of range. Real world data shows significant variations for longer ranges, which is an attribute of the aforementioned uncertainties in the environment, sensor, and target panel. These real world variances can account for the observable trends in the mean intensities.

Though simulation data accounts for some target surface degradation by incorporating statistically varying surface normal boundaries, it does not account for all aspects of a real world data capture. Previous discussions on the target panel variability and composition is largely what is producing the significantly lower mean intensities compared to simulation.

The plot in Figure 20 shows the result from the mean intensity from lidar scenario 0. In this test case, the intensity is approximately double that of the real world lidar scenario 1 data for similar range, sensor type, and target reflectivity.

The relative error for mean intensity at the 12m range for lidar scenario 0 compared to simulation is 7%. This result confirms the differences in the target panel composition, sensor type, and measurement environment can produce significantly different results but are still commensurate with simulation.

Lidar scenario 2 - Plexiglass before target panel

For the second set of experiments, we conducted the same tests as lidar scenario 1, but now with a mounted transmissive plexiglass plate, as shown in Figures 21a and b.

The plate is placed at a defined distance of 1m in front of the reflective target panel between the ego vehicle and the reflective target panel. The goal is to evaluate DRIVE Sim’s ability to mimic multiple returns. We also want to measure the transmissive capabilities of the simulated material in addition to the lidar simulation model. We used the same metrics for this experiment as in lidar scenario 1.

Figure 21a: Plexiglass mounted on a 1m tall frame.

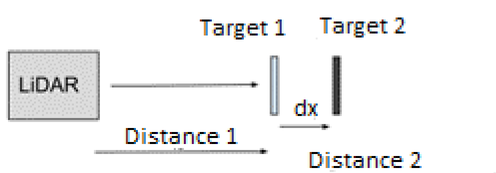

Figure 21b: Lidar reflective panels experiment with transparent plexiglass obstacle.
Target1 - transmissive plexiglass object with known material properties for transmission and reflection
Target2 - non transmissive object with known reflectivity
dx - known constant distance between the two targets (e.g. 1m)
Distances - fixed distances for validation study (approximately 7.5m, 12.5m, 24m, 52m, and 103m)

Plexiglass before 5% target panel

Figure 22 shows the point cloud captured from the real world recorded data for lidar scenario 2. We extracted only the portion of the data spanning the 5% reflective target panel and used it for the comparison to simulation data.

We omitted any returns from the transparent plexiglass panel from the analysis to maintain consistency in comparing simulated and real world data from known reflective target panels.

Real lidar scenario 2

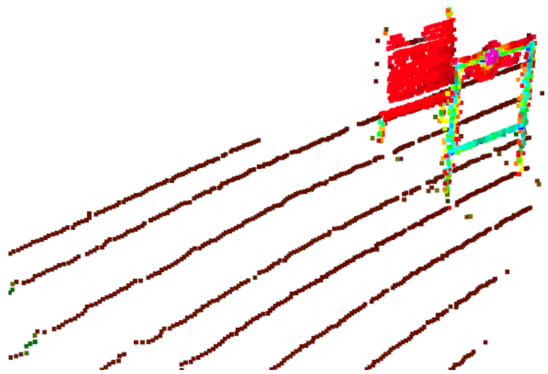

Figure 22: Real world point cloud images from recorded data of 5% reflective target with plexiglass target in front.

In Figure 23, the image shows the point cloud captured from the simulated recorded data for lidar scenario 2 evaluation.

Sim Lidar Scenario 2

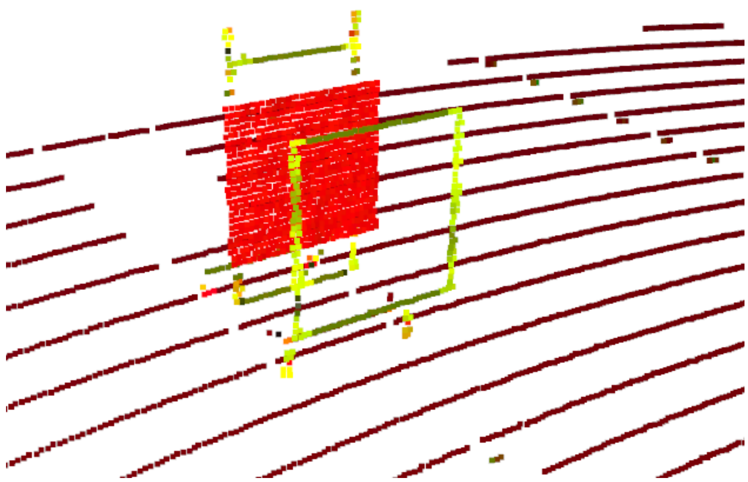

Figure 23: Simulation replica point cloud of real world scene of 5% reflective target with plexiglass target in front.

Figures 24a-c and 25a-c detail the same analysis of the raw point cloud data metrics (hit ratios, intensity, and range variations) as in Figures 17a-c and 18a-c. The only difference in this scenario is the transparent plexiglass panel in front of the reflective target panel.

An overall consequence of adding the panel, as the data will show below, is the reduced intensities and hit ratios due to the lidar beams' round trip transmission through the plexiglass.

24m range

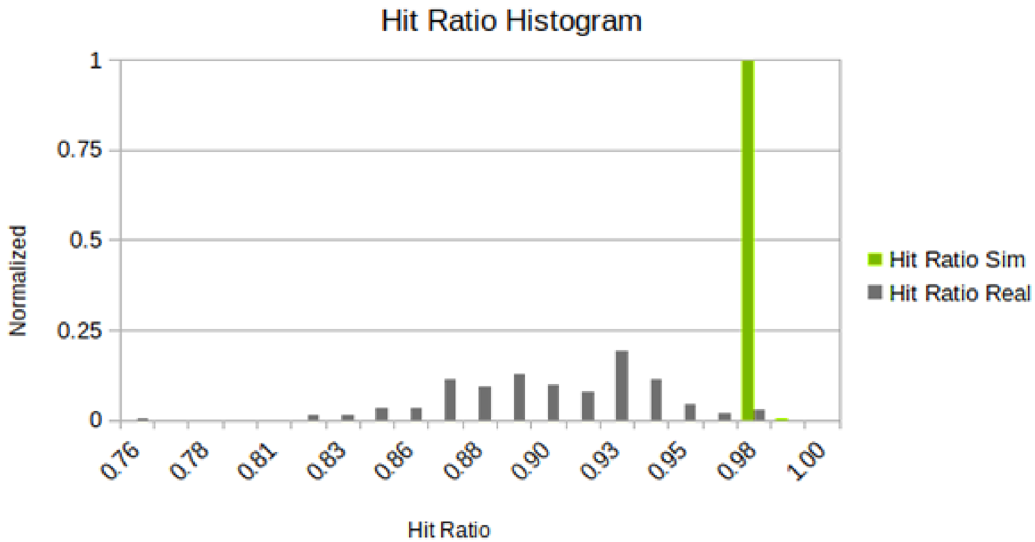

Figure 24(a): Hit ratio of simulated and real data at 24m range.
Real mean - 0.912, Real STD - 0.072
Simulation mean - 0.97, Sim STD - 0.001

The plexiglass panel introduced significant variations in the hit ratios. This result is a consequence of the transmission through the surface as well as solar reflection noise in the real world data, in addition to the uncertainties of the rotating lidar. Solar-based noise is not yet modeled in DRIVE Sim materials and is the primary reason behind the variances observable in the simulation vs. real world comparison.

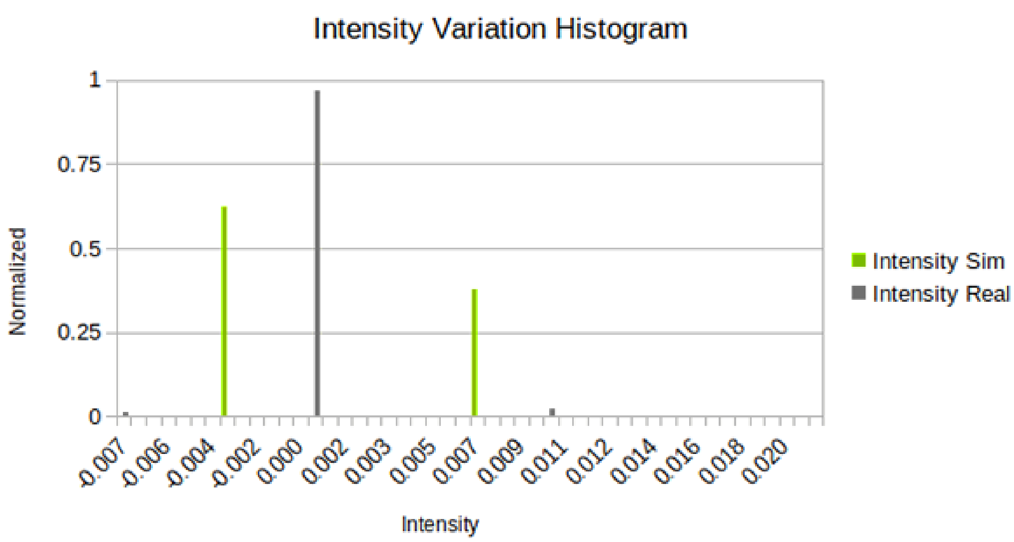

Figure 24(b): Intensity variation of simulated and real data at 24m range.
Real STD - 0.002; Simulation STD - 0.0049

