P2 Graduate Plan

Mels Smit
student #4601866
melssmit@tudelft.nl

1st supervisor: Edward Verbree
2nd supervisor: Martijn Meijers
external supervisor: Robert Voûte

January 05, 2021

1 Introduction

Large buildings and other indoor environments such as shopping centres, airports and office buildings suffer from the lack of up-to-date 3D Models (Nikoohemat et al., 2019). When something happens, having a digital twin of the building can be of great help with all kinds of processes such as evacuation, recovering damages and resource management. In recent years there have been numerous improvements in the field of Mobile Laser Scanning (MLS) making it possible to scan these larger complexes in just a few hours by traversing the buildings themselves with a wearable scanner (Lehtola et al., 2017). This solution produces point clouds and images as output in a time-efficient manner, which capture the current status of the building (Nikoohemat et al., 2019), thereby functioning as a digital twin. These digital twins are, however, more than likely not complete. While LiDAR point clouds are a great medium to convey 3D spatial information, which can be both time-efficient using the Mobile Laser Scanning approach or accurate using a more traditional although slower Terrestrial Laser Scanning (TLS) approach they are still held back by the physical properties of their capturing medium, namely laser.

Laser scanning which is the foundation of LiDAR is done with light at particular wavelengths to get the best results, taking into account all kinds of factors, such as reflectivity, general absorption by surroundings and possibly most importantly preventing the possibility of harm to human eyes. The human eye is capable of seeing all so-called "Visible light" which means light of wavelengths ranging from 380 to 780 nanometers (Schott AG, 2020). To ensure that the Scanning process is not annoying or even hurtful for humans most LiDAR scanners tend to scan with a wavelength that is in the Near-Infrared (NIR) region of the colour spectrum with wavelengths between 800 nm and 2500 nm, although often the smaller wavelengths are taken as the larger the wavelength the more energy the light has, which can potentially cause damage to eyes. Light with a wavelength above 1400 nm can be absorbed by the cornea and lens of the eye after which it is converted to heat which can cause damage to the eyes(las), so even though some wavelengths higher than 1400 nm are still allowed to be used via the FDA eye-safety standard IEC 60825 (for instance some scanners from Velodyne are capable of measuring using 1550 nm wavelengths (VELODYNE LIDAR, 2018)) it is safer and thus more common to use lower wavelengths.

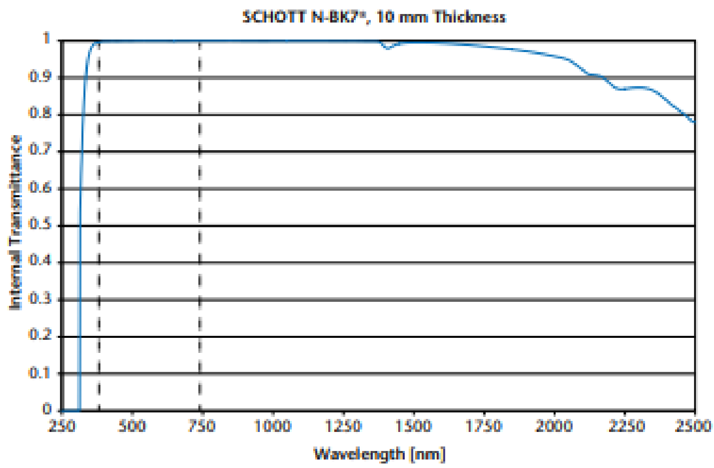

Figure 1: The internal transmittence of light in a 10mm thick SCHOTT N-BK7 glass plate. The blue line indicates the transmittance to waves with the wavelength on the x-axis and the dotted lines show the extent of the visible light spectrum. Acquired from (Schott AG, 2020)

This choice, although very much understandable, is the source of not being able to detect certain materials when using LiDAR as these materials are tricky to identify using light at these wavelengths. In this research, a look is taken at one of these materials which is very prevalent in almost all buildings that exist nowadays, namely glass. Glass is has a fairly unique property in that it is transparent for all visible light and a large part of the infrared spectrum. How the type of glass behaves does depend on the type of glass and its condition but to illustrate a look is taken into the internal transmittance of the glass type SCHOTT N-BK7 in Figure 1, where a transmittance of 100% can be seen for visible light and NIR.

It does however not only transmit laser but can also reflect it directly back, reflect it under an angle or absorb it. The extent to which this happens once again all depends on the wavelength of the laser and the condition and type of glass, but for the wavelengths most lasers use when compared to more standard glass panes, most of the light is transmitted through the glass, a small portion is reflected and an even smaller portion is absorbed.

When receiving returns from laser that interacted with a glass pane, the points are then found mostly in three different ways. In Figure 2 we can see a return that directly reflected from the glass surface, although this return does not have the same intensity as the laser which was sent out due to absorbance and transmittance of the laser. This return can only happen when the laser is fired in a very perpendicular angle to the glass surface, as otherwise the return would not find its way back to the sensor in the scanner on its own.

