

Refined Non-Rigid Registration of a Panoramic Image Sequence to a LiDAR Point Cloud

Arjen Swart^{1,2}, Jonathan Broere¹, Remco Veltkamp², and Robby Tan²

¹ Cyclomedia Technology BV, Waardenburg, the Netherlands
{aswart, jbroere}@cyclomedia.com,

² Utrecht University, Utrecht, the Netherlands
{Remco.Veltkamp, robbytan}@cs.uu.nl

Abstract. The combination of LiDAR data with panoramic images could be of great benefit to many geo-related applications and processes such as measuring and map making. Although it is possible to record both LiDAR points and panoramic images at the same time, there are economic and practical advantages to separating the acquisition of both types of data. However, when LiDAR and image data is recorded separately, poor GPS reception in many urban areas will make registration between the data sets necessary. In this paper, we describe a method to register a sequence of panoramic images to a LiDAR point cloud using a non-rigid version of ICP that incorporates a bundle adjustment framework. The registration is then refined by integrating image-to-reflectance data SIFT correspondences into the bundle adjustment. We demonstrate the validity of this registration method by a comparison against ground truth data.

Keywords: Mobile Mapping, LiDAR, Multi sensor, Registration, Point Cloud, Panorama

1 Introduction

Panoramic images have become widely used by both governments and businesses for a variety of geo-related processes. Among the many applications are the assessment of permits and inspection of security risks, real estate valuation, and the tracking and managing of objects in the public space, such as traffic signs, lampposts, park benches and greenery. Panoramas are also being used for measuring and map-making. Many of these processes could be greatly simplified or enhanced by the addition of 3D information from a mobile LiDAR system.

There are several mobile mapping systems currently in use that record LiDAR points and color imagery at the same time. However, there are some advantages to separating image and LiDAR acquisition. In many countries, mornings and evenings in winter are too dark to produce high-quality panoramas. And of course large-scale photography at night is simply not feasible at all. LiDAR sensors on the other hand do not require an external light source to operate. Because LiDAR devices suitable for mobile mapping applications are very expensive,

they should be idle as little as possible. This is especially true if the goal is to cost-effectively acquire point clouds for very large areas of the country, rather than for several relatively small projects. But even for smaller projects, separate acquisition can be economically advantageous. In the Netherlands for example, annually renewed parallax-free panoramic images are already available for the entire country. Combining these with separately recorded LiDAR data could save the potentially costly integration of the LiDAR sensors with a parallax-free panoramic camera system.

Separate acquisition of panoramas and point clouds does mean that the positioning information from GPS/INS is acquired separately as well. Especially in environments with poor GPS reception, like areas with many trees, or urban canyons, there could be a significant shift between the two data sets. The direction and size of this shift can change over the course of the trajectory. It is therefore necessary to register the sequence of panoramic images to a LiDAR point cloud of the same environment in a non-rigid manner.

Algorithms for the registration of color images to a 3D model or point cloud can generally be statistics based or feature based. Statistics based algorithms attempt to use a statistical dependence between the image and the point cloud. Mastin et al. (2009) register aerial images to LiDAR data by maximizing the mutual information between the image and the grayscale encoded height. When present, they use LiDAR reflectance values as well. If multiple images are available, a 3D model can also be registered using the photo-consistency principle (Clarkson et al., 2001). When the model is correctly registered, 3D points are likely to have the same color in multiple images. This principle can be used to optimize a cost function.



Fig. 1. Part of a panorama projected onto the triangle mesh generated from a LiDAR point cloud, before and after registration.

Feature based algorithms find matching features between images and a point cloud. For urban environments, Liu and Stamos (2005) find corresponding rectangular parallelepipeds between range scans and images. They also make use of line parallelism and orthogonality to estimate camera parameters. Böhm and Becker (2007) have used the SIFT feature detector (Lowe, 2004) to find corre-

sponding interest points between the reflection image from a laser range scan and a photograph in order to register the two. Zhao et al. (2005) applied a dense 3D reconstruction technique to change the problem of image-to-point cloud registration into a 3D-to-3D registration. ICP (Chen and Medioni, 1991; Besl and McKay, 1992) is used to register a dense point cloud generated from a pair of key frames in an aerial video sequence to a 3D model or point cloud. The transformation computed by ICP for each pair of key-frames is then added as an observation in a bundle adjustment framework. The bundle adjustment interpolates the registration of the non-key frames and can partially correct large errors in the ICP registration. Li and Low (2009) use a semi-automatic method to find an initial registration of indoor photographs and laser range data. After the initial registration, they match 3D points from a sparse 3D reconstruction with the closest plane detected in the LiDAR cloud. These correspondences are added to a bundle adjustment framework, which is used to refine the registration in order to handle non-uniform spatial distortion.