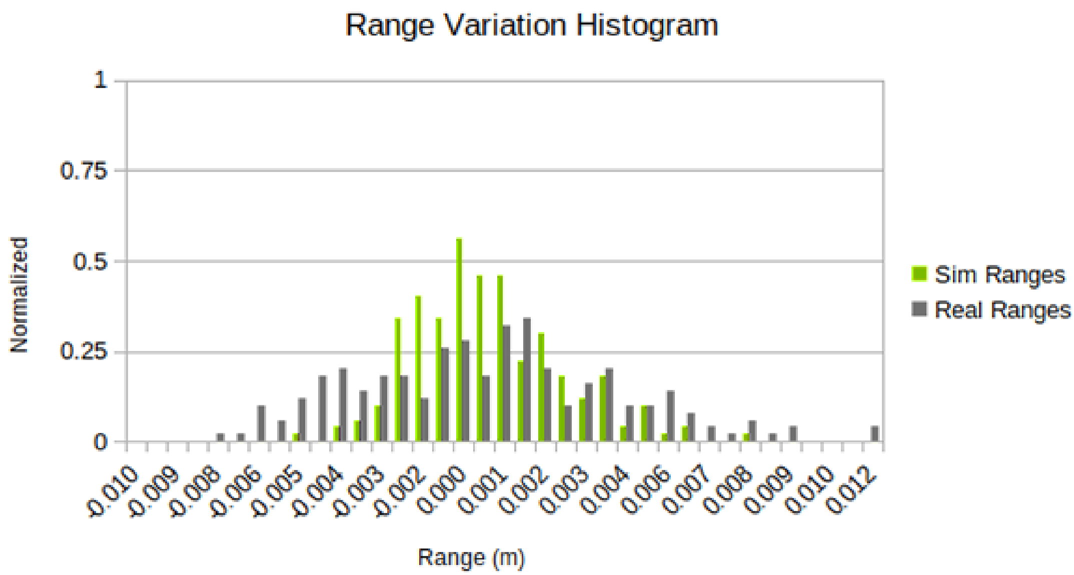

Figure 24(c): Range of simulated and real data at 24m range.
Real STD - 0.004; Simulation STD - 0.002

Figures 25a-c are the same metrics, but at a range of approximately 103m.

103m range

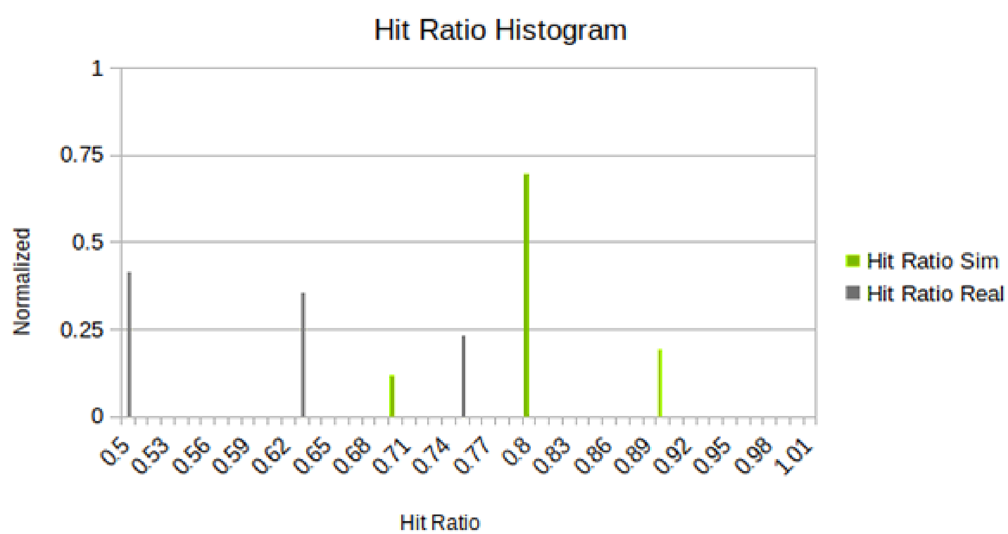

Figure 25(a): Hit ratio of simulated and real data at 103m range.
Real Mean - 0.588, Real STD - 0.241
Simulation Mean - 0.779, Simulation STD - 0.038

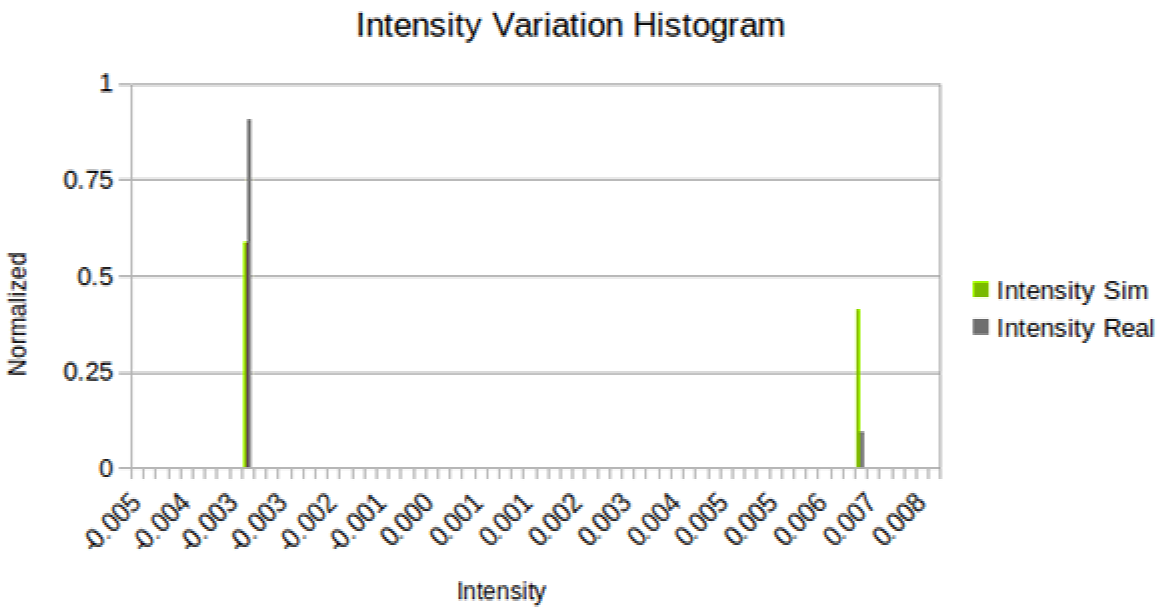

Figure 25(b): Intensity of simulated and real data at 103m range.
Real STD - 0.003; Simulation STD - 0.005

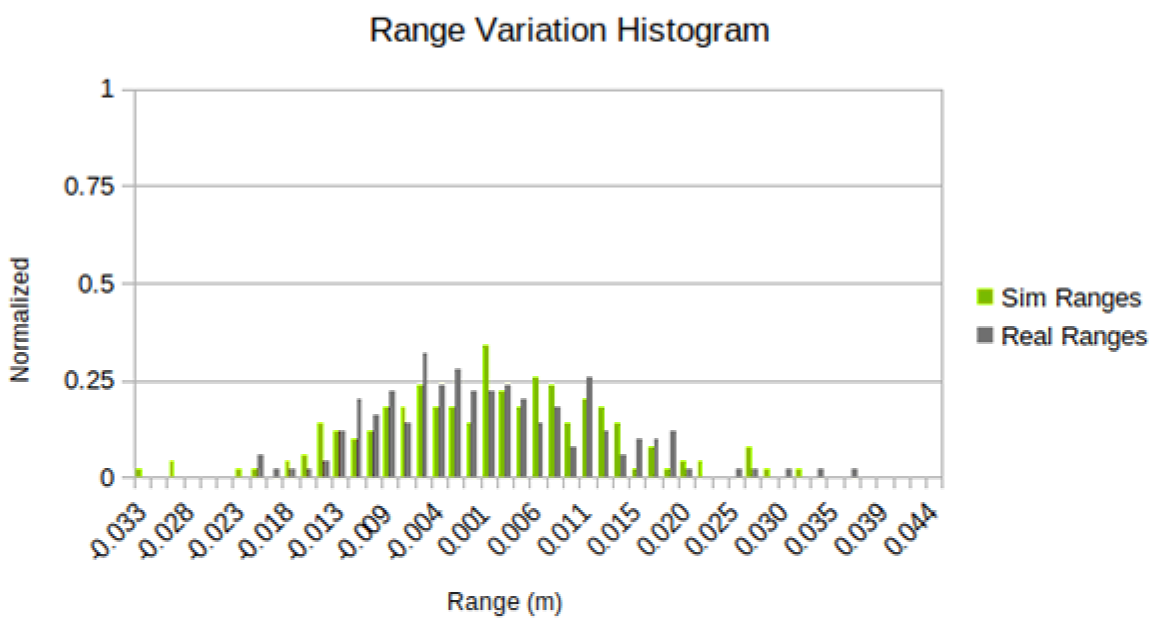

Figure 25(c): Range of simulated and real data at 103m range.
Real STD - 0.01; Sim STD - 0.011

Figures 25a-c show how range degrades the measurements. All three metrics show greater variations, with the hit ratio depicting smaller means. This result is a consequence of the longer range having fewer points on the region of interest, and the aforementioned effects of firing pattern uncertainty and solar reflection noise from the plexiglass surface.

Thus, the real data shows slightly larger variance trends than simulation. Overall, the results show that simulation trends are following expected real world trends.

Tables 2a and b in the data table appendix depict the real world and simulated data statistics for the set ranges identified for lidar scenario 2 on the 5% reflective target panel with a transparent plexiglass panel in front of the target. The data presented here is visualized in the following figures.