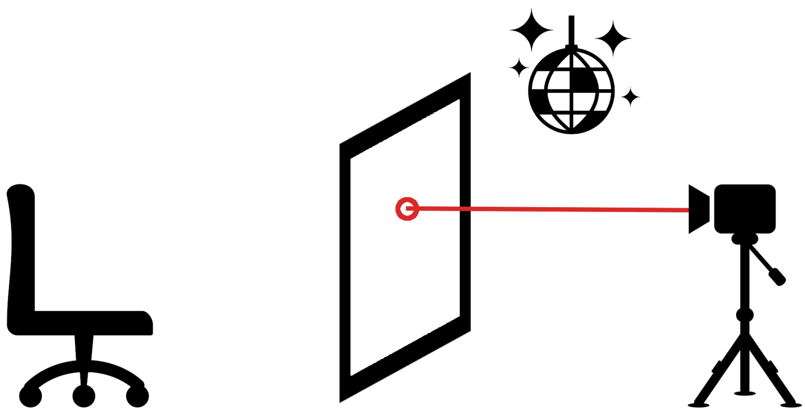

Figure 2: The laser partially gets reflected directly back to the scanner from the glass pane resulting in a point on the glass pane itself

It is, however, possible for a laser to return to the scanner via another object which is shown in the Figure 3. The laser in this case bounces from the glass pane to the disco ball, which would reflect it to the mirror and then again to the scanner, losing intensity with each reflection. When the laser has returned to the scanner it however only measures how long the laser was on its way, assuming it moved in a straight line. This, in the end, results in a false second disco ball forming behind the glass. The point registered in the result will also not be positioned in the red circle in Figure 3 but in the blue circle.

The last of the common ways points can be found after interacting is when the laser simply is transmitted by the glass. This is the most common option of the three and the one with the least intensity lost as most of the laser goes through glass without much issues. The result as shown in Figure 4 is a point as expected on the object behind the glass pane almost as if it was not even there.

As the last of the three options is by far the more common case, glass became a material that is often not captured in point clouds even-though office buildings nowadays have plenty

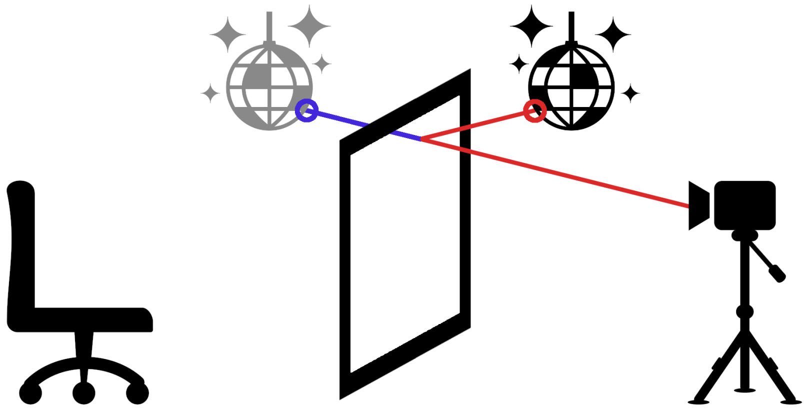

Figure 3: The laser (in red) partially gets reflected by the glass and hits a disco ball. This reflects the laser back to the scanner via the window. The scanner however measures only the time the laser spent before returning, assuming that the blue path is taken when reflecting on the glass falsely creating a second disco ball behind the glass

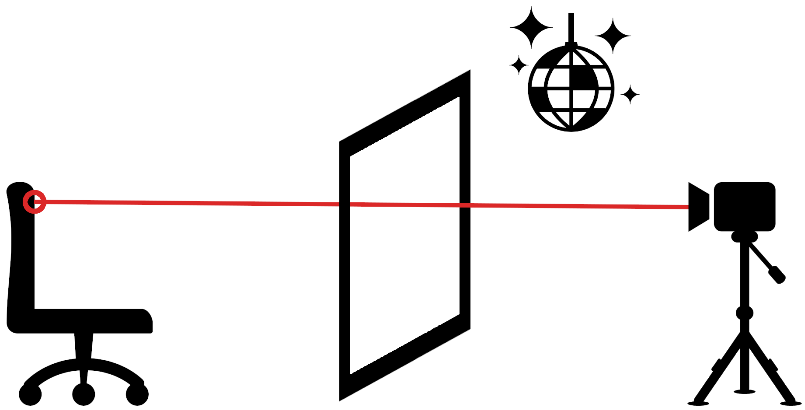

Figure 4: The laser goes through the glass with almost no interference and registers a point location on the chair behind it without problems.

of large windows, glass doors and even full glass walls. Should these point clouds be used for planning evacuation routes or as support for first responders, this lack of information could present routes in places where no such passage is available costing valuable time in life-threatening situations. A similar lack of information would also be prevalent, when remodelling the building as some bigger furniture pieces could not be transferred due to unforeseen lack of space or when buying glass cleaning supplies in bulk as some parts of the building are neglected as they were not properly inventoried.