Our main contribution in this paper is the application of non-rigid ICP to the problem of registering image and LiDAR trajectories by means of a bundle adjustment framework. A similar approach was taken by Lothe et al. (2009). However, we also integrate the use of feature point correspondences between images and LiDAR reflectance data into this framework. Results obtained by our method are verified by comparison to a ground truth data set.

High quality line or point correspondences between image and LiDAR data can be difficult to find in many real-world environments. This is especially so with mobile LiDAR systems, which don't reach the same spatial resolution possible with stationary scanners. Therefore, our system uses 3D-to-3D registration of a reconstructed sparse point cloud and the LiDAR data as its core. We build on the idea of Li and Low that the observations of reconstructed 3D points in the images can be used in bundle adjustment to constrain a non-rigid deformation of the sparse cloud itself. This applies especially to the type of data found in outdoor sequences. In a sparse 3D reconstruction from a long trajectory, the majority of the reconstructed points will only have been detected in several neighboring images. This causes the resulting network to be fairly weak and incapable of correcting errors that slowly build up over time. However, the connections are sufficiently strong to function as regularization. Due to IMU drift under sub-optimal GPS conditions, the non-rigid deformation of the trajectories can easily be too severe to handle with a single application of bundle adjustment. Therefore, we propose to wrap the bundle adjustment in a non-rigid ICP procedure. Each iteration, a rigid ICP step is performed first on successive segments of the trajectory in order to find the required correspondences. This is followed by the bundle adjustment stage, which updates the camera parameters and the sparse point cloud for the next iteration. The process is described in section 3.1.

After this registration, the now greatly improved initialization makes it easier to find more direct correspondences between the images and the LiDAR cloud. It is not necessary that these feature correspondences can solve the entire

registration problem by themselves. By adding the information to the bundle adjustment stage, they can contribute wherever possible. We believe this is the key to building a robust system for registration that can handle a wide variety of environments. In section 3.2, we describe our method for adding SIFT interest point matches to the procedure. The results of our experiments are presented in section 4. We provide a comparison of the output of the registration procedure to a ground truth data set constructed using manually selected correspondences for each panorama.

2 Data

Our main test site is an approximately 750m long trajectory in the Dutch city of Enschede. At this site, LiDAR data was recorded in December 2008 and panoramic images in September 2009. This means that roughly nine months passed between the dates that both systems recorded their data. During this time, many things changed in the scene, making the registration task more challenging. These changes included, but were not limited to, the relocation of small street furniture, temporary road/traffic signs, and the absence or presence of parked vehicles at various locations. Both data sets contain several moving objects such as cars and pedestrians. A Lynx mobile mapping system from Optech (Optech, 2007) was used to obtain LiDAR data of the area. The mapping system was equipped with two LiDAR scanners that can each produce up to one hundred thousand points per second.

Parallax free spherical panoramic images were obtained by a CycloMedia DCR7 recording system (van den Heuvel et al., 2008). Location and orientation data are available for each panorama from the integrated GPS/INS system.

The CycloMedia recording vehicle did not drive along the exact same trajectory as the Lynx system. Rather, a panorama trajectory of 147 recordings at 5 m intervals that approximately follows the LiDAR point cloud was pieced together from several tracks recorded at different times during the day.

3 Algorithms

3.1 Initial Registration

The goal of the initial registration is to compute a non-rigid deformation of the panorama trajectory so that, at least in the local area, each panorama fits well to the LiDAR point cloud. This is accomplished by a non-rigid Iterative Closest Point (ICP) process that includes both rigid ICP and bundle adjustment (BA). We use a Structure from Motion (SfM) technique based on SIFT interest point matching between consecutive panoramas to obtain a sparse reconstruction of the environment. It is this reconstructed point cloud that is initially registered to the LiDAR cloud. The positions and orientations of the panoramas in the trajectory are adjusted along with the sparse reconstruction. Each iteration of

the non-rigid registration procedure first uses rigid ICP on segments of the trajectory to find correspondences between the sparse cloud and the LiDAR data, and then a bundle adjustment stage to estimate the transformation. The process is initialized with the GPS/INS positions and continues until convergence is reached.