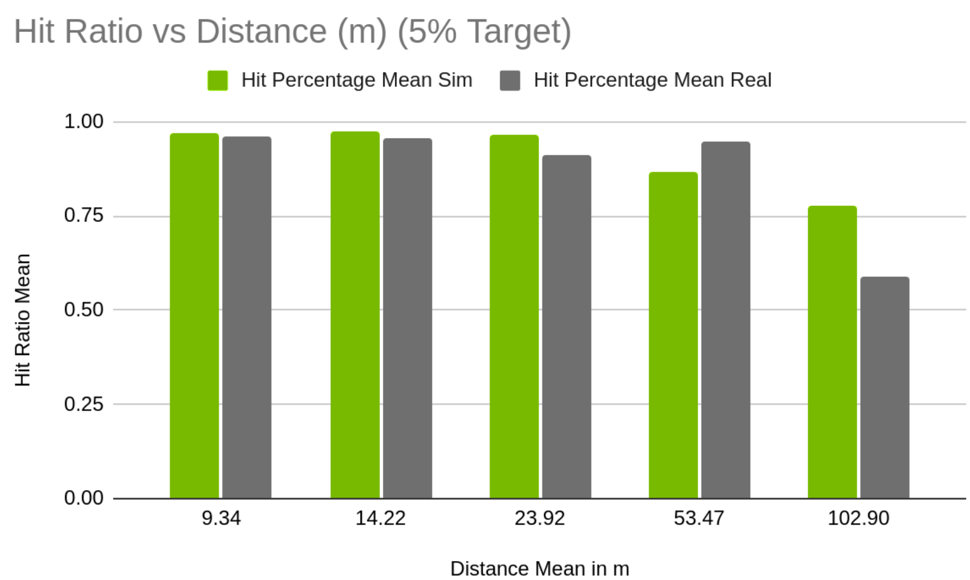

Figure 26: 5% reflective panel hit ratios as a function of range with plexiglass panel.

The hit ratios show the level of consensus between simulation and the real world. The data depicts that the ratios are similar for the 5% reflective target in lidar scenario 1, but exhibit additional dropped points at longer ranges.

We expect this pattern now that an additional plate, albeit transparent, is in front of the reflective target. The results highlight that, in addition to the fidelity of the data generated by simulation for a rotating lidar, the simulation model can properly model transparent obstacles that can cause degradations from these layered panel interactions.

The small divergence in the comparison at longer ranges shows the previously mentioned uncertainties are exacerbated by the addition of a reflective panel in front of the target panel. The relative error between simulation and real world is less than 10% for ranges shorter than 100m. Ranges greater than 100m demonstrate significant solar glint noise that produces greater relative errors.

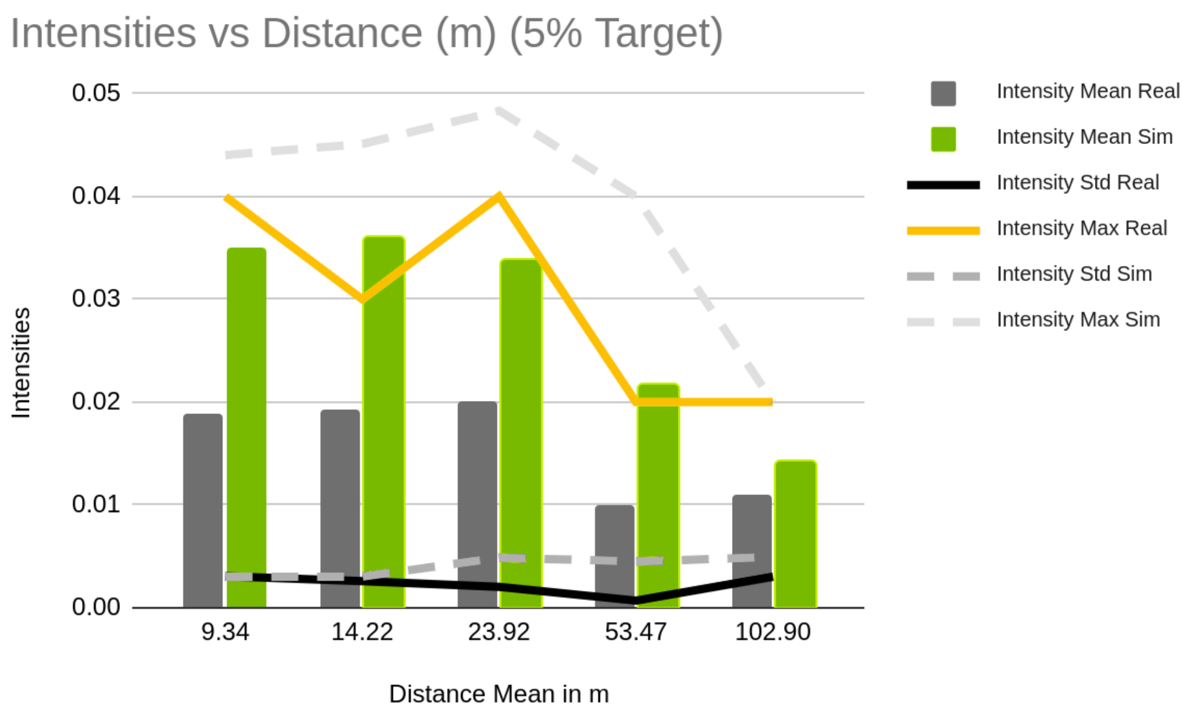

Figure 27: 5% reflective panel intensities as a function of range with plexiglass panel.

The details presented in Figure 27 include intensity mean, max, and STD for real world and simulation. The data is similar to lidar scenario 1, but does not include the ground intensity mean nor the mean intensity from lidar scenario 0. Rather, the focus is on data for the 5% reflective target that is obstructed by the clear plexiglass panel.

Discussion of scenario 2 results

Overall, there is observable correlation between simulation and the real world. There are apparent differences, but overall, similar data trends and amplitudes are comparable in both the simulated and real data.

The intensity precision is within a factor of 2. Just as in the metrics and validation section, previous discussions with the target panel and environmental uncertainties account for the observable differences in the intensities. However, the data is still showing credible similarities between simulation and the real world.

An overall comparison of the results with lidar scenario 1 show similar trends for all ranges, including the flat mean for shorter ranges and then a monotonic drop for longer ranges in simulation.

There are also increases in the STD as a function of increasing range. Real world mean intensities are consistently lower than simulation, as well as a falloff in the real world intensity followed by an uptick at a longer range.

These trends are discussed in lidar scenario 1 and apply in this test case. With all the uncertainties aside, the trends and STDs are comparable, showing that the simulation model is producing data similar to what real world physics produces.

In regard to the handling of material boundary and transmission factors, the lidar simulation model accounts for the full path, forward and back propagation. The simulation material model handles transmission and losses in both directions through the plexiglass material to produce a cumulative drop in the lidar signal returns.

The modeling of these material interactions gives a more accurate simulation scene representation and generates data more comparable to real world data. This fact is most evident by observing the lower mean and max intensities across all data ranges when compared to the 5% target panel in lidar scenario 1.

There are also key differences in comparison to lidar scenario 1. The decreases in the mean intensities compared to scenario 1 is a consequence of the plexiglass panel in front of the reflective target panel.

In addition, the STDs are larger in lidar scenario 2 than in lidar scenario 1. Variations in light path deviations from the first surface and then again from the second surface affect the lidar returns recorded by the detector. Two sources of light interactions are producing variances compared to one in the first scenario. The STDs for both simulation and the real world reflect this change. These factors contribute to the larger observable relative error compared to lidar scenario 1.

Conclusion

This report shows a quantitative study of a lidar sensor evaluating three core metrics of known reflective target panels as a function of range.

In these scenarios, five ranges are evaluated for two different reflective target panels (5% and 84%). The data for the 84% reflective panel can be found in the appendix. The two scenarios conducted involved taking each target at a prescribed distance and measuring metrics like hit ratio (the measured number of points divided by the theoretical number of expected points for a given region of interest), intensity, and range.

The second scenario was the same as the first, but with a transparent target panel in front of the known reflective target panel under test. This experiment is done for simulation and real world, and we collected and evaluated the data for accuracy and precision.

The results from this validation study show that simulation and real world data display similar means and trends as a function of range. The standard deviations show the expected precision trends for all three metrics and correlation between simulation and real world data. Increases in range and the addition of a transmissive panel in lidar scenario 2 depict larger STDs for all metrics.

Though there are observable differences with the mean metrics, reasonable similarities are apparent between simulated and real datasets, demonstrating the accuracy of simulation. In this comparison, the errors that are part of the real world data include uncertainties from the rotating lidar, target panel deformations affecting reflectivity measurements, less-than-expected reflectivities from the target panels, sensor type intrinsic details and calibration differences, and non-idealized experimental conditions (outside influences such as the sun).

Though these differences existed, other data metrics with ground plane comparisons and other target panel data from a different experiment with the same lidar type show more accurate data trends found in the simulation model.

The results demonstrate the depth of the lidar simulation model's ability to model firing patterns, wave propagation, radiometry, BSDF reflection/transmission, detector efficiencies and photon transfer, and noise, characterizing what can be expected in real world lidar scenarios.

In addition, the lidar simulation model can be tuned with configuration parameters to better follow real world data trends, helping combat outside factors. For comparability reasons, the model aspects and lidar intrinsics are held constant throughout this evaluation.