1.1 Objective

What I aim to achieve in this Master Thesis will be to detect or deduce the presence of glass in point cloud data captured using a LiDAR sensor to enhance the value of indoor point cloud data without the need of a large and expensive extra setup.

This is done in conjunction with the company CGI, which over the last few years have had a contribution in the work of multiple Theses to capture indoor semantics in a better understandable way. One of the first steps in the road towards a semantically fully understandable indoor environment has been performed by Staats et al. (2017), where point cloud was analyzed to find the navigable space inside. This includes the classification of stairs, slopes and horizontal surfaces, so routes could be planned for differently mobile individuals. This research was then continued by Flikweert et al. (2019), who looked into the connectivity of different indoor spaces and finding the spots where these indoor spaces connect (e.g. doors). The method proposed in this research for finding doors uses properties of openings in point cloud facades and requires the user of the MLS device to walk through openings for 100% certainty that the door found is not permanently closed or a glass pane. This is where my research will continue with the project as by finding and placing the position of said glass panes in point cloud the need for certain extra steps or restrictions in similar methodologies would be eliminated as there would be certain of whether the opening in the point cloud would be navigable or not. This extra information will enhance the total semantic understanding of indoor spaces and provide extra clarity in planning routes for indoor navigation using such a medium.

As can be seen, the result of incomplete data can cause problems in such situations making its connection to Geomatics. As a Geomatics student, I follow the Geomatic process from data capture to geo-management. This project mostly focusses on different parts of the Geoinformation chain (see Figure 5) as it looks at how the data is captured by the LiDAR sensors and how this can be improved to show the presence of glass. This data is then stored where seemingly irrelevant data for other projects can be of the greatest importance for this task as low-intensity readings and reflected points can hold great value for the result, which has as a consequence that the total dataset is larger and therefore needs to be efficiently handled. The analysis of the data is probably the most important part of the project as by analysis I hope to be able to place the position of glass panes into the point cloud. Points that are found to be directly on the glass pane would then be classified as such and by using deduction other points of the glass plane that have not been captured properly can still be approximated finishing the obstruction of space. Using analysis, other low-intensity data such as reflections and points on the outside of the building captured through glass could also be classified as such so they could be removed at a later stage should this be wanted. Reflections that can be matched to the original data could furthermore also be used to deduce the position of other glass panes that caused this reflection, showing just how much analysis of the data can bring for this project. Finally, this data will then need to be visualized to show the results intuitively and the quality of this result will need to be tested as well, meaning this project will go over almost the entire Geoinformation chain from start to finish.

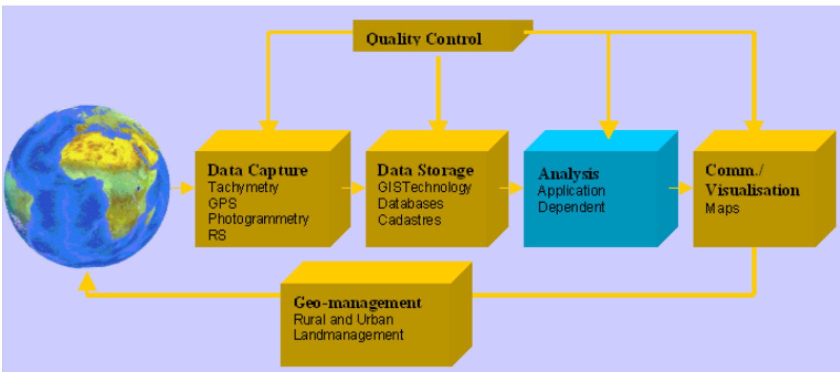

Figure 5: The Geoinformation chain which contains the general overview of all steps involved in a geomatics project. Acquired from the course GEO1001.

2 Related work

As glass has been a prevalent problem when capturing the indoor scene for quite some time, my research is not the only research that has been done in trying to find its presence. One of the most informative works I have found on the topic as of now is the work of (Ye et al., 2015). In their related work section, they distinguished the topic of detecting glass in 4 major categories, which coincide with my own findings.

1. **Physical manipulation:** The phenomenon of physical manipulation is possibly one of the more straightforward approaches that can be taken when trying to identify glass. It involves changing the geometries in the real would to improve their detectability, which in the case of glass could involve putting paint on the surface, placing a thin cardboard layer or using a spray to coat the surface which would be detectable with the scanner. While this approach may be valid it does have its downsides as physical manipulation is not always possible or desired, costs time and resources and also cuts off data which was originally in view by completely blocking the signal at the glass surface.

2. **Active illumination:** These methods involve illuminating the glass surface with special illumination methods, such as structured light or coded illumination. By the analyzing distortions in the patterns, these methods output the position of the distortion source (i.e. glass) can be retrieved. The main downsides that active illumination approaches have is that they are often very dedicated in setup. This means that a specific setup to show a specific property is needed to get results, which hinders the scalability of these approaches and requires prior knowledge of the scene that would be captured. Even if these conditions are met it can also still be that case that scanning larger buildings will still prove rather difficult due to lower efficiency when working with large data.