Rigid Iterative Closest Point We use rigid ICP to solve for rigid registrations of the two point clouds. Because the non-rigid changes along the trajectory can be large, the LiDAR cloud is split into several segments containing a fixed number of points. The sparse point cloud is then registered to each sequential segment. This makes it possible to find a better local fit with only a rigid transformation. Remaining errors will be handled in the next stage.

In every rigid ICP iteration, we find the closest corresponding LiDAR point for each 3D point in the sparse reconstruction using a Kd-Tree. To make ICP more robust against outliers, the only points that are used are those within a Relative Motion Threshold e_t (Pomerleau et al., 2010) that scales based on how well the solution is converging. A minimum threshold ϵ is added to e_t , in order to handle noise in the data, and to allow for residual errors due to the assumption of rigidity.

Approximately planar structures like building facades, pavement, traffic signs and bus stops are common in urban environments. This makes our problem suitable for minimizing the point-to-plane distance, which leads to improved convergence compared to the point-to-point distance. The required normal vectors of the LiDAR points are calculated beforehand based on their local neighborhood. Each LiDAR point can now be considered as a locally valid plane.

No closed form solution for point-to-plane based registration exists, so after having matched the points within the current threshold, the 6 rigid transformation parameters are computed with the Levenberg-Marquardt algorithm. For increased robustness, errors are re-weighted according to the Huber cost function (Huber, 1973). With this adjusted error function, outliers will have less influence on the solution.

When the sparse point cloud has been registered to all LiDAR segments, the computed rigid transformations are no longer necessary. Only the final point-to-plane matches are required for the next stage.

Bundle Adjustment Bundle adjustment is executed to solve for non-rigid deformations of the trajectory and to smooth transitions between the different ICP sections. It optimizes both the camera parameters and the SfM points by minimizing the cost function

$$K = K_{rpe} + K_{ptpe} + K_{ce}, \quad (1)$$

where K_{rpe} represents the summed re-projection errors, K_{ptpe} the point-to-plane correspondences as observations, and K_{ce} cameras as observations.

The re-projection error component is defined as

$$K_{rpe} = \sum_i \sum_j a \left\| f(\hat{\mathbf{C}}_i, \hat{\mathbf{X}}_j) - \mathbf{u}_{ij} \right\|^2, \quad (2)$$

where \mathbf{u}_{ij} is the observation of 3D scene point j in camera i , and $f(\hat{\mathbf{C}}_i, \hat{\mathbf{X}}_j)$ defines the projection of the scene point estimate $\hat{\mathbf{X}}_j$ in camera i using the estimated camera parameters $\hat{\mathbf{C}}_i$. The squared error is weighted by a . The re-projection error term preserves the relative motion between the panoramas as determined by SfM. Due to the sequential structure of the network, these links are strongest between panoramas that are close to each other in the sequence. This term will propagate information from the point-to-plane matches to the nearby area. It also serves as a smoothness constraint by allowing some deformation of the panorama trajectory in order to follow deformations in the LiDAR trajectory, but resisting large changes. The amount of deformation allowed depends on the weight a .

The point-to-plane correspondence term K_{ptpe} is added to the bundle adjustment as

$$K_{ptpe} = \sum_i b \left\| (\hat{\mathbf{X}}_i - \mathbf{P}_i) \cdot \mathbf{N}_i \right\|^2, \quad (3)$$

where $\hat{\mathbf{X}}_i$ is an estimated 3D scene point from the sparse reconstruction, \mathbf{P}_i a LiDAR point matched with scene point i during the rigid ICP stage, \mathbf{N}_i is the surface normal at point \mathbf{P}_i , and b the weight used for point-to-plane correspondences. This term ties the LiDAR and panorama trajectories together. If a mistake was made in creating a certain match by the rigid ICP stage, this will not necessarily lead to a large error in the cost function. As long as the LiDAR point was on the correct, or a similar plane, the scene point can slide along this plane with little or no extra cost. In the next iteration, ICP can then select a better match using the output of the bundle adjustment stage as the initialization.

The different camera positions and orientations from the GPS/INS based trajectory data have also been added as observations:

$$K_{ce} = \sum_i \sum_j c_{ij} b \left\| (\hat{\mathbf{C}}_{ij} - \mathbf{C}_{ij}) \right\|^2 \quad (4)$$

Here we sum over each camera \mathbf{C}_i , where c_{ij} is the weight for observed camera parameter \mathbf{C}_{ij} , and $\hat{\mathbf{C}}_{ij}$ is the estimate for parameter \mathbf{C}_{ij} . This term prevents the cameras from drifting too far from their original absolute positions. However, since the deformations in both the panorama and the LiDAR trajectories must be compensated for, the weights are set to a very low value.

