

Detection of Specular Reflections in Range Measurements for Faultless Robotic SLAM

Rainer Koch¹, Stefan May¹, Philipp Koch¹, Markus Kühn¹, and Andreas Nüchter²

¹ Technische Hochschule Nürnberg Georg Simon Ohm
Kesslerplatz 12, 90489 Nürnberg, Germany
{rainer.koch}@th-nuernberg.de
<http://www.th-nuernberg.de>

² Informatics VII – Robotics and Telematics
Julius-Maximilians University Würzburg
Am Hubland, 97074 Würzburg, Germany
{andreas}@nuechti.de
<http://www.uni-wuerzburg.de>

Abstract. *Laser scanners are state-of-the-art devices used for mapping in service, industry, medical and rescue robotics. Although a lot of work has been done in laser-based SLAM, maps still suffer from interferences caused by objects like glass, mirrors and shiny or translucent surfaces. Depending on the surface's reflectivity, a laser beam is deflected such that returned measurements provide wrong distance data. At certain positions phantom-like objects appear. This paper describes a specular reflectance detection approach applicable to the emerging technology of multi-echo laser scanners in order to identify and filter reflective objects. Two filter stages are implemented. The first filter reduces errors in current scans on the fly. A second filter evaluates a set of laser scans, triggered as soon as a reflective surface has been passed. This makes the reflective surface detection more robust and is used to refine the registered map. Experiments demonstrate the detection and elimination of reflection errors. They show improved localization and mapping in environments containing mirrors and large glass fronts is improved.*

Keywords: SLAM · error-free mapping · multi-echo laser scanner · reflectance filter · specular reflection · reflective objects

1 Introduction

Mapping is an essential task in mobile robotics. It is used in service robotics, e.g., in industrial, medical, and rescue applications. Before training or exploration, the environment is partly or completely unknown. SLAM (Simultaneous Localization and Mapping) is one of the most frequently applied approaches to provide an environmental representation to service robots. Nevertheless, customizing most environments is necessary to reduce interferences from objects with specular

reflective and transparent surfaces, e.g., glass, mirrors, and shiny metal. This is one reason why robots are still not ad-hoc integrable in most applications.

In addition, glass surfaces are reflective or transparent depending the laser beam’s incident angle. Objects behind the glass surface are only occasionally visible. Even worse is the aspect that the glass surface is rated as volatile (disappearing) object or not seen at all. This carries the risk of navigating a robot into it. Hence most service robots are reliant on a second sensor principle, like ultrasonic arrays, to respect these situations. Despite the fusion with other sensor

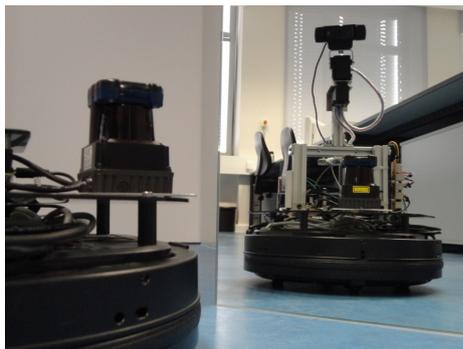


Fig. 1: Robot equipped with laser scanner facing a unframed mirror.

principles, it is difficult to register a laser scan based map without reflection influences. Therefore the environment is modified. If the laser beams hit the surface in an angle associated to total reflection, returned measurements provide wrong distance data. At certain positions phantom-like objects appear in the map. Figure 2 depicts this effect for three state-of-the-art SLAM approaches using the same dataset: CRSM-SLAM (Critical Rays Scan Match-SLAM), Hector-SLAM and TSD-SLAM (Truncated Signed Distances-SLAM). Phantom-like areas are marked with a red rectangle. The location of the mirror is marked with a blue rectangle and magnified on the top left. Hector-SLAM creates a static map, i.e., points added once to the map remain ad infinitum. The mirror is partly recognizable in the Hector-SLAM map due the fact that at some positions the laser beam was not deflected. In comparison, CRSM- and TSD-SLAM build a dynamic map. Changes in the environment are respected in both approaches, e.g., if objects are moved. Therefore, the mirror disappears if its surface is not measurable at certain perspective views. This is likely the case for passing a mirror.