3. **Passive methods:** Passive methods can directly generate 3D reconstructions from captured imagery or other data, without needing an active additional interaction with the scene, unlike physical manipulation and active illumination. For example, mirror shapes can be reconstructed by observing the distortions of known patterns inside the scene, provided that the mirrors can refract light significantly. An upside of this method is that no extra actions have to be taken to get a result, making it cost-efficient. A downside of this is, however, that a lot of prior knowledge of the scene is needed to calculate the results and that the scene needs to be suited for this kind of approach, to begin with as this method is fundamentally dependent on the condition of the scene.

4. **Sensor fusion:** The last of the categories is sensor fusion which quite literally means the combining of multiple sensors to see the world more accurately. Lots of different sensors have been combined with LiDAR data over the last few years to enhance it. For example in their own work Ye et al. (2015) looked into the addition of an ultrasonic sensor to a Microsoft Kinect camera to give the depth camera an additional reference to range with (e.g. sonar) which is capable of seeing glass. Other researchers have also done similar approaches of combining LiDAR data with Sonar data to navigate robots indoors in an office building with glass walls (Wei et al., 2018). By using this approach to enhance the original LiDAR data, transparent object can be placed in the scene removing the weakness of the first approach. A downside of this approach is however the cost of an additional sensor and the need to make software that is capable of joining the datasets and decide which is the right one.

Aside from pure glass (or transparent object) detection, there have also been some more methodological approaches for detecting glass windows that have appeared over the years.

Windows in building often have the attribute that they are placed regularly, with repeating patterns and similar (often rectangular) shapes. In a synthesis project from previous Geomatics students, Kaniouras et al. (2019) worked out a method for detecting holes in a point cloud surface which then presumably should be the place where windows appear in the side of buildings. In this way they turned a lack of information into an indication of new information, although detecting whether there was glass present versus there being a literal hole was hard to do in this case.

A somewhat similar approach was also performed by Tuttas and Stilla (2011) but instead of inspecting holes in facades in point cloud data, they assumed the repeating structure of windows in these facades and therefore looked at the points behind the facade where the indoor points captured behind the transparent windows would be captured and possibly missing windows could be included by filling in missing parts of the pattern. In this case they are more robust to closed windows are sparse input data with the trade-off of being more dependent on the structure of the buildings itself as if there is no pattern to see in the buildings then the approach would not work.

Approaches that take into account a different property of windows have also emerged recently. As windows transmit light and energy it can let this energy into the building but also is a spot where the energy can escape from the building. This property of windows and other openings in facades was used by Maas et al. (2020) to detect openings as the spots where the most heat left the building. Using a supervised machine learning approaches these spots in the input thermal images were then classified showing that thermal information can also be a solid approach.

3 Research questions

For my research, I plan to find the position of glass in LiDAR point cloud and saving this position back into the same point cloud to enhance it with further information. To make this research a step forward from previously done research I am focusing on an alley that has not been explored as much before, namely using only the information from the scanning medium to approximate the position of glass in the point cloud acquired for said medium. This approximation will then be turned into 3D point positions to be stored in the input point cloud itself, enhancing its earlier version with new information. To concretely encapsulate this I formulated the following research question:

How can the position of glass be detected/deduced as 3D points using only information acquired from the scanning medium?

To answer this question and achieve the proper results the following research questions will also have to be answered:

- *What properties are required from the scanning medium to be able to detect glass?*
- *What characteristics does data that interacted with glass have to make them distinguishable?*
- *How reliable/accurate is the deduction of the position of glass from 3D point cloud data?*
- *How can different types of points that interacted with glass be distinguished from each other?* (Different types are for instance points behind glass, points on glass and reflected points)
- *How can reflected points best be matched to the original data?*

3.1 Scope

For the scope of the project, I aim to do the following things as shown using the MoSCoW[1] method. The MoSCoW method is used to indicate the priority for tasks in four different categories, which here represent the scope of the project as well for me. The Must-haves are of vital importance to the solution and have to be finished for the project to be a success. The Should haves are properties that are wanted for a proper end result but should these not be completed in the end then there is still a result achieved with value for the scene. The Could haves are interesting angles for the project to also look at, but can be scrapped without leaving a lot of impact on the main result should time be limited. The Wont haves are a hard scope limitation for the project where it might be nice to have in the end, but for which it has been decided that this area will not be explored in the time frame of the project.