Appendix

Lidar Scenario 1 - 84% Reflectivity Target Panel

Real world Data

Figures 28a and b below show the point cloud captures from the real world recorded data for lidar scenario 1 and 0. The target panel reflectivity is ideally 84% in these measurements.

Lidar Scenario 0 (real)

Lidar Scenario 1 (real)

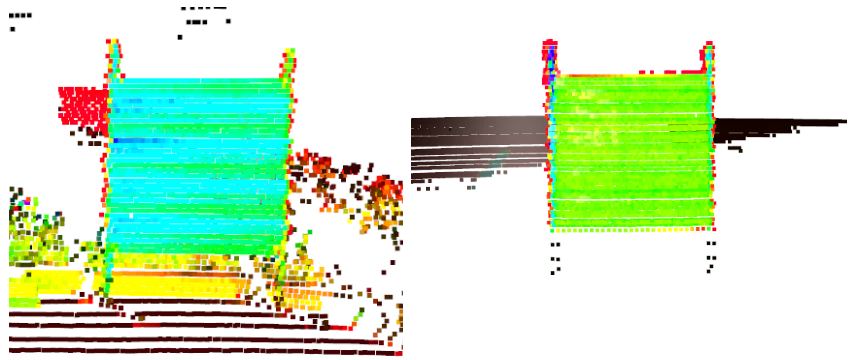

Figure 28(a) Figure 28(b)
Real world point cloud images from recorded data of 84% reflective target

Simulated Data

Figure 29 below shows the point cloud captured from the sim recorded data for lidar scenario 1 on a simulated 84% reflective target panel.

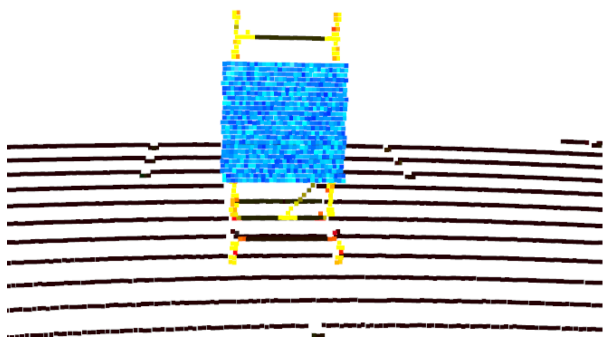

Figure 29: Sim replica point cloud of real world scene of 84% reflective target

The sequence of figures in 30a-c and 31a-c show the 3 raw data metrics as in the previous LiDAR scenario 1 case now with the 84% reflective panel.

22.5m Range

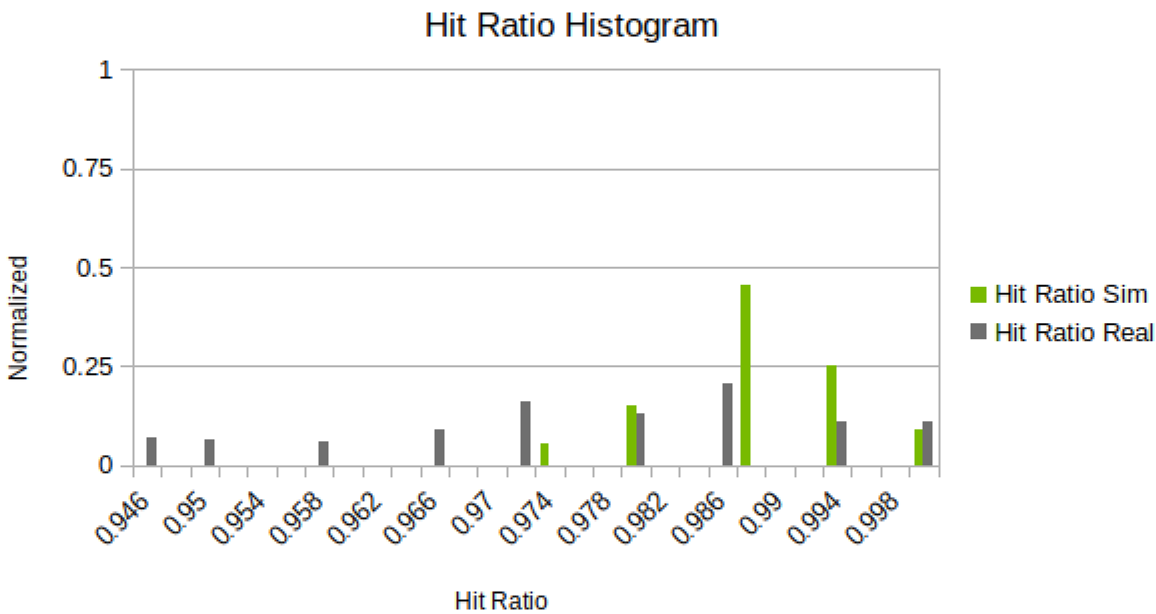

Figure 30(a): Hit ratio of Sim and Real at 22.5m range.
Real Mean - 0.975, Real STD - 0.035
Sim Mean - 0.958, Sim STD - 0.006

The mean hit ratios show a high degree of correspondence to theoretical point coverage of the region of interest. The real data exhibits greater deviation, as seen in the other test cases, due to the uncertainties introduced by the rotating lidar. However, the mean ratios are above 95% for both sim and real.

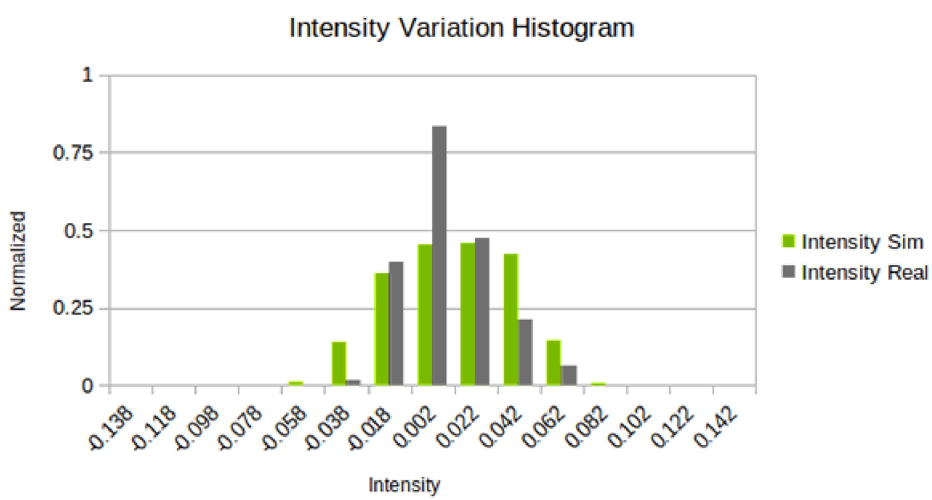

Figure 30(b): Intensity Variation of Sim and Real at 22.5m range.
Real STD - 0.02; Sim STD - 0.027

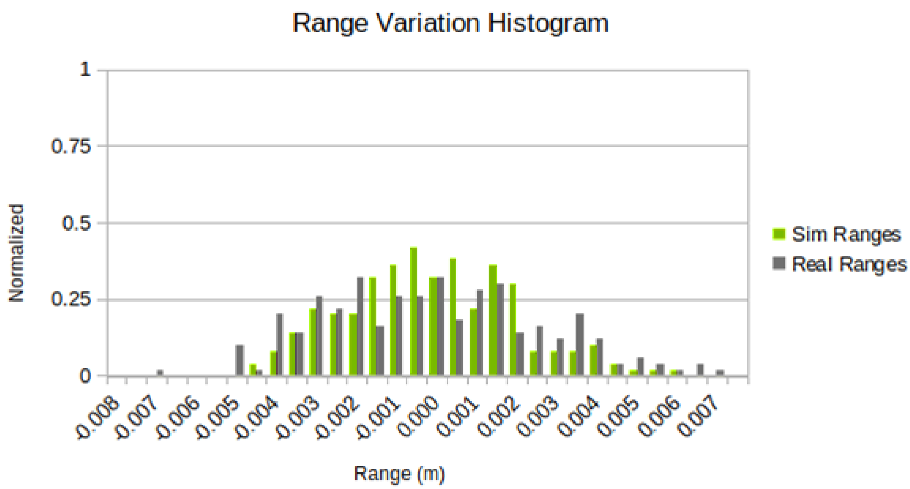

Figure 30(c): Range Variation of Sim and Real at 22.5m range.
Real STD - 0.003; Sim STD - 0.002

Both figures 30b and c show good variation correlation between sim and real. The standard deviations match up very well for intensities and range.