In the following, we present a reflectance detection approach applicable to multi-echo laser scanners in order to remove above mentioned effects. Section 2 outlines related work. Section 3 describes the two filter stages used for mirror detection. In Section 4 experiments demonstrate the applicability to environments with a large proportion of reflective and transparent surfaces. Finally, Section 5 summarizes results and gives an outlook for future work.

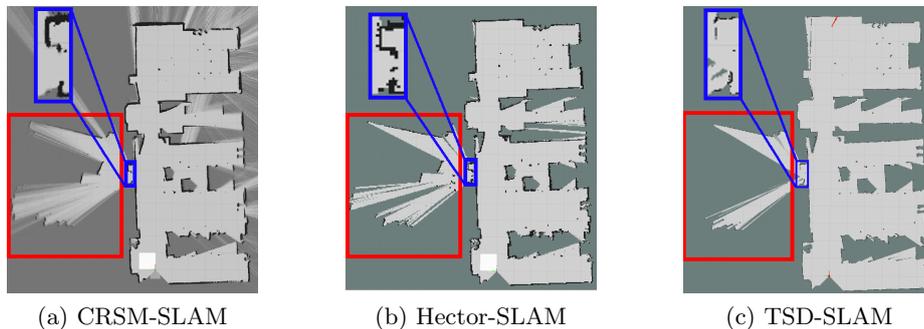


Fig. 2: Maps registered with the same dataset and with different SLAM approaches in environment containing a mirror. The mirror is marked blue and the reflections are marked red.

2 Related Work

As far as reflection is concerned, there are two different strands of research - stationary and mobile systems. Often, the environment is adapted to prevent influences when working with stationary systems. Therefore, research in this field has less impact on mapping. Covering objects is unwanted when mapping with mobile systems because it requires a lot of effort to deal with all specular reflective and transparent objects, especially when operating outside.

To avoid the need to cover surfaces for mapping, Yang et al. [1] presented an approach which fuses a laser scanner with an ultrasonic sensor. Two individual grid maps are created. With the assumption that reflective objects are flat and framed, the data from the two sensors are compared w.r.t. consistency. Mirrors are detected and tracked online, while resulting errors are recalculated only offline. In further research Yang et al. [2] extended their algorithm for advanced mirror detection and identification of mirror images. Each gap in the wall is assumed to be a specular reflective object. Therefore, no ultrasonic sensor is required anymore. Once such a mirror candidate is detected, the space behind the gap is analysed for a mirrored image, i.e., the search for similarity between both sides of the opening. Objects with symmetry w.r.t. a line might be identified wrongly.

Another online applicable approach was implemented by Forster et al. [3]. At specific angles reflections can be identified based on the returning intensity of the laser. A subset of these angles is tracked - on occurrence mirrors are assigned in dependency of the laser beam's intensity. An object with diffuse reflectivity causes false identification if it is placed directly behind the transparent object.

Käshammer et al. [4] presented an approach which recognises framed mirrors with a predefined size in 3D point cloud data. A panorama range image is generated and searched for jumping edges. In case of a positive search, the contour of the mirror frame will be extracted. Finally, objects are verified by considering

their size and shape. This only applies to framed squared mirrors with a known size. Glass or other objects are not considered.

Tatogulu et al. [5] use the best fitting illumination model to modulate the surface. Lambertian diffuse reflection models, Blinn-Phong models [6], Gaussian models [7] and Beckmann specular reflection models [8] are fitted to the data set to identify the characteristics of the scanned surface. While this system is quite effective for diffuse surfaces, it does not handle specular reflections.