The MoSCoW priorities I set out for my thesis are shown in Table 1. All points in the Must and Should haves have made it into my Methodology for the project as I strive to complete all of these at least in the project. The Could have I have placed on the test case with visualization spray is also a nice addition to said methodology, but would mostly serve as a separate test case to compare results with and is dependent on whether this spray can be brought along on the second day of capturing data, hence its placement as it is of less importance to the overall success of the project. As for the Wont have of not trying a sensor fusion approach this is decided for 2 main reasons, one being the pricing and management issues to be able to realize such a sensor fusion on short notice and two being the innovation of the research I want to perform. As shown in section 2 sensor fusion to find glass has already been performed in different ways, so what could be innovated here is taking such an approach and placing the data created back in point clouds to enhance them, but in my opinion this is not enough innovation for a Master Thesis. This is why I decided that this approach will not be pursued in this project.

[1]https://www.agilebusiness.org/content/moscow-prioritisation

Must Haves	Should Haves
* A way to deduct glass pane positions * A way to detect points positioned on glass * A way to recognize reflections * An algorithm that matches reflections with other data points * A way to label points as glass/reflection * A result set that can be visualized to show the differences * Documentation on which properties can be exploited to deduct/detect glass using the scanning medium	* A way to create points that represent glass in the point cloud * A way to evaluate the accuracy of the glass placement * A way to recognize outside objects as separate entities and label them * Documentation on how outside objects can be distinguished from reflections and points positioned on glass
Could Haves	**Wont Haves**
* A test case where glass visualization spray is used before scanning	* An integration with extra sensors in a sensor fusion approach

Table 1: The MoSCoW priorities set for the scope of the project. Note here that the intent is to complete at least both the Must and Should haves during the project. The Could haves are something I hope to be able to get to should time and resources allow it. The Wont haves here are out of scope and will not be pursued in this project due to time constraints.

4 Methodology

At the time of writing this report, the first steps of the research have already been set in motion for the project. On December 15th a test was held in conjunction with Leica who CGI and the TU Delft helped me get in contact with. During this test, we captured data in the orange hall at the faculty of Architecture and the Built Environment of the TU Delft, with the Leica BLK360, Leica BLK2GO and Leica RTC360 scanners. In this experiment, some interesting information was found which made its way into the methodology that is to be used for this project, as we found that the scanners were able to register glass directly, as well as capture reflections via the glass, as can be seen in Figure 6a and Figure 6b, by simply following the regular scanning procedure. These points were found in the lower intensity parts of the data as interaction with glass causes the signal to not fully return and are thus usually filtered out, as users often assume these points to be noise or inaccurate readings which can happen when a scan hits an edge or corner of an object. However, in this project, this is of greatest importance.

The methodology to be used for this project will therefore be the following pipeline as shown in Figure 7. The individual parts of the methodology will be further discussed in the following subsections.

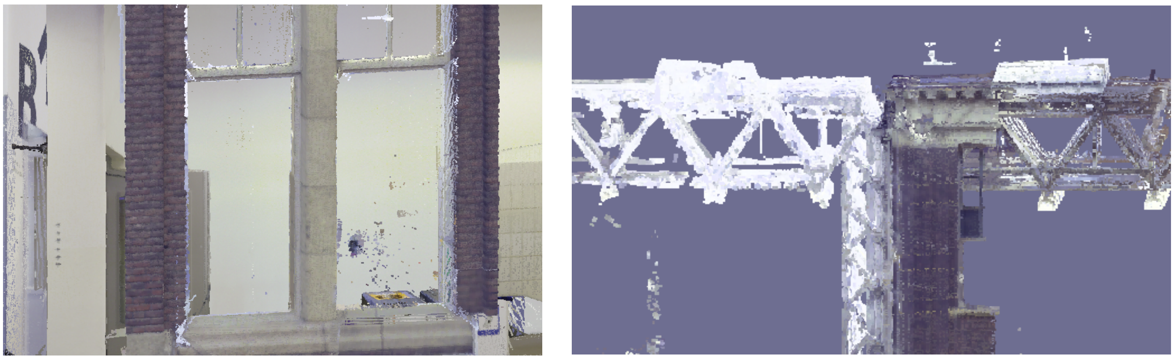

(a) The points visible in the window are reflected points that the laser registered

(b) The points visible in white on the left are reflections of the colored points on the right

Figure 6: Cases that are interesting in approximating the position of glass in the scene