The same results are presented in figures 31a-c below but now at a range of approximately 101.5m

101.5m Range

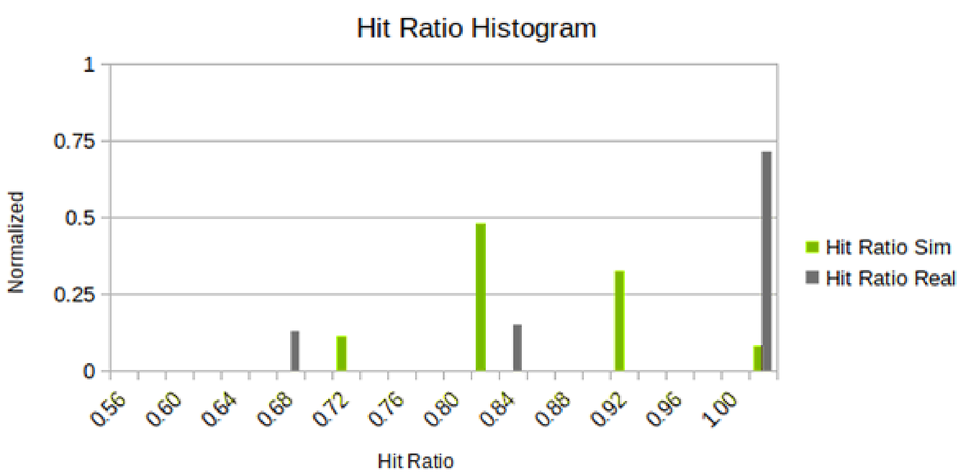

Figure 31(a): Hit ratio of Sim and Real at 101.5m range.
Real Mean - 0.926, Real STD - 0.055
Sim Mean - 0.86, Sim STD - 0.047

Figure 31a shows that hit ratios suffer as a function of range. The uncertainty is greater in these cases as the region of interest is covered by fewer LiDAR returns. The standard deviations confirm this fact.

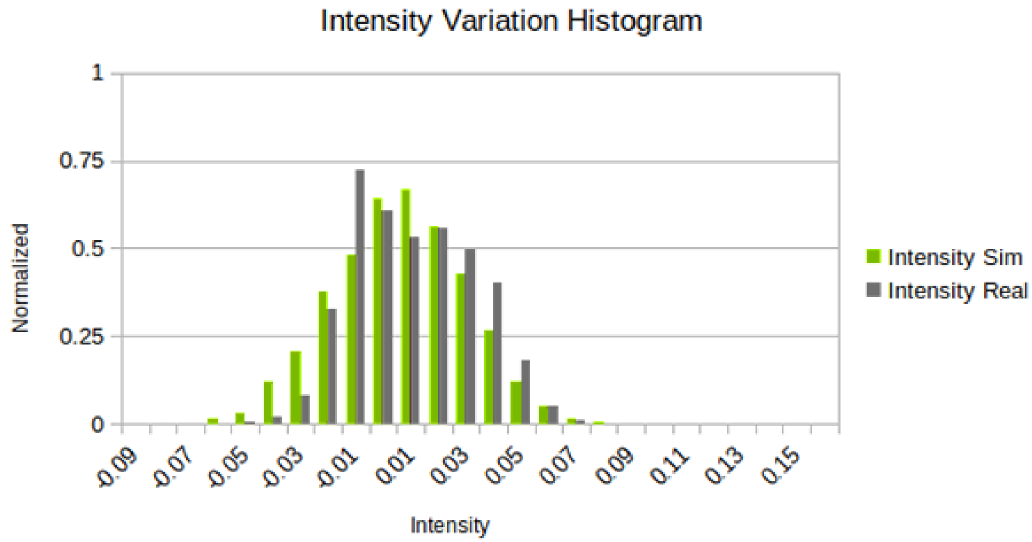

Figure 31(b): Intensity Variation of Sim and Real at 101.5m range.
Real STD - 0.0245; Sim STD - 0.024

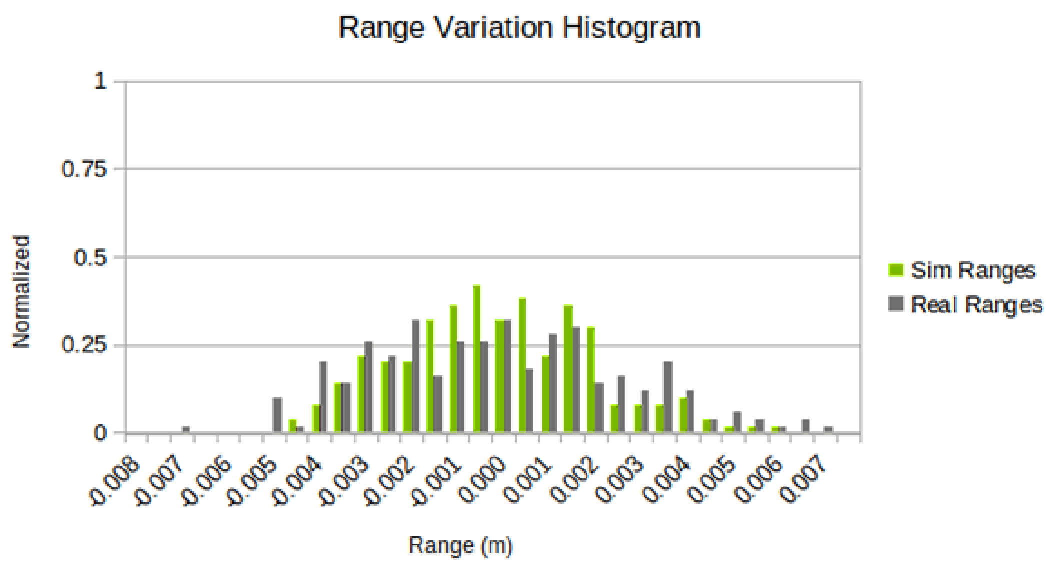

Figure 31(c): Range Variation of Sim and Real at 101.5m range.
Real STD - 0.006; Sim STD - 0.01

The tables 3a and b in the data table appendix detail the real world and sim data statistics for the set ranges identified for lidar Scenario 1 on the 84% reflective target panel. This data presented here is visualized in the following figures below.

The plots in figures 32 and 33 correspond to the data tabulated above.

Real vs. Sim

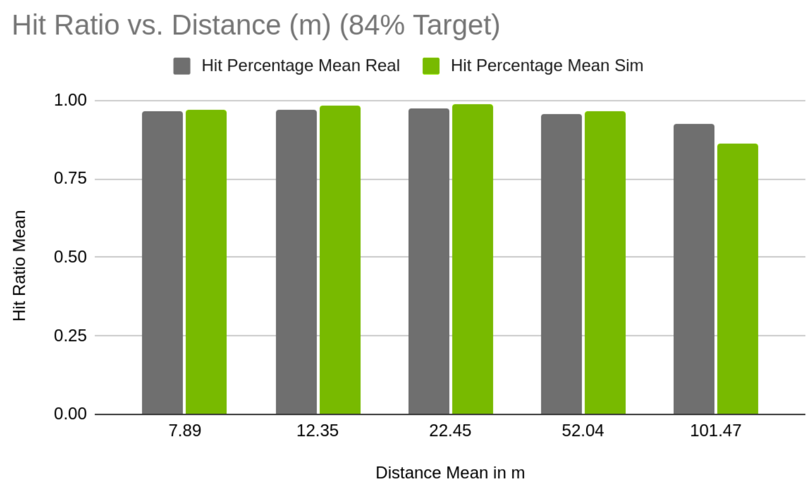

Figure 32: 84% reflective panel hit ratios as a function of range.

The hit ratios between real and sim data are quite accurate. All ranges below 100m exhibit a relative error less than 2% and 7% at the 101.5m range. This demonstrates that the theoretical hit point coverage on the sim and real target panels for the 84% reflective panel are in very good correlation. This highlights the fidelity of sim generated firing patterns and how range dependence can affect the number of returned points affecting the hit ratio.

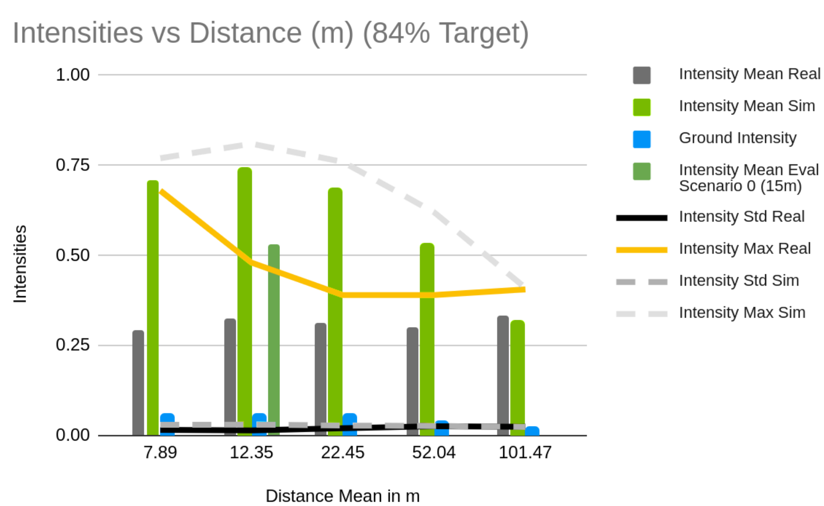

Figure 33: Intensity data for 84% reflective panel as a function of range.

The same data trends and conclusions detailed in the 5% target panel of lidar scenario 1 can be concluded here. Overall, good correlation is observable across the sim and real target panel data.The standard deviations show similar scales and trends. Mean intensities do exhibit differences but can be attributed to the variations of the aforementioned uncertainties in measurements from the target panel, sensor type, and environment. Just as in the case of the 5% target panel, the differences in real world data from lidar scenario 0 and lidar scenario 1 show significant differences confirming the uncertainties in real world intensity measurements.

Lidar Scenario 2

Plexiglass before 84% Target Panel (real, with reflected sun noise)

The figure below shows the point cloud captured from the real world recorded data for lidar scenario 2. The setup is the same as the lidar scenario 2 for the 5% reflective panel but now with an 84% reflective panel. The raw data is collected from both sim and real environments for comparison.

