Best Practices for Indoor 3D Mapping

Junaid Khan, Thomas Watson, Mazharul Hossain, Eddie Jacobs, Lan Wang

Map901 Project, University of Memphis

We have developed these practices based on our experiences operating our 3D LiDAR scanner: the GVI LiBackpack and Insta360 PRO 2 camera. We believe these practices are generally applicable to every mobile 3D scanner, but the performance of any scanner will vary widely depending on its technology, software, and the space it is scanning. Nevertheless, you can use these practices to help identify any issues you may experience with your scanner, understand their cause, and investigate strategies to avoid them.

Every mobile 3D scanner today relies on a Simultaneous Localization and Mapping (SLAM) algorithm to build the scan. The scanner estimates its path through space by observing how the scene it is scanning moves around it, much like humans. For example, you, and the scanner, can tell that you are moving closer to an object when walking because it appears bigger. A correct path estimation is the single most important factor in the resulting scan quality, as this is how the scanner places each piece of the scene into the overall picture.

Multi-lens cameras (a requirement for 360°; our Insta360 PRO 2 has six lenses), or scanners with multiple cameras, must integrate (or “stitch”) the image from each lens into a single coherent view. Generally, they also do this by observing how the scene moves. Stitching problems can result in a severely distorted image, which negatively impacts the coloring of the scan.

A. Data Collection - General Quality Issues

- Moving objects/humans:

  If an object moves during the scan, it will appear multiple times because the scanner will see a new object each time it observes it. Reconfiguring the scene during scanning by e.g. moving a chair or opening a door will cause duplicate e.g. chairs or doors to appear.

  Humans should not move around the scene (except for the scanner operator) when scanning. Before starting the scan, ensure the operator can navigate to all relevant areas without having to move objects, and that doors are opened or closed as desired (generally, opened is preferable). Take care to avoid bumping objects during scanning.

- Indoor usage and tight spaces:

  Limitations of the equipment cause operational difficulties when scanning indoor spaces. Tight spaces and small rooms only increase the difficulty. The common VLP-16 LiDAR sensor has a very poor ±15° vertical field of view. If the sensor is six feet above the ground, and level, then it can only see parts of the ground that are 22 or more feet away. Narrow corridors, close walls, and featureless rooms present extreme difficulty to the SLAM algorithm because it cannot see enough movement to effectively track its position.

We recommend that the operator stays about 5 feet away from walls when possible, or near the center of the room if not. This ensures the LiDAR is not so close to the near wall that the scanner is confused yet is far enough away from the far wall to capture a good view. In some cases, the scanner operator can lean side to side, or manually detach the sensors, to provide better coverage. Many of our scans simply lack a floor (and have very detailed ceilings) because the room was not large enough for the LiDAR to see the floor.

- Doors, Windows, and Mirrors:

  The LiDAR has a maximum range of hundreds of feet, so it can easily see into other rooms through open doors, or outside through uncovered windows. Though helpful for registration (discussed later), this can result in the scan containing far away garbage which is not relevant to the scene. The LiDAR's laser beam can be reflected by mirrors, and the LiDAR cannot determine that this has happened, so a mirror in the scene will create an unwanted duplicate scene behind it in the scan. This effect can also be caused to a lesser degree by shiny metal, glass (like a glass table), and other reflective surfaces.

  If the garbage is a concern (note that it can be easily removed during post-processing), consider covering windows and closing doors before scanning. Cover up any mirrors in the scene and consider covering other reflective surfaces.

- Camera Image Quality:

  Good lighting is important to the quality of the camera image. In dark rooms, the image will be low quality and noisy, and a multi-lens camera may be unable to properly stitch it. In rooms with significant variations in lighting, the camera's auto-exposure function will adjust image brightness and result in inconsistent coloring being applied to the scan; this is particularly obvious on walls and results in a splotchy appearance. Excessively bright lights, and the sun, can further harm coloring quality by causing lens flares.

  Ensure all available lights are turned on before starting the scan. Naturally dark rooms must be lit up with stationary light. Do not turn on or off lights while scanning. Consider covering windows if the sun is shining into the scene.