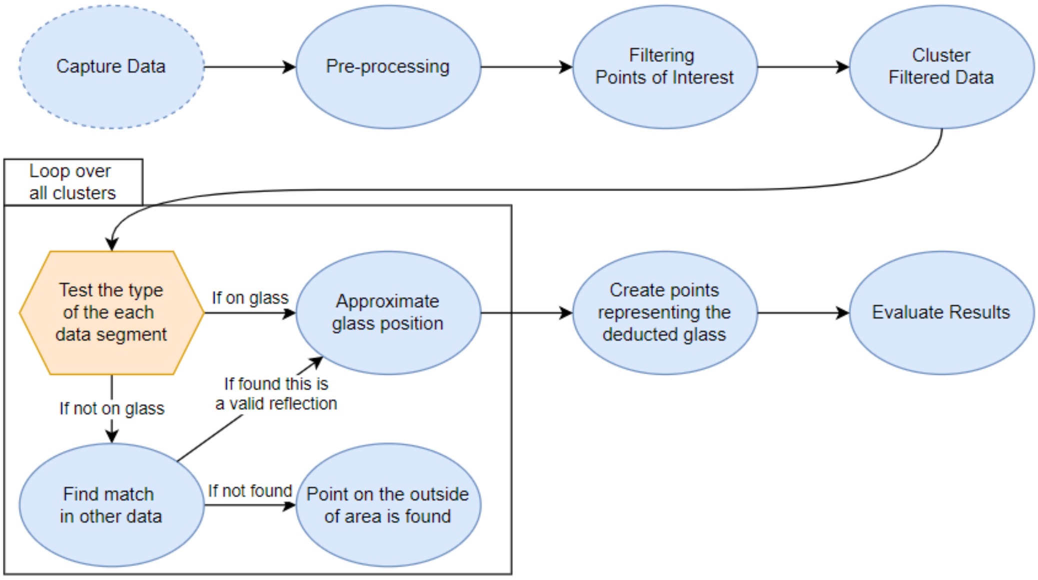

Figure 7: Methodology pipeline for the project

4.1 Capture Data

The first step of the methodology is acquiring data for the research itself. As the data of interest is often removed from point clouds to save space this is something that needed to be done once more without removing the low-intensity points so a suitable test can be run. Should dataset be delivered with all captured results then this step can be skipped and that data can be used as input for the next step. When capturing data we found that the points that were directly visible on glass only appeared on windows that were (almost) perpendicular to the scan position, most likely due to the way the laser refracted on the window surface. Should it be wanted to directly deduce the position of glass windows from the scanned data then it might be necessary to have a scan position perpendicular to these windows to find proper data for the approximation.

4.2 Pre-Processing

As the datasets are fairly large to handle and a lot of the data in the scene can be ignored in parts of the project, pre-processing can be applied in the form of temporarily removing points outside of the region of interest, thereby limiting the number of points that need to be processed. In addition to this, there might also be properties of points that are in the end not that relevant such as the numerical identifier of the points. These values will be removed to keep the data that needs to be processed as small as possible.

4.3 Filtering and Clustering

After the data has been made ready for evaluation we start by filtering the points of interest. This is done by creating a distinction between the directly captured points and the points that have interacted with glass. As glass does not properly reflect all of the signal, as it is partly transmitted, partly absorbed and partly reflected in another direction, the signals that came into contact with glass have a lower intensity than directly captured points (Wei et al., 2018). By exploiting this property an initial filter can be made to get closer to the points of interest, but there are more reasons that points can have a lower intensity such as points on the corners of objects so other filters might also be needed. After the unwanted points are filtered out the data are then clustered into groups that are positioned close to each other for inspection as a neighbourhood in the next steps.

4.4 Loop over all clusters

When the clusters have been constructed, a loop is made over all of them and it is determined whether the cluster is thought to be part of a glass surface directly or if the cluster is detected via the glass making it a possible reflection of a part elsewhere in the dataset or points behind the glass registered because of its transmissive property. This determination is done based on properties of the neighbourhood itself, as a glass window for instance has a vertically oriented planar neighbourhood, whereas the other objects can have a lot more diverse shapes. If the cluster is thought to be a possible reflection then its match needs to be found in the other data by comparing the neighbourhood with it. Should a good enough match appear between these, then the reflection is considered a valid one, but if this match does not occur then the cluster is assumed to be an object registered on the outside of the scanned area or noise. All clusters that have gone through this test are classified in the process and are labelled accordingly, to enhance its semantic value for later users of the data. Valid reflections and clusters assumed to be parts of windows are then used to approximate the position of the glass. For the planar window parts, we simply expand the plane to where more intensely registered points occur that indicate the edge of the window. For the reflections, we approximate the position of the glass via the position of the scanner, the position of the real points and the position of the reflected points to also approximate a plane of glass.

4.5 Create points for glass

After the planes of glass have been deducted and the missing areas have been approximated, each plane needs to be properly saved as points in the point cloud to enhance the original format and allow for a visible obstruction of space during visualization. These points can be extracted by taking several points that exist on the rectangle that has been approximated and taking these as the points representing the glass window. This process should match the density of the original point cloud and the points found in this process should be labelled as having been generated to distinguish between the original and new points, meaning all points in the point cloud will get a new binary value after this process. Other values of the generated points also have to be replicated, but do not have any real meaning aside from making the data complete.

4.6 Evaluate results

The first evaluation of the results will be done by manual inspection of the scene itself. In the test scene, the position of glass is easily verifiable for a human so doing so is the easiest first step. However, to see if the result is as intended for the whole scene, a separate test set needs to be made to evaluate against. This will probably be done by making a ground truth 3D model of the scene with glass added in the correct places and testing if the new points are within an error margin from the ground truth. From this, a final accuracy metric can then be calculated for the placing of the points to evaluate the results.