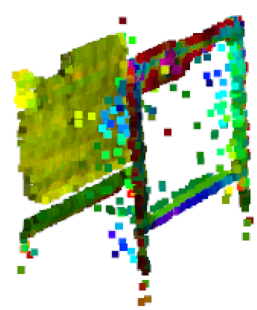

Figure 34: Real world point cloud images from recorded data of 84% reflective target with plexiglass target in front. The spurious points are noise generated by solar glint off the plexiglass panel.

Plexiglass before 84% Target Panel (Sim)

The figure below shows the point cloud captured from the sim recorded data for lidar scenario 2 with an 84% reflective target panel behind a transparent plexiglass panel.

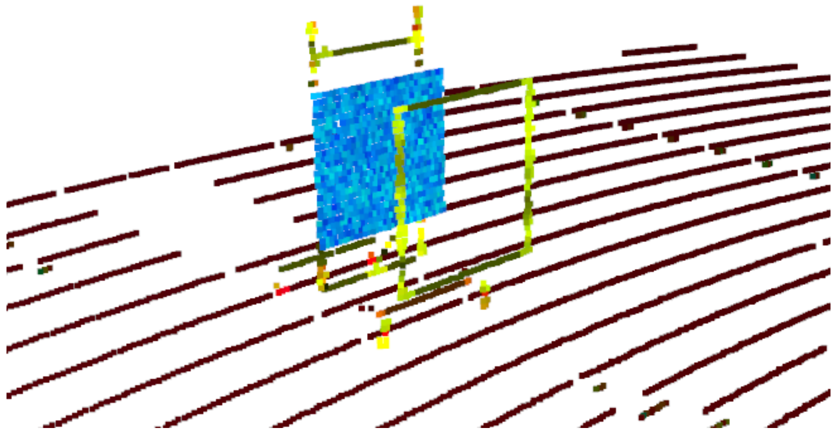

Figure 35: Sim replica point cloud of real world scene of 84% reflective target with plexiglass target in front.

Figures 36a-c and 37a-c show the 3 metrics for this experiment configuration at 24m and 103m ranges.

24m Range

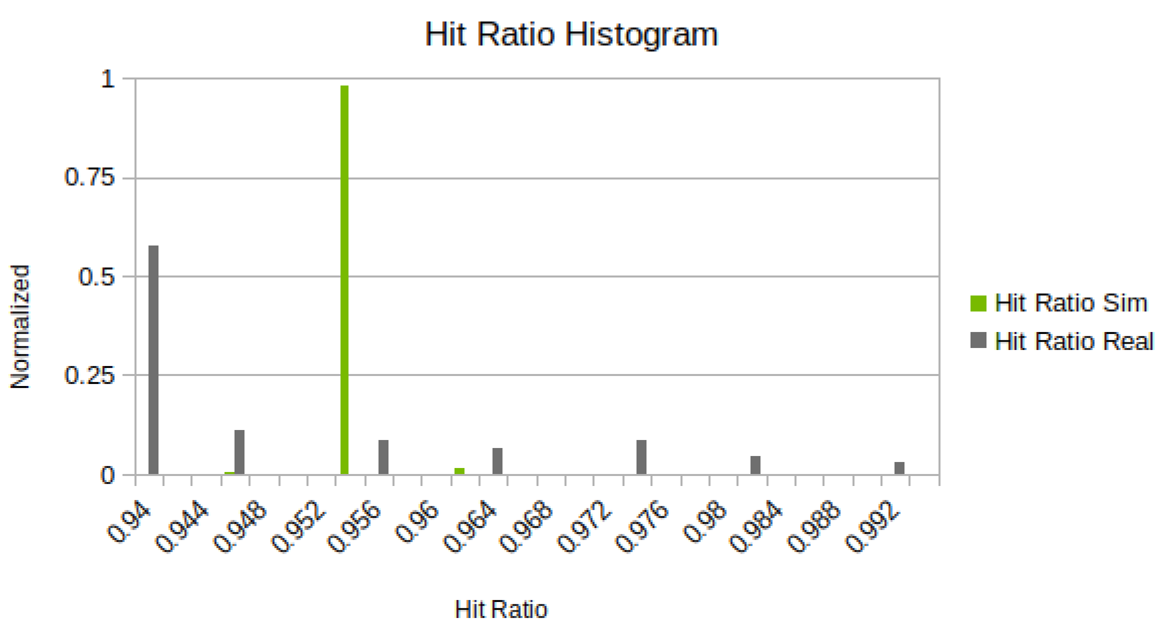

Figure 36(a): Hit ratio of Sim and Real at 24m range.
Real Mean - 0.92, Real STD - 0.07
Sim Mean - 0.967, Sim STD - 0.005

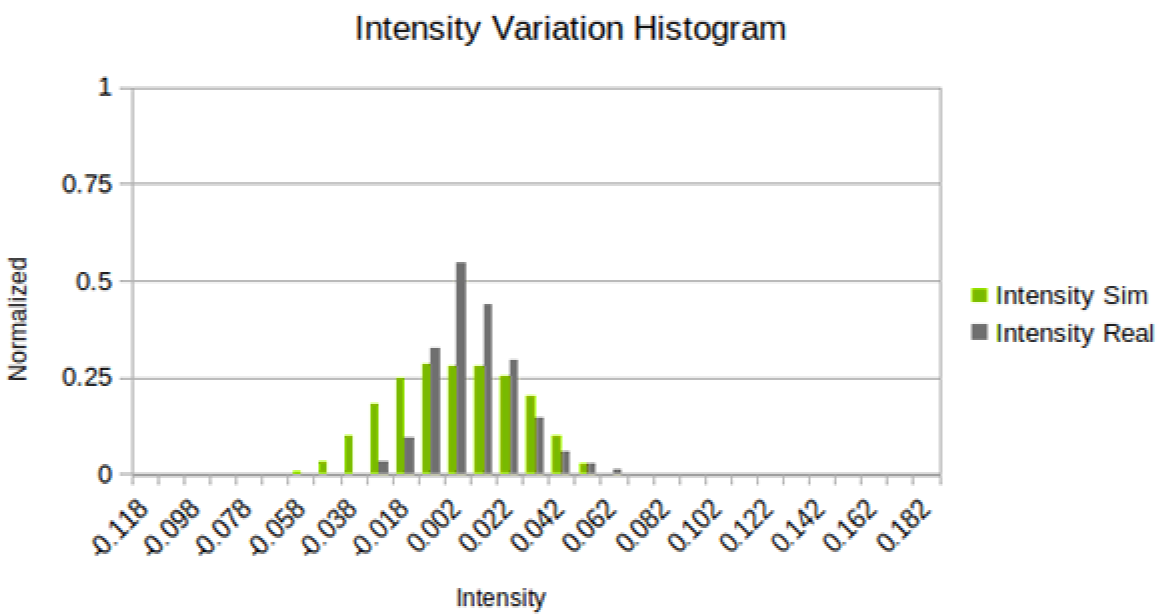

Figure 36(b): Intensity Variation of Sim and Real at 24m range.
Real STD - 0.022; Sim STD - 0.023

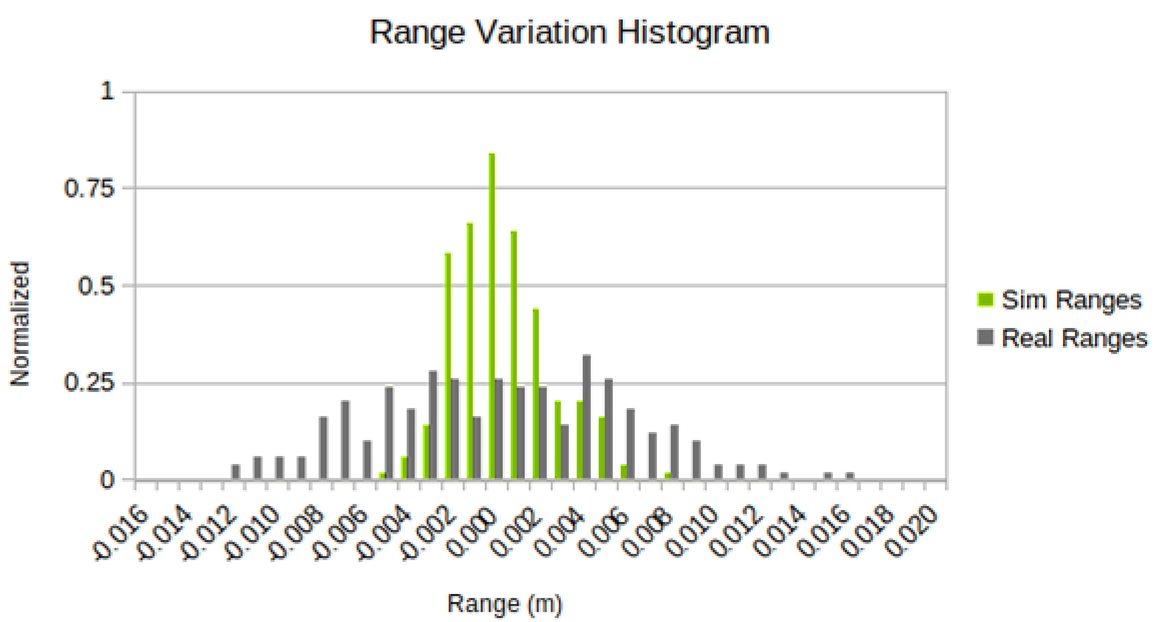

Figure 36(c): Range of Sim and Real at 24m range.
Real STD - 0.006; Sim STD - 0.002

Similar to the trends found in the lidar scenario 2 for the 5% reflective panel, the addition of the plexiglass panel created greater variances in all the datasets. The real world data is affected even more by the sun glint making it more noisy.