- Multi-lens Camera Image Distortion:

Objects closer to the camera are generally harder for the stitching algorithm to deal with. Things like close walls or wires on the camera body will cause significant distortion. Distortion usually shows as a seam in the image, though sometimes you may see a “water ripple” effect around problematic objects as the camera moves through the scene. In some cases, the distortion is more subtle, and objects will appear complete but be shown at exaggerated sizes or impossible perspectives.

Avoid getting close to walls. Ensure the camera has a clear view to the scene and that parts of the scanner (like the frame or cables) or other objects close to the camera cannot, under any circumstances, be seen by it. Objects directly above or below the camera, like the operator, are usually okay because these parts are not used to color the scan, and the highest distortion occurs at those points anyway. The camera may have a function to calibrate its stitching algorithm, but the calibration will be highly scene specific. Off-camera stitching software will always produce the best results.

- Maximize coverage

If the scene to be scanned is large enough, the operator can walk through it in a loop so that they cross over their path and traverse the same area multiple times. This assists the SLAM algorithm and can increase quality. The image below shows the path that might be taken by the operator through a large area. By walking it twice, the scene will be revisited, and the SLAM algorithm can improve its path estimates.

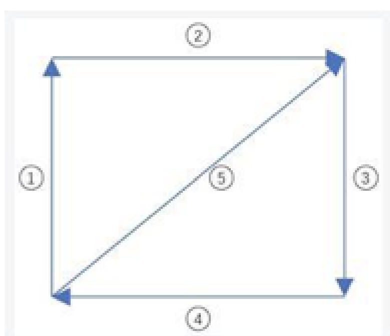

B. Data Collection - SLAM issues

- SLAM errors:

If the SLAM algorithm loses track of its path badly enough, the estimate will be completely wrong. This usually appears as multiple copies of the room superimposed onto each other with random angle or position offsets. This can happen if the SLAM algorithm cannot understand the scene (the space is small, narrow, or featureless) or if the algorithm mishandled a loop in the path (a loop is any place the path crosses over itself).

SLAM errors cannot be corrected once they have happened. The scan must be discarded, and the scanner must start over. Sometimes avoiding loops, or deliberately making loops, can reduce the chance of an error. Some spaces may simply be un-scannable if the algorithm can't make sense of them.

- SLAM drifts:

  Because it relies on accumulating all the movement it's seen so far, the SLAM algorithm's path estimation gets more and more inaccurate as the scan progresses. This usually appears as a room that is bent (so it is not quite level) or that is not quite square. Hallways and corridors are most susceptible.

  Consider scanning smaller areas of large spaces (with large overlap), then joining them together in postprocessing. Walking in loops can allow the algorithm to correct its estimation, but this may cause SLAM errors if there is too much drift. In some cases, the scan can be cut up and manually registered during postprocessing, but rescanning the area is usually easier.

- Walking speed:

  The SLAM algorithm needs lots of data to do its job most effectively. Walk slowly and methodically during scanning (about half normal walking speed) and avoid sudden movements to reduce the chance of SLAM errors.

- Doors and stairs:

  When the operator walks through a door or takes some stairs to another room, the SLAM algorithm cannot cope with the sudden change to unfamiliar scenery and a SLAM error is likely to result. Walking through the same door twice can create a loop that the algorithm is unable to handle.

  We recommend avoiding walking through doors or taking stairs that are less than 10 feet wide. Instead, scan the two spaces separately (and the stairs, if applicable) then join the scans together during postprocessing. If 10 feet wide or greater, so that the scanner can see a large part of the other room through the door or stairs, walk into the room through one side, then out through the other so that a loop is not created.

C. Data Postprocessing - Fusion

- Extrinsics and Parallax Error:

By using the sensor extrinsics, or position and rotation of the camera relative to the LiDAR, the scanner calculates which pixel of the camera is looking at each point captured by the LiDAR in order to determine the point's color. This calculation will produce incorrect results if the scanner doesn't know the correct extrinsics, or for very close objects where both sensors can't physically see the same point. For example, if the scanner is close to a cabinet, then the LiDAR may look over the top of it and see the wall while the camera will only see the cabinet. Inaccurate extrinsics will produce blurry or completely meaningless (if bad enough) colors.