5 Time planning

My time planning focuses mostly on my methodology and the deadlines that come with the P's as can be seen in Figure 8. Note that the second day of capturing data is not added as this date is not set in stone yet although somewhere is February is expected. Performing extra research and writing of the thesis is also not added as this is something that will be integrated in my weekly schedule and felt redundant to put in for every week. Now to go over the Gantt Chart. My methodology is a fairly linear one in its execution, so I made the plan for it similar in linearly going through each task. The task of matching the reflected data gets a week more time as I expect this to be one of the biggest challenges in the project. I planned the methodology with quite long timespans per task to allow me to try different approaches to know which is the most efficient one, read up on possible methods for each task, write parts for the thesis when it is finished and test the results in intermediate stages. If all goes well the final version should be fully functional at the start of April leaving more than a month for additional testing, evaluation of the results and writing a paper to go along with the research.

The P's with exception of P2 have all been indicated with a range instead of a date since these are not final yet. To allow for sufficient time of preparation for the upcoming P's I planned 2 weeks of preparation with P5 getting even more time as I plan to have finished most other tasks by P4.

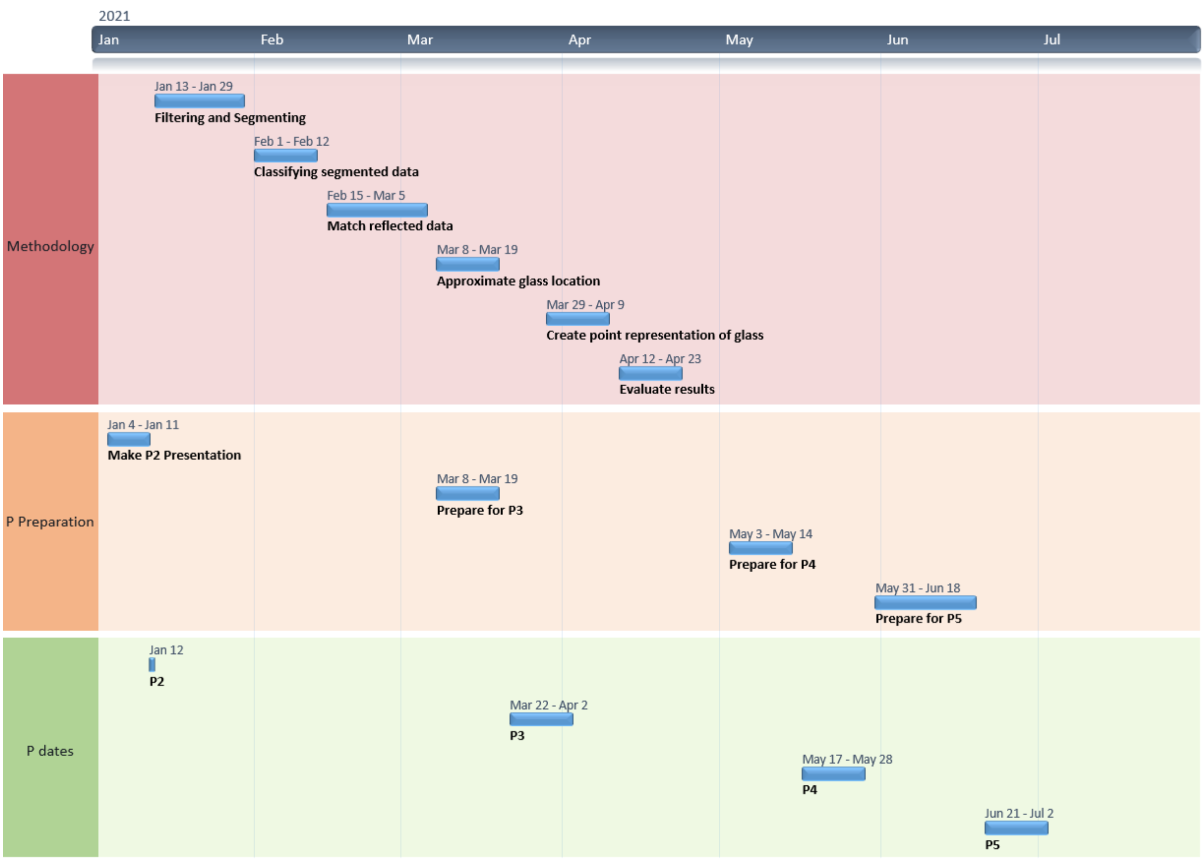

Figure 8: Gantt Chart for the rest of the project

6 Tools and datasets used

The data that is used for this project is from a few different sources. For initial testing of approaches TU Delft and CGI have both provided anonymized LiDAR datasets. These datasets included the orange hall in the faculty of Architecture and the Built environment, rooms in the CGI main office and data of the office of Veiligheidsregio Rotterdam-Rijnmond. The provided datasets might also be used as extra test cases in the end but might prove unsuited for the presented methodology due to the removal of low-intensity data.