103m Range

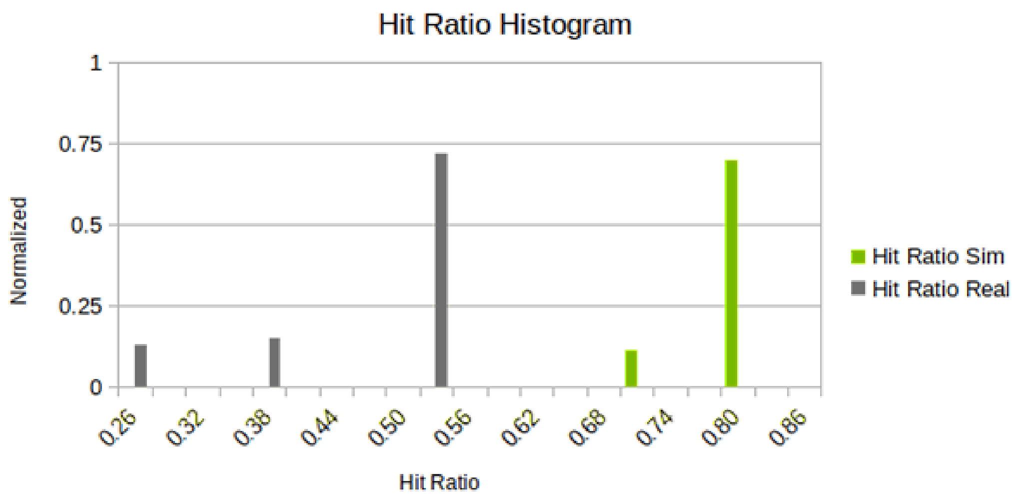

Figure 37(a): Hit ratio of Sim and Real at 103m range.
Real Mean - 0.46, Real STD - 0.7
Sim Mean - 0.78, Sim STD - 0.042

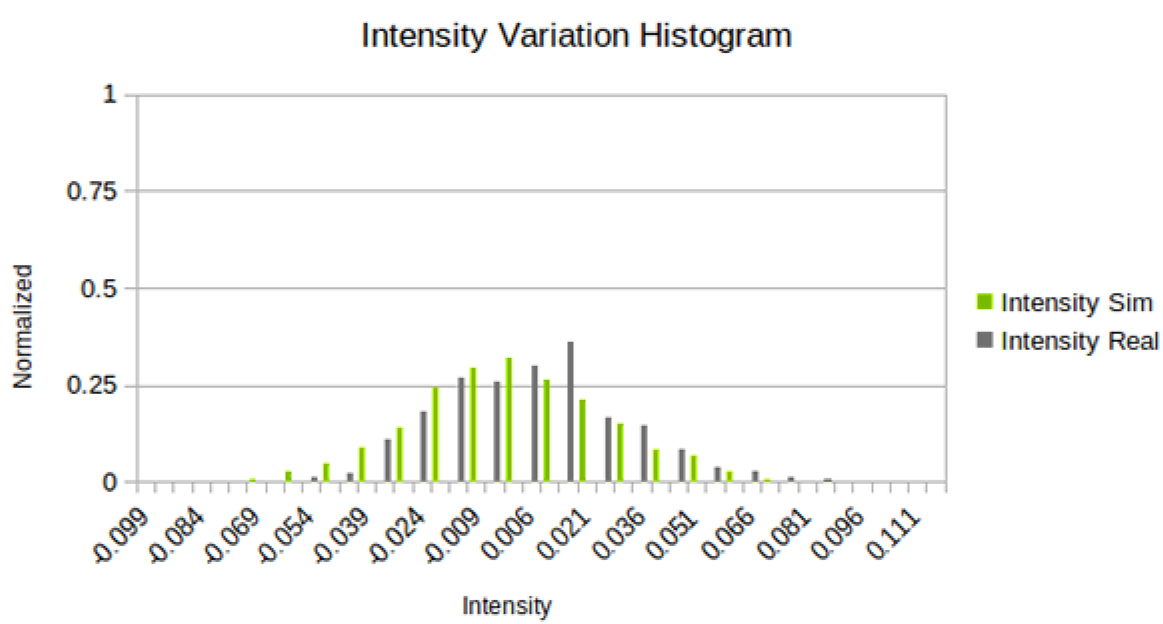

Figure 37(b): Intensity of Sim and Real at 103m range.
Real STD - 0.037; Sim STD - 0.026

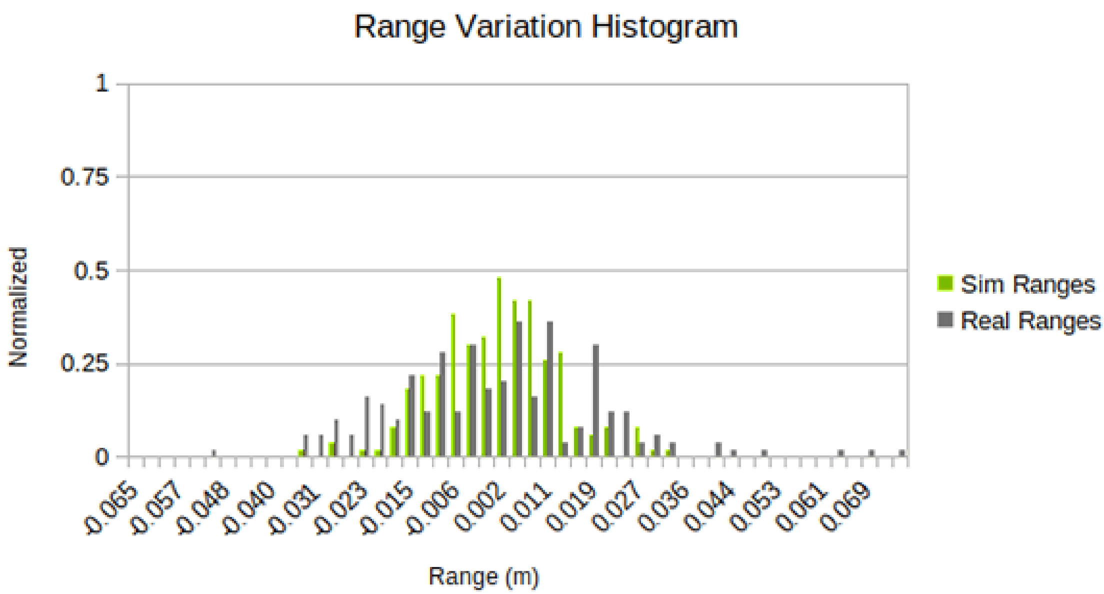

Figure 37(c): Range of Sim and Real at 103m range.
Real STD - 0.019; Sim STD - 0.011

As anticipated, the larger ranges exacerbate the standard deviations and reduce mean values further for the raw data metrics.

The tables 4a and b in the data table appendix depict the real world and sim data statistics for the set ranges identified for lidar scenario 2 on the 84% reflective target panel with a transparent plexiglass panel in front of the target. This data presented here is visualized in the following figures of 38 and 39.

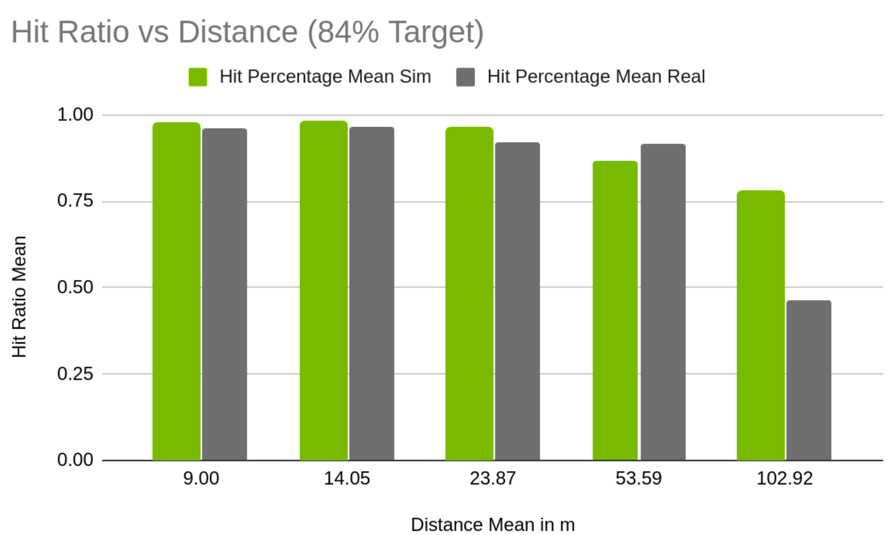

Figure 38: 84% reflective panel hit ratios as a function of range with plexiglass panel.

Just as in the previous experiment, the hit ratios show good correlation for the different ranges. Both sim and real show smaller ratios at longer ranges. Ranges under 100m exhibit a relative error no larger than 5% and 41% for ranges greater than 100m. One common reason for this is the additional losses and resultant dropped points with the addition of the plexiglass panel in front of the target. The other factor is from the already discussed solar glint noise. The increased noise produced by the stray reflections at larger ranges reduces the points recorded from the target panel. This increase in stray light noise is not modeled in the simulation which accounts for the clear increase in hit ratio for sim compared to real at the farthest range.