We present a specular reflectance detection approach applicable to multi-echo laser scanners in order to identify and filter mirrored objects. In contrast to above mentioned approaches, recognition of specular reflections is possible, also for frameless and free-standing objects.

3 Approach

The mirror detector uses a Hokuyo 30LX-EW multi-echo laser scanner. For each data take the Hokuyo records up till three echoes of the returning light wave, including distance and intensity. The first two echoes s_1 and s_2 , which are further called scan tuple s , are taken to detect reflective objects:

$$s = \{s_1, s_2\} \quad (1)$$

with

$$s_1 = \{d_{1,i} | i = 1 \cdots N\}, \quad (2)$$

$$s_2 = \{d_{2,i} | i = 1 \cdots N\}, \quad (3)$$

where $d_{1|2,i}$ are distance measurements and N is the number of measurement points.

Differences in scan messages indicate surface reflection properties. While a specular reflective object causes differences in both scan messages, diffuse reflective objects provide consistency. The problem in detecting specular reflections is that it depends on the laser beam’s incident angle to the surface and the refractive index. If the angle is too big, the light will be totally reflected according to the reflection law. Hence, the robot will only detect the mirrored object. If the angle is smaller, there are three potential cases of measurements also depending on the material. For a transparent object the robot can receive a point on the surface, a point behind the surface, or a mirrored point. If the object is nontransparent, the last case disappears. Therefore, the robot has to pass the surface to ensure that it is seen at least once from the “right” perspective. Hence, it is not possible to eliminate all reflective errors on the fly. That is why the mirror detector is set up in two filter stages and two mapping stages, cf. Figure 3. The pre-filter runs on the fly and filters current scans. The post-filter is triggered after a reflective object has been passed, e.g. by an passing-algorithm or a loop-closure, to reduce remaining errors.

Figure 3a shows the processing chain of the pre-filter with its mapping stage. The resulting map is without reflection errors, which are detectable in a single data take. Hence, the map include less erroneous data than a map

using data directly provided by the Hokuyo. But the map is not completely cleaned of reflective influences. The mapping stage assigned to the post-filter

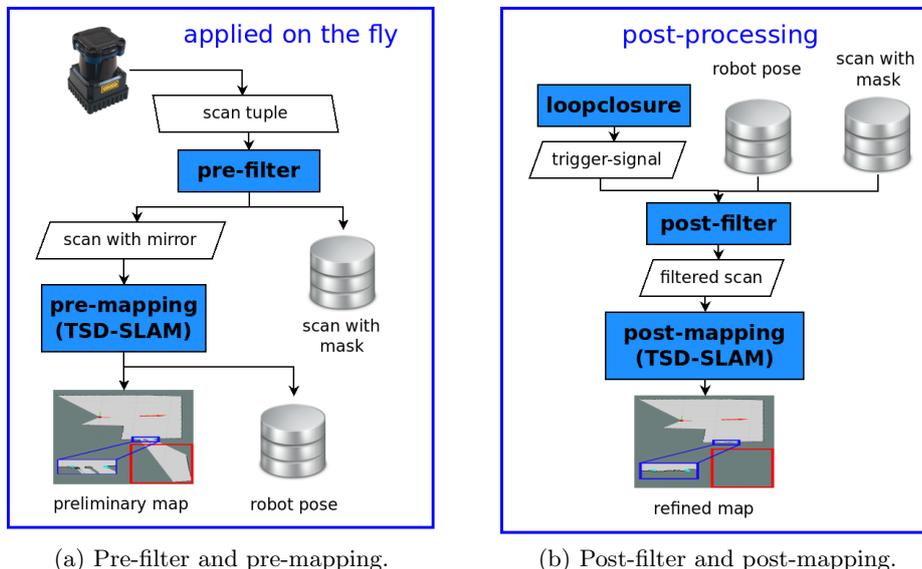


Fig. 3: Processing chains of the mirror detector: Pre-filtering removes affections on the fly. Post-filtering refines the resulting map after a trigger signal.

chain, cf. Figure 3b, considers a set of scan tuple. All measurements originating from specular reflective objects are filtered in retrospect, i.e., specular reflective objects detected in any scan tuple are propagated to the whole history. This supplies a map free of any reflection errors available at new trigger events, e.g., from a loop closure module. All modules are implemented as ROS-nodes and are publicly available as open-source packages at http://www.github.com/autonohm/ohm_mirror_detector.git.