For the main testing of my approach, the data will therefore be captured in the testing day with Leica that already happened on December 15th 2020 and the day of capturing data in the Erasmus Hospital that is scheduled to be in about a month from now. This data has parameters that are fully known to me for a better understanding of the result and that is confirmed to have low-intensity points which are the foundation of the methodology.

Aside from these datasets I might also take a look at testing my approach on more common benchmark datasets such as the datasets from ISPRS found on `https://www2.isprs.org/commissions/comm4/wg5/dataset/`, should these have data that is compatible with the methodology to have a better-known dataset in my result sets.

As for the tools I will be using, these are split up in software tools and hardware tools. For hardware tools, I am mostly dependent on dedicated data capture days where I have an expert use the scanners to capture the environment and provide me with access to the data. Some of the scanners I know I will get/have acquired data from are the Leica BLK360, Leica RTC 360, Leica BLK2GO, Leica Pegasus Backpack, ZEB REVO and ZEB HORIZON.

As for the software I will be using, most of the project will be investigating point cloud data and modifying it. For viewing the data in a 3D setting I am using CloudCompare as a free Open Source option for investigating the data, which is also able to deal with both las and e57 files which I both have access to. With the data I got from the test day with Leica I also got files in lgs format, which is their own filesystem. To be able to access this data and work with it I therefore also got access to the JetStream Viewer application, which Leica developed themselves as their home file viewer.

The language I use for the coding in this project is Python, for which libraries such as Laspy are used to load the data from the point clouds. Python has powerful libraries to its disposal such as the numpy and sklearn packages for quick analysis of point neighbourhood data, making it a suitable choice for this project. Should performance ever become an issue then other options can be considered but as a lot of the Python libraries use C++ on the back end and this code is not needed in real time this will probably not be the case.

7 Meetings

During the project contact with the mentors involved and CGI has been of great importance to me, for getting in contact with experts on this topic, enabling opportunities for capturing data and brainstorming with other students that are graduating at CGI for fresh looks at the project. This is contact has been done and will continue for the rest of the project, in the form of a weekly meeting of half an hour with the mentors to keep everyone involved up to date. In addition to this I have a meeting with the students graduating at CGI on Monday, Wednesday and Friday to conclude the working day and to still get the feeling of working with a company during this thesis. All of these meetings are kept digitally until further notice due to the pandemic and as such my workplace is also at home for the project except for days we capture data in the field.

References

Laser safety – module 4: Bioeffects. URL https://www2.lbl.gov/ehs/training/webcourses/EHS0302/story_content/external_files/Module%204%20Bioeffects.pdf.

P. Flikweert, R. Peters, L. Díaz-Vilariño, and R. Voûte. Automatic extraction of an indoorgml navigation graph from an indoor point cloud. 2019. URL http://resolver.tudelft.nl/uuid:b11f5b57-5362-4b45-bed6-d5bc154d86aa.

P. Kaniouras, M. Moscholaki, J. van Liempt, K. Jarocki, L. Zhang, E. Verbree, and M. Meijers. Direct analysis on point clouds: Geomatics synthesis project 2019. 2019. URL http://resolver.tudelft.nl/uuid:4f610e66-01b5-409a-aaf9-9f03d96c7889.

V. V. Lehtola, H. Kaartinen, A. Nüchter, R. Kaijaluoto, A. Kukko, P. Litkey, ..., and M. Kurkela. Comparison of the selected state-of-the-art 3d indoor scanning and point cloud generation methods. In *Remote Sensing*, volume 9, 2017.

H. G. Maas, D. Lin, and M. Jarząbek-Rychard. Supervised detection of façade openings in 3d point clouds with thermal attributes. In *Remote Sensing*, volume 12, 2020.

S. Nikoohemat, A. Diakité, S. Zlatanova, and G. Vosselman. Indoor 3d modeling and flexible space subdivision from point clouds. In *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, volume 4, 2019.

Schott AG. Transmittance of optical glass. 2020. URL https://www.schott.com/d/advanced_optics/5b1f5065-0587-4b3f-8fc7-e508b5348012/1.1/schott-tie-35-transmittance-of-optical-glass-february-2020-row-20022020.pdf.

B. Staats, S. Zlatanova, A. Diakité, and R. Voûte. Identification of walkable space in a voxel model, derived from a point cloud and its corresponding trajectory. 2017. URL http://resolver.tudelft.nl/uuid:6a827a88-ce09-43f1-adb9-da886448a1fc.

S. Tuttas and U. Stilla. Window detection in sparse point clouds using indoor points. In *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume 38, 2011.

VELODYNE LIDAR. A guide to lidar wavelengths. 2018. URL https://velodynelidar.com/blog/guide-to-lidar-wavelengths/.

H. Wei, X. Li, Y. Shi, B. You, and Y. Xu. Fusing sonars and lrf data to glass detection for robotics navigation. In *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2018.

M. Ye, Y. Zhang, R. Yang, and D. Manocha. 3d reconstruction in the presence of glasses by acoustic and stereo fusion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015.