It is likely that there will be outliers in the data, especially because of the significant changes in the environment in the time between the recording of both sets. The Huber cost function has been used to re-weight the residuals of the

re-projection and point-to-plane correspondence components. After performing the bundle adjustment once, the potential outliers exceeding a threshold are removed, and the procedure is repeated.

The bundle adjustment was implemented using a heavily modified version of the SBA sparse bundle adjustment software package (Lourakis and Argyros, 2009).

3.2 Refinement

Having improved the panorama positions using the initial registration algorithm, there is often still room for improvement by exploiting the reflectance information from the LiDAR data (see section 4 for quantitative examples). Our method for improved registration is based on the idea that after the initial registration, color and reflectance images as observed from the same position should be very similar.

However, while a vehicle mounted LiDAR system can produce reflectance data per point, it cannot directly produce a reflectance image that can be used to find matching feature points with a panorama image. It would of course be possible to project the points into an image, but simple Z-buffering would not be sufficient to handle occlusions, unless very large pixels were used. Since the density of the point cloud is not uniform, this would needlessly waste resolution in some parts of the image. Therefore, we generate a triangle mesh from the LiDAR point cloud using the method described by Carlberg et al. (2008). Intensity images can then be produced by rendering the mesh from the point of view of a virtual camera.




Fig. 2. Color and reflectance images created with the initial registration.

The SIFT feature detector is robust enough to handle some of the differences between the panorama images and LiDAR reflectance data. However, it works best when perspective deformation is minimal. For that reason, the positions from the initial registration are used to render both color and reflectance images from the same virtual camera position. If there were no mis-registration at all, the feature points derived from LiDAR reflectance and panorama data would have the same location. So if we assume the remaining registration error is small, like shown in Fig. 2, the search area for matching feature points can be limited.

We orient the virtual camera towards the ground, because the point density is highest near the trajectory of the recording vehicle. This is important for good localization of the features.

For each panorama, the matched points from the corresponding color and LiDAR reflectance images are transformed back into panorama-coordinates. Since the 3D positions of the LiDAR reflectance SIFT points are known, the result of each matching pair is a 2D point in the panorama matched to a 3D point from the LiDAR. With enough point correspondences, it would be possible to filter out erroneous matches and use a spatial resection algorithm. However, this would force us to throw away points when a panorama has less than three intensity matches, or if the points are in a configuration unsuitable for resection. Since many of the reflectance intensity SIFT points are found on road markings, it is certainly possible to only have matches on a single line. These points would not be enough to solve for the registration of the panorama by themselves, but combined with other information, they can still contribute to the solution. For this reason, the matches are only filtered by a simple distance threshold, and we rely on the next stage of the algorithm to remove outliers.

In order to combine all of this information in a single framework, we include the reflectance matches in the process described in section 3.1 as an extra term to Equation 1, as

$$K_{ipe} = \sum_i \sum_j d \left\| f(\hat{\mathbf{C}}_i, \hat{\mathbf{Y}}_j) - \mathbf{v}_{ij} \right\|^2 + \sum_i e \left\| \hat{\mathbf{Y}}_i - \mathbf{Y}_i \right\|^2, \quad (5)$$

where \mathbf{v}_{ij} is the SIFT point observed in the color image of camera i , corresponding to 3D point \mathbf{Y}_j as obtained from the reflectance image. This term adds the reflectance points to the bundle adjustment as control points. Like the points created by SfM, they are subject to filtering by the error detection procedures. The entire non-rigid ICP process is repeated from the start with the new matches.

4 Results

To validate our approach, we have compared the results of our algorithm with a ground truth data set. Because the panorama and LiDAR data sets were recorded at different times under moderately difficult GPS conditions, no actual ground truth was available. Therefore, we have generated a ground truth by manually selecting point correspondences between the LiDAR reflectance data and the panoramas. The registration was solved by using bundle adjustment with these correspondences as control points. Naturally, point-to-plane and reflectance matches were not included in this. To check the validity of the ground truth, the manual selection was performed twice, each time by a different person. The RMS error between the resulting camera position sequences was approximately 1.9 cm. Although the algorithm was run on the full trajectory, only the first part (about 250m) of the full trajectory contained enough clear LiDAR reflectance features to consistently find correspondences for manual registration. Therefore,

we can only visually confirm the results on the second part, and just the first part is shown in comparisons with the ground truth data. This causes a bias in the results favoring the refinement stage, since good LiDAR reflectance features can be used by SIFT as well. Nevertheless, a comparison of the results with ground truth data generated in this manner can reveal patterns that might be less obvious with only visual validation.