3.1 Pre-filter

The pre-filter receives a scan tuple from the laser scanner. First, it removes sparse points, i.e., isolated points without other points nearby. These points are likely to be artefacts for example from jumping edges. This happens when neighbouring measurements cross an object edge and provide discontinuity in depth. Further, the corresponding points in the scan tuple are subtracted and analysed. A difference between s_1 and s_2 points out that the laser beam was reflected. This happens when it hits, e.g., a reflective or transparent surface. The first echo has to be from the error causing object, since it was hit first by the laser beam. The second echo includes a point more fare away, an affected

point. It is identified according to the distance between the scan messages:

$$\Delta d_i = d_{2,i} - d_{1,i}, \quad (4)$$

$$f(\Delta d_i) = \begin{cases} d_{1,i} \rightarrow \text{valid}, & \text{if } \Delta d_i \leq \text{threshold} \\ d_{1,i} \rightarrow \text{mirror}, d_{2,i} \rightarrow \text{affected}, & \text{if } \Delta d_i > \text{threshold}. \end{cases} \quad (5)$$

Glass fronts and mirrors are assumed to be planar surfaces. They relate to line segments in a scan message. The two distinct end points \mathbf{c}_1 and \mathbf{c}_2 are determined with a RANSAC-based algorithm, i.e., finding the line parameters fitting best to the set of mirror points in Cartesian space.

With robot position \mathbf{p} a sector is spanned up, cf. Figure 4. Finally, the scan is checked again to classify the points into three groups: valid, mirror plane, and affected points. Valid points are located in the green hatched area. They are free of reflection influences. The second group contains points on the mirror, glass, or reflective plane, e.g., blanc metal. They are found in the solid blue area. All other points remain in the red crossed area and shall be assigned to the third group. The pre-filter processing time is less than 4.5ms, which caused mainly by the RANSAC-algorithm.

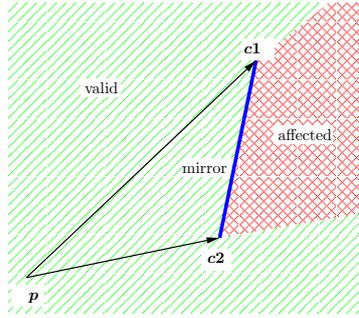


Fig. 4: Classification of points based on the mirror line corners.

3.2 Pre-mapping

The first two point groups, containing valid and mirror points, are forwarded as scan messages to the pre-mapping module. As a result, the preliminary map is generated on the fly. The pre-mapping employs an unadjusted version of TSD-SLAM as described in [9]. Nevertheless other mapping algorithms can be integrated as well. In order to do so, the SLAM module must provide the scanner's pose.

3.3 Loop closure

A simple loop closure algorithm compares the current robot position with the previous robot positions. Therefore, it records a complete history of the robot pose. If the new position is within a limited range to previous positions, a trigger signal is broadcasted to the post-filter.

3.4 Post-filter

A mirror is only detectable, if the incoming laser beam hits the mirror in a particular angle. Hence, the post-filter builds a history of all pre-filtered scans and poses with dynamic length. If a dataset includes mirror points the post-filter identifies the line associated to the mirror plane. The corners of the line are simply the outer points in the scan because the dataset is ordered according the increase of angle and the scan has been cleaned in the pre-filter. Received mirror corners are added at the first occurrence. Corners nearby existing mirror corners are fused together.