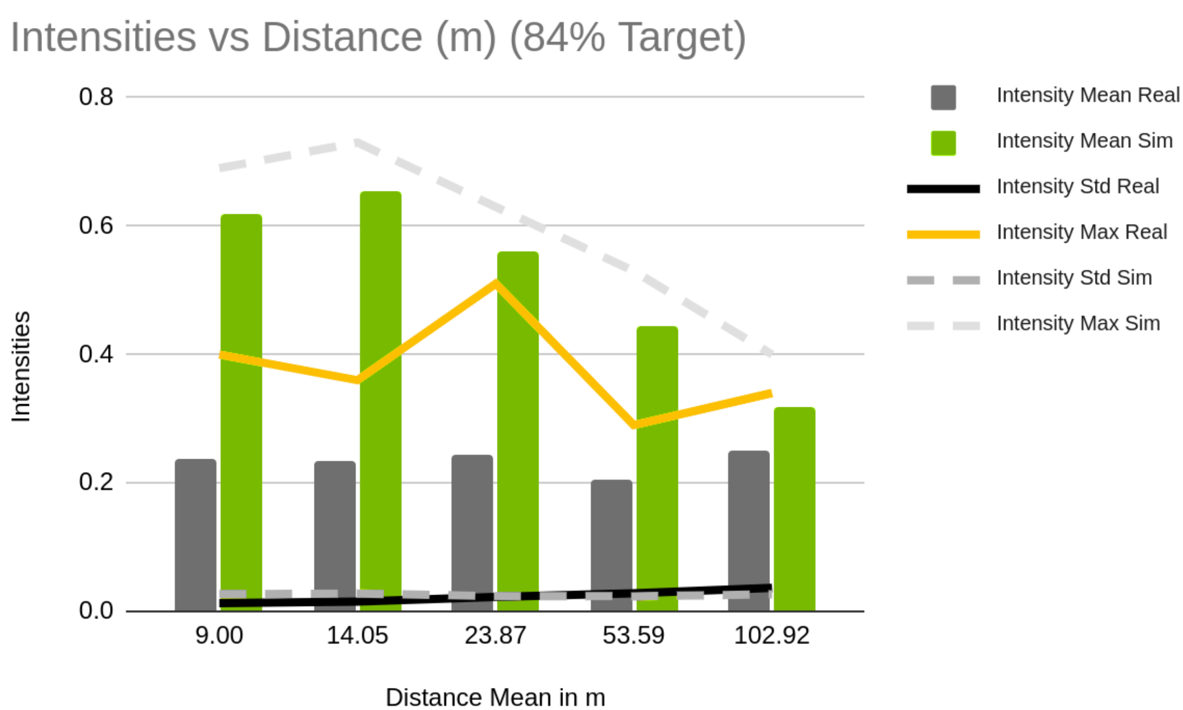

Figure 39: 84% reflective panel intensities as a function of range with plexiglass panel.

Similar conclusions from the 5% reflective panel case in lidar scenario 2 can be concluded here for the 84% reflective panel. The data shows good correlation for the intensity standard deviations with differences less than a factor of 2. This highlights the precision between sim and real. The mean intensities of sim vs real are plausible despite the observable differences. Overall, there is a good level of correspondence for all the statistics between real and sim in this experiment.

Data Table Appendix

5% Reflectivity Target Panel

The tables below depict the summary statistics for the 5% reflective target panels of lidar scenario 1 and 2.

Real Sensor Evaluation for lidar scenario 1:

Table 1a - Real world data for 5% reflective target in lidar scenario 1

Distance Mean	Hit Percentage Mean	Intensity Mean	Intensity Std	Intensity Max	Intensity Min
7.81	0.967294197	0.02101267454	0.003017057198	0.03999999911	0.01999999955
12.25	0.9238344035	0.02202648263	0.004019725946	0.02999999933	0.01999999955
22.02	0.9208126035	0.02437971051	0.004961374862	0.02999999933	0.01999999955
52.10	0.96	0.01957754586	0.002247032732	0.02999999933	0.009999999776
101.53	0.9333333333	0.01887499958	0.003252631646	0.02999999933	0.009999999776

Sim Sensor Evaluation for lidar scenario 1:

Table 1b - Sim data for 5% reflective target in lidar scenario 1

Distance Mean Sim	Hit Percentage Mean Sim	Intensity Mean Sim	Intensity Std Sim	Intensity Max Sim	Intensity Min Sim
7.665318785	0.973097969	0.04198277962	0.003987153166	0.05003260075	0.02999742933
12.07152951	0.9581236674	0.04527468314	0.003162068569	0.05156088879	0.03478847741
22.03852106	0.9730272109	0.03232530427	0.004243442663	0.03999999911	0.01999999955
52.1548336	0.9375	0.02710999939	0.005352373187	0.04200000075	0.009999999776
101.4751033	0.9566666667	0.01554064523	0.004919266579	0.03149	0.002465

Real Sensor Evaluation for lidar scenario 2:

Table 2a - Real world data for 5% reflective target in lidar scenario 2

Distance Mean	Hit Percentage Mean	Intensity Mean	Intensity Std	Intensity Max	Intensity Min
9.336848153	0.9627697368	0.0189564096	0.003064400877	0.03999999911	0.009999999776
14.21689759	0.9582886905	0.0192819537	0.00258464644	0.02999999933	0.009999999776
23.91789613	0.9128703704	0.02009429235	0.002029813031	0.03999999911	0.009999999776
53.46670504	0.9472727273	0.01004798442	0.0006910455134	0.01999999955	0.009999999776
102.9045694	0.5875	0.01101063805	0.003014132181	0.01999999955	0.009999999776

Sim Sensor Evaluation for lidar scenario 2:

Table 2b - Sim data for 5% reflective target in lidar scenario 2

Distance Mean	Hit Percentage Mean	Intensity Mean	Intensity Std	Intensity Max	Intensity Min
9.329848043	0.9688016796	0.0350808584	0.003029910259	0.0440305884	0.0260111284
14.2437297	0.973457223	0.0360980177	0.00302622269	0.0451156857	0.0270152896
23.93343602	0.9674731183	0.03374622505	0.004866587693	0.04834598815	0.01924646955
53.47365521	0.8656818182	0.02168679395	0.004502335741	0.03999999911	0.009999999776
102.9206415	0.779375	0.01419406544	0.004934619429	0.01999999955	0.009999999776

84% Reflectivity Target Panel

The tables below depict the summary statistics for the 84% reflective target panels of lidar scenario 1 and 2.

Real Sensor Evaluation for lidar scenario 1:

Table 3a - Real data for 84% reflective target in lidar scenario 1

Distance Mean	Hit Percentage Mean	Intensity Mean	Intensity Std	Intensity Max	Intensity Min
7.89	0.9657218845	0.2911447903	0.01610156535	0.6800000072	0.05000000075
12.35	0.9677146172	0.3263315659	0.01430384788	0.4799999893	0.1099999994
22.45	0.9750715286	0.3125859883	0.02064251943	0.3899999857	0.2599999905
52.04	0.9551162791	0.300064526	0.02582240236	0.3899999857	0.09000000358
101.47	0.9266666667	0.3323415471	0.02452006597	0.4059015	0.2587813495

Sim Sensor Evaluation for lidar scenario 1:

Table 3b - Sim data for 84% reflective target in lidar scenario 1

Distance Mean Sim	Hit Percentage Mean Sim	Intensity Mean Sim	Intensity Std Sim	Intensity Max Sim	Intensity Min Sim
7.924411908	0.9712347418	0.7094463309	0.03007601358	0.7699999809	0.6200000048
12.37685046	0.9824193548	0.7407663757	0.03055611283	0.8100000024	0.6499999762
22.03854156	0.9874827586	0.6857364715	0.02789937758	0.7599999905	0.6100000143
52.15455138	0.9647457627	0.5306280704	0.02679109349	0.6200000048	0.4399999976
101.4749397	0.8603846154	0.3167545843	0.02412193672	0.4099999964	0.2399999946

Real Sensor Evaluation for lidar scenario 2:

Table 4a - Real data for 84% reflective target in lidar scenario 2

Distance Mean	Hit Percentage Mean	Intensity Mean	Intensity Std	Intensity Max	Intensity Min
8.997084077	0.9617892157	0.2371291493	0.012842473	0.400000006	0.01999999955
14.04962671	0.9655135952	0.2333889248	0.01500745524	0.3600000143	0.1299999952
23.87406773	0.9215765766	0.2436722216	0.02248548116	0.5099999905	0.1199999973
53.59222828	0.9145652174	0.2036985973	0.02827775593	0.2899999917	0.05999999866
102.9196488	0.4614450516	0.2500789466	0.03671360216	0.3400000036	0.1700000018

Sim Sensor Evaluation for lidar scenario 2:

Table 4b - Sim data for 84% reflective target in lidar scenario 2

Distance Mean	Hit Percentage Mean	Intensity Mean	Intensity Std	Intensity Max	Intensity Min
9.031115017	0.9740079365	0.6199410396	0.02671286967	0.6899999976	0.5299999714
14.07533663	0.9780706522	0.6530133655	0.02766231311	0.7300000191	0.5600000024
23.89326576	0.9614802632	0.5600287351	0.0237964778	0.6299999952	0.4799999893
53.61367766	0.8646590909	0.442696804	0.02345183192	0.5299999714	0.3600000143
102.9206415	0.776375	0.3195669634	0.02640086702	0.400000006	0.2300000042

Resources

[1] https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7765299/#B6-sensors-20-07186

[2] https://ieeexplore.ieee.org/abstract/document/8317864

[3] https://ieeexplore.ieee.org/abstract/document/8691584

[4] https://dl.acm.org/doi/abs/10.1145/3206025.3206080

[5] http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.605.5992