It should be noted that the LiDAR cloud itself is not deformed in the manual registration. Just like the automated registration algorithms, its goal is to find a good local fit. Registration errors for LiDAR points far from the position of a panorama are less visible and not as important.

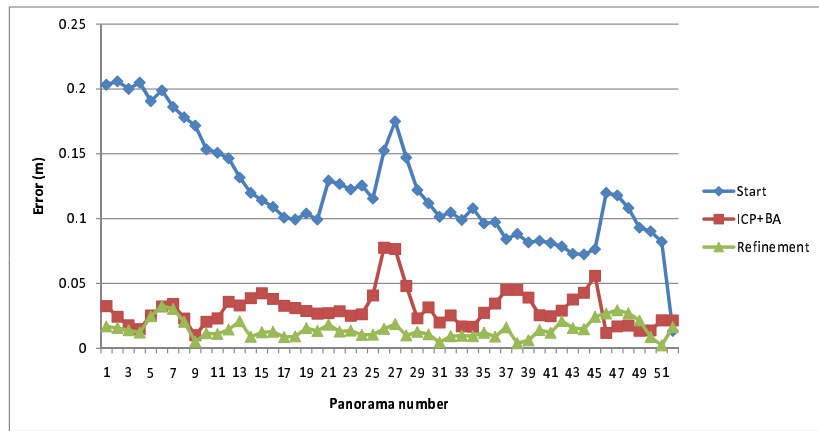


Fig. 3. Distance (m) to the ground truth of each panorama in the sequence.

In Fig. 3, we show the distance of successive panoramas in the trajectory to the ground truth. The initial registration with ICP and BA is already a large improvement over the starting positions from GPS/INS. Refinement brings the error down to about 2 cm. This is approximately the precision level of the ground truth, so a further reduction would not be meaningful.

Some of the peaks visible in the graphs are explained as follows. Between panoramas 26 and 27, and also 45 and 46, there is a gap in the data set. The reason is that the full panorama trajectory that follows the LiDAR trajectory is actually pieced together from different panorama sequences. Because panoramas closer than 5 m to another panorama are not stored by the recording system, there are often gaps at crossroads. When the recording vehicle comes in from an area with different GPS conditions, the error relative to the LiDAR will be different as well. The gaps cause a decrease in the number of SIFT matches between the panoramas, thereby weakening the connection just where it is needed most. With enough direct matches between the panoramas and the LiDAR reflectance in the area near the gaps, the inter-panorama links become less important.

A displacement of just 15 cm can already have a significant impact on the usability of the data. But in practice, it is quite possible to have much larger differences. We have simulated bad GPS conditions by adding an artificial additional error to the initial positions in the panorama trajectory. The added error gradually increases (and later decreases) with time in order to simulate IMU drift. All processing, including the structure from motion estimation was repeated on the new data. The results of registration with an artificial initial er-




Fig. 4. Distance (m) to the ground truth of each panorama in the sequence. The starting positions have been given an additional, increasing offset (following a gaussian curve) in order to simulate drift in adverse GPS conditions..

ror are shown in Fig. 4. We can see that the registration procedures are capable of correcting quite large non-rigid deformations in the panorama-trajectory. Fig. 5 illustrates the results of the registration process on one part of the scene. Our method has also been tested on several other data sets, but due to the difficulty of obtaining manual correspondences, only visual confirmation of the registrations could be made. Fig. 1 shows a small part of the results on a triangulated



Fig. 5. A panorama onto the triangle mesh, before and after registration.

point cloud in Ghent (Belgium). The results on the Ghent data set appeared to be slightly worse than on the Enschede set, possibly due to the fact that many LiDAR points on the ground were missing.

5 Conclusions and Future Work

The results presented in this paper show that the automated registration of separately acquired panorama and LiDAR trajectories is feasible. This was achieved using a novel combination of non-rigid ICP and panorama-to-reflectance data interest point matching. The bundle adjustment framework used can be expanded to include other types of correspondences.