While building the history, the post-filter awaits a trigger signal. It is provided for instance by the loop closure module or any other external trigger. Afterwards the post-filter starts to refine the scans in the history. It uses the same algorithm as in the pre-filter module to span up a sector and classify the points into the three groups. This will be repeated for every set of distinct end points. The processing time of the post-filter differs. Without a running publisher the processing time is less than 30s. Afterwards it rises up to maximum of 1.5ms. Finally, the refined scans are transferred to the post-mapping module. Thus, the sensor's final pose is determined by the SLAM-module on the basis of the refined scans.

3.5 Post-mapping

Post-mapping also applies an untouched version of the TSD-SLAM approach. It delivers a refined map. The SLAM module can be replaced by any other approach. The map is registered with refined scans from the history and therefore without reflection errors.

4 Experiments and Results

This chapter consists of three sections to qualify the experiments with the mirror detector. The first experiment uses a "sandbox setup" to test the mirror detector on a defined scene. Experiment 2 has been performed in an office-like environment. Experiment 3 applies the approach to a corridor with a large glass front.

4.1 Experiment 1: "Sandbox"

The "sandbox" is used to evaluate the mirror detection approach on a simple scene with a defined mirror location. The setup is shown in Figure 5 and the mirror is marked with a blue square.

As already described, the mirror cannot be seen from every pose, even it is in the field of view of the laser scanner. Only if the angle of the incoming laser beam hits the mirror plane in a particular angle, the scanner will broadcast different values in its echoes. As a result, the mirror is identified. If a mirror plane was

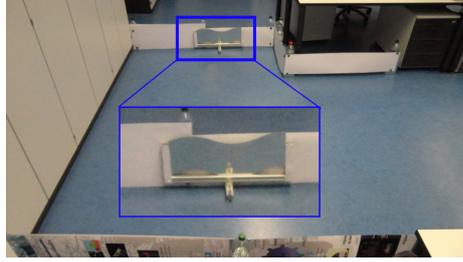
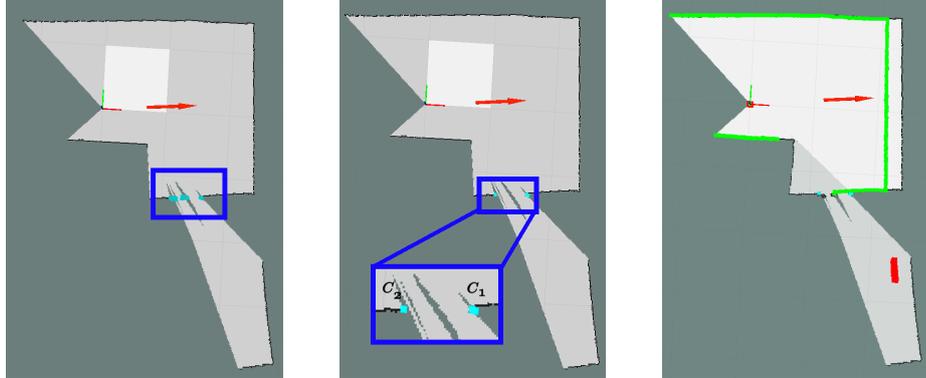


Fig. 5: "Sandbox"-experiment with mirror and simple scene.

detected, cf. Figure 6a, the history is searched for erroneous points. Therefore, the post-filter stores the history of all scans.

We assumed a planar surface of the reflective object. That is why, two boundary points c_1 and c_2 , cf. Figure 6b, are enough to describe the subject. These corners will be updated when more points on the mirror plane are determined or added as new ones. When receiving a trigger signal, the post-filter, cf. Figure 6c, masks all points in the history. Subsequently, it broadcasts a refined scan, including valid and mirror plane points (green) and a scan with erroneous points (red). The post-mapping uses the valid scans to build up a map including the mirror plane. The erroneous points are not used yet, but it is our intent to recalculate them to their true position and embed them in mapping. Figure 7a evinces the preliminary map and Figure 7b the refined map to show the difference.