It is vital that the scanner know its extrinsics accurately and that they do not change during a scan, i.e. that the camera is rigidly fixed to the LiDAR. Although it cannot be completely corrected, a common symptom of imperfect extrinsics is blue trees: the LiDAR sees the leaves, and the calculation pulls the color from the sky visible between them. For scanners with integrated cameras, inaccurate extrinsics are not likely to be a problem (and cannot easily be adjusted) so long as the scanner is not disassembled or damaged.

- Temporal Alignment:

In addition to knowing the right pixel to get a point's color from, the scanner also must determine the right frame to look at. If the scene changes between the capture of the point and its corresponding pixel, i.e. the scanner is moving, the calculation will be incorrect. This shouldn't ever be a problem for scanners with integrated cameras unless there is a software bug. However, if the camera and scanner are independent, accurately determining the time alignment between them is vital. Slightly incorrect alignment will produce blurry color and grossly incorrect alignment will produce meaningless color.

We found that the alignment must be within 0.05 seconds (i.e. the camera must capture a given point not more than 0.05s before or after the time the LiDAR captured that point) for the best result, as the scanner moves a lot. Our system uses the time that recording started on both the LiDAR and the camera to determine the alignment to within a couple seconds, which requires that both have a correct clock. It then achieves fine alignment by correlating the accelerometer data from both the LiDAR and camera. Because both are rigidly fixed, they experience the same movement, and the maximum correlation must be when they are aligned correctly. This works about 90% of the time, but the other 10% it produces a meaningless result and must be manually corrected. Ideally both could be fed accurate time from an external source, like a GPS, but our equipment did not have this feature.

D. Data Postprocessing - Registration

- Manual Registration:

  After all the rooms are scanned, they must be manually placed correctly in the overall building structure. Using software such as LiDAR360 or CloudCompare, two scans can be registered, or lined up with each other, by selecting three or more points in each scan you know to be in the same place. For example, you might use a doorknob as something that is in the same place, and easily identifiable, in both scans. After selecting the same point on the doorknob in both scans (and at least two other point pairs), the software will rotate and translate one scan so that all the selected pairs line up as much as possible, and by extension align the rest of the scan. If correctly registered, both scans together will now appear to be one complete scene.

  Registration accuracy can be improved by selecting more than three points and by selecting points that are far apart in each scan. For example, selecting three points on the same doorknob would produce a poor result. Each registration introduces error, so limit the "chain" as much as possible. For example, if you have several rooms and a hallway, register each room to the hallway. If you register the first room to the hallway, then the second to the first's wall, then the third to the second's wall, etc., the last room's registration will be impacted by the cumulative error of all the previous registrations.

- Scanning for Registration:

  The more overlap there is between scans, the easier registration will be for you. How the building is scanned can also have a significant impact on registration ease and quality. Scanning each room multiple times allows you to select the best quality or easiest to register scan. Knowing where each room belongs in the building (using e.g. room numbers and CAD drawings) makes it easier to lay out the rooms in preparation for registration. Taking the time to capture a high-quality scan of the hallways simplifies registering the rooms along them.

  Allow as much overlap as possible between scans. Consider placing an easy to identify object in a place that can be seen in both scans, like a box in a doorway. Keep a good record of which room is captured in each scan (we write down the current time, scan data filename, and room number at the start of every scan) and how they fit together.

- Automatic Registration:

  Both LiDAR360 and CloudCompare feature implementations of the ICP (Iterated Closest Point) registration algorithm. This algorithm is highly accurate, but it depends heavily on a large overlap in the scans (at least 50%, ideally 100%) for best results. It also requires that the scans are manually registered first. If there is insufficient overlap or a poor initial registration, it will produce a meaningless result. This algorithm works very well for registering multiple parts (or scans) of the same room. It still introduces error, so limiting chains is still important.