As noted in section 3.1, only the panorama-trajectory is currently updated by our algorithm. It can be expected that the quality of the LiDAR trajectory will often be higher than that of the panorama trajectory. Environments with challenging GPS reception conditions, such as city centers can be scheduled at night, when there is less traffic. The LiDAR recording vehicle will then be able to move faster through these areas with poor GPS reception, thus minimizing IMU drift. Still, the LiDAR trajectory will occasionally be worse, and if we are interested in absolute positioning rather than just registration of the data sets, then the independent GPS measurements from both trajectories should be taken into account with the appropriate weighting. To prevent unreasonable deformations of the LiDAR cloud, it would be necessary to constrain the LiDAR trajectory. Besides any GPS information, the INS data could be used to enforce smoothness. In fact, integrating raw GPS/IMU/DMI measurements into the adjustment procedure may be advantageous for the panorama trajectory as well.

As mobile LiDAR scanners capable of producing higher densities at reasonable driving speeds become available, direct interest point or line matching between images and LiDAR reflection data or geometry is likely to contribute more to the solution. Higher level features such as lampposts and windows can help as well. For the registration to work well in a wide range of environments, a combination of several different types of features should be integrated into the bundle adjustment.

Acknowledgements Part of this research has been supported by the GATE project, funded by the Netherlands Organization for Scientific Research (NWO) the Netherlands ICT Research and Innovation Authority (ICT Regie). The authors thank TopScan GmbH and Teccon bvba for the use of their data sets.

References

- Besl, P., McKay, N.: A Method for Registration of 3-D Shapes. *IEEE T. Pattern Anal.* 14(2), 239–256 (1992)
- Böhm, J., Becker, S.: Automatic Marker-Free Registration of Terrestrial Laser Scans using Reflectance Features. In: 8th Conference on Optical 3D Measurement Techniques. pp. 338–344 (2007)

- Carlberg, M., Andrews, J., Gao, P., Zakhor, A.: Fast Surface Reconstruction and Segmentation with Ground-Based and Airborne Lidar Range Data. In: 4th International Symposium on 3D Data Processing, Visualization and Transmission. pp. 97–104 (2008)
- Chen, Y., Medioni, G.: Object Modeling by Registration of Multiple Range Images. In: 1991 IEEE International Conference on Robotics and Automation. pp. 2724–2729 (1991)
- Clarkson, M., Rueckert, D., Hill, D., Hawkes, D.: Using Photo-Consistency to Register 2D Optical Images of the Human Face to a 3D Surface Model. *IEEE T. Pattern Anal.* 23(11), 1266–1280 (2001)
- van den Heuvel, F., Beers, B., Verwaal, R.: Method and system for producing a panoramic image from a vehicle. European Patent EP1903534. (2008)
- Huber, P.: Robust Regression: Asymptotics, Conjectures and Monte Carlo. *Ann. Stat.* 1(5), 799–821 (1973)
- Li, Y., Low, K.: Automatic Registration of Color Images to 3D Geometry. In: Proceedings of the 2009 Computer Graphics International Conference. pp. 21–28. ACM, New York (2009)
- Liu, L., Stamos, I.: Automatic 3D to 2D Registration for the Photorealistic Rendering of Urban Scenes. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. vol. 2, pp. 137–143 (2005)
- Lothe, P., Bourgeois, S., Dekeyser, F., Royer, E., Dhome, M.: Towards Geographical Referencing of Monocular SLAM Reconstruction using 3D City Models: Application to Real-Time Accurate Vision-Based Localization. In: 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 2882–2889 (2009)
- Lourakis, M., Argyros, A.: SBA: A Software Package for Generic Sparse Bundle Adjustment. *ACM. T. Math. Software* 36(1), 1–30 (2009)
- Lowe, D.: Distinctive Image Features from Scale-Invariant Keypoints. *Int. J. Comput. Vision* 60(2), 91–110 (2004)
- Mastin, A., Kepner, J., Fisher, J.: Automatic Registration of LIDAR and Optical Images of Urban Scenes. In: 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 2639–2646 (2009)
- Optech: Lynx mobile mapper data sheet. <http://www.optech.ca/pdf/LynxDataSheet.pdf> (2007)
- Pomerleau, F., Colas, F., Ferland, F., Michaud, F.: Relative Motion Threshold for Rejection in ICP Registration. In: Howard, A., Iagnemma, K., Kelly, A. (eds.) *Field and Service Robotics: Results of the 7th International Conference*. pp. 229–238. Springer, Heidelberg (2010)
- Zhao, W., Nister, D., Hsu, S.: Alignment of Continuous Video onto 3D Point Clouds. *IEEE T. Pattern Anal.* 27(8), 1305–1318 (2005)