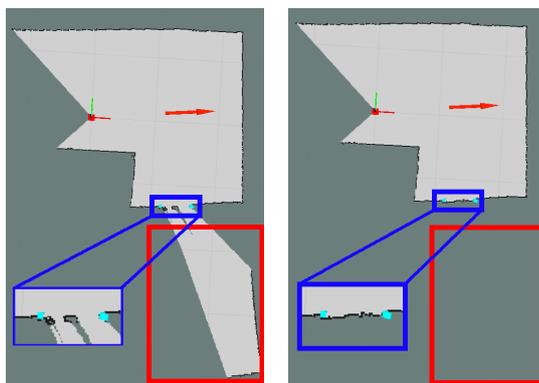


(a) Exp. 1: Detection of mirror plane.

(b) Exp. 1: Resulting mirror corner points.

(c) Exp. 1: Post-filtered scan (green), erroneous points (red).

Fig. 6: Exp. 1: Different steps of the mirror detection.



(a) Exp. 1: Preliminary map without the mirror surface and including reflections. (b) Exp. 1: Refined map including the mirror surface and free of reflections.

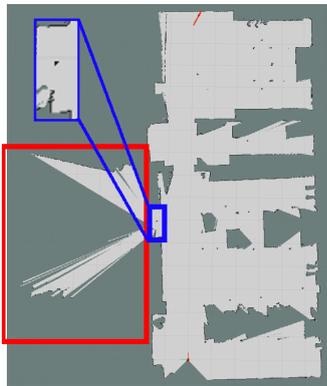
Fig. 7: Exp. 1: Comparison between preliminary map and refined map. The mirror is marked blue and reflections are marked red.

4.2 Experiment 2: Office-like Environment

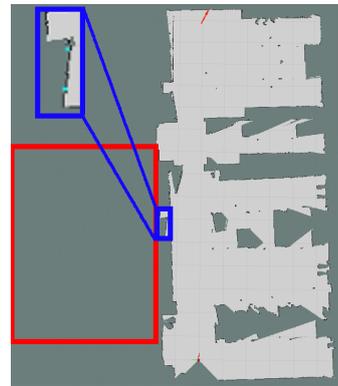
The map of experiment 2 contains three office rooms with a mirror, cf. blue squared in Figure 8. It is necessary to note that the mirror is not planar to the wall, as visible in the magnified area of the refined map. Comparing the preliminary map in Figure 9a and the refined map in Figure 9b the effect of the mirror detector is visible. There are no remaining reflections in the refined map (red square). In addition the mirror plane is completely mapped (blue square). In this case the robot will not try to navigate through the mirror plane.



Fig. 8: Exp. 2: Office room with mirror.



(a) Exp. 2: Preliminary map without the mirror surface and including reflections.



(b) Exp. 2: Refined map including the mirror surface and free of reflections.

Fig. 9: Experiment 2 - office-like environment. The mirror is marked blue and reflections are marked red.

4.3 Experiment 3: Room with Corridor

In this experiment a corridor with a large glass front was mapped, cf. Figure 10. Such glass fronts are a major reason for erroneous measurements. Therefore, they are normally covered by hand.

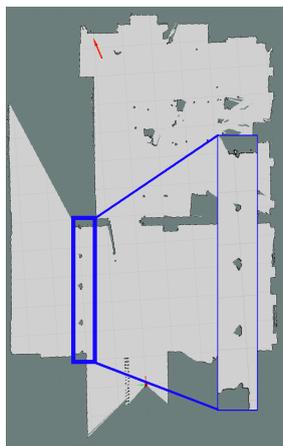


Fig. 10: Exp. 3: Corridor with glass front.

Figure 11a displays the preliminary map of the mirror detector. The glass front is marked with a blue square. Above the glass front another part of the corridor is mapped. This is correct and also visible in Figure 10. As previously described, this is a possible result if the robot faces a transparent surface. The

mirror detector is not able to distinguish between mirror and glass yet. Therefore it erases the points behind the detected surface, cf. Figure 11b. In addition, it marks the points on the surface and therefore the post-mapping includes them. This is wanted, since it will prevent the robot to think there is a free path. Still some area behind the glass remained in the refined map. This is caused by the fact, that the laser beams had not yet hit the surface in the desired angle. In this particular case, the glass behind the open door was previously seen, but when passing, the door blocked the laser beam from reaching the desired angle. That is why, the glass area cannot be identified as such.

This experiment emphasizes that further investigation is required. It is necessary to distinguish between transparent and mirror surfaces. Thereby, a discrimination of the points behind the surface is possible. Further, there is a need to improve the broadcast of the trigger signal. A trigger module is needed to assure that a certain position before the surface has been reached. In that case it is verified, that it is not a reflective area.



(a) Exp. 3: Preliminary map without the glass surface.



(b) Exp. 3: Refined map including the glass surface.

Fig. 11: Experiment 3 - room with corridor. The glass front location is marked blue.

5 Conclusions and Future Work

The mirror detector identifies transparent and specular reflective objects like glass, mirrors, or shiny surfaces such as blank metal. Free standing unframed objects of different size are detectable. Two filter stages are implemented. The first filter reduces errors in the current scans on the fly. A subsequently applied SLAM module builds a preliminary map and provides the robot poses for the

post-filter. The post-filter records the pre-filtered scans and the robot pose in a history. Besides, it calculates the corner points of the mirror plane. As soon as a trigger signal from the loop closure occurs, the post-filter delivers refined scans to the second mapping stage. The result is a refined map without reflection errors. The TSD-SLAM module allows to map in a dynamic world. However, the influences of moving reflective objects are not tested yet.

Future work is concentrated on advanced loop closure techniques, classification of transparent and reflective materials, as well as detecting shaped reflective objects. Hence, the plane detection algorithm will be replaced. These will help to archive some difficult but interesting tasks, such as mapping modern glass galleries or historic buildings. Castles and palaces are full of reflective and transparent objects, e.g., chandeliers, golden artwork, or mirror cabinets. The aim is to support robust localisation and mapping in such areas.

References

1. Yang, S.W., Wang, C.C.: Dealing with laser scanner failure: Mirrors and windows. *IEEE International Conference on Robotics and Automation ICRA* (2008) 3009–3015
2. Yang, S.W., Wang, C.C.: On solving mirror reflection in lidar sensing. *Mechatronics, IEEE/ASME Transactions on* **16**(2) (2011) 255–265
3. Foster, P., Sun, Z., Park, J.J., Kuipers, B.: Visagge: Visible angle grid for glass environments. In: *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. (2013) 2213–2220
4. Käshammer, P.F., Nüchter, A.: Mirror identification and correction of 3d point clouds. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **XL-5/W4** (2015) 109–114
5. Tatoglu, A., Pochiraju, K.: Point cloud segmentation with lidar reflection intensity behavior. In: *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. (2012) 786–790
6. Blinn, J.F.: Models of light reflection for computer synthesized pictures. In: *Proceedings of the 4th Annual Conference on Computer Graphics and Interactive Techniques. SIGGRAPH '77, New York, NY, USA, ACM* (1977) 192–198
7. Torrance, K.E., Sparrow, E.M.: Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.* **57**(9) (Sep 1967) 1105–1112
8. Beckmann, P., Spizzichino, A.: *The scattering of electromagnetic waves from rough surfaces*. Pergamon Press (1963)
9. Koch, P., May, S., Schmidpeter, M., Kühn, M., Martin, J., Pfitzner, C., Merkl, C., Fees, M., Koch, R., Nüchter, A.: Multi-robot localization and mapping based on signed distance functions. In: *Autonomous Robot Systems and Competitions (ICARSC), 2015 IEEE International Conference on*, <https://www.waset.org/conference/2015/03/istanbul/ICARSC> (03/2015 2015)