A GUIDED REGISTRATION STRATEGY EMPLOYING VIRTUAL PLANES TO OVERCOME NON-STANDARD GEOMETRIES – USING THE EXAMPLE OF MOBILE MAPPING AND AERIAL OBLIQUE IMAGERY

P. Jende [1]*, F. Nex [1], M. Gerke [2], G. Vosselman [1]

[1] Dept. of Earth Observation Science, Faculty ITC, Universiteit Twente, The Netherlands –
(p.l.h.jende, f.nex, george.vosselman)@utwente.nl
[2] Institute of Geodesy and Photogrammetry, Technische Universität Braunschweig, Germany –
m.gerke@tu-bs.de

Commission II, WG II/1

KEY WORDS: Mobile Mapping, Feature Matching, Image Orientation, Oblique Imagery

ABSTRACT:

Mobile mapping (MM) is an intriguing as well as emerging platform and technology for geo-data acquisition. In typical areas of interest for MM campaigns, such as urban areas, unwanted GNSS multipath, non-line-of-sight effects, and IMU drifts may lead to deteriorated position fixes. In this work, we are proposing a novel technique to register MM and aerial oblique imagery. As aerial platforms are not affected by GNSS occlusions and are able to collect very-high resolution images, a co-registration of the data sets enables a) an independent verification of the platform's accuracy and b) an adjustment of the MM data's pose. Both data sets depict the scene from an entirely different perspective, which complicates the matching problem. Our approach is based on the assumption that common visible entities in both images are available, e.g. façade surfaces. By determining planes coinciding with these visible entities in object space, variances can be overcome. As the orientation of the data sets is known – MM data has an unknown accuracy – derived planes are employed to support a visibility hypothesis while storing image information for image registration in object space. This enables constraining search space to support the actual registration. Although the inhomogeneity of the data sets poses a challenge to a successful registration, we can show that our stepwise strategy of finding common visible entities, exploiting them to increase the resemblance of the data sets, and utilising accurate registration methods renders this matching scenario possible. In this paper, the algorithm is explained in detail, experimental results of significant steps will be shown, and possible extensions are discussed.

1. INTRODUCTION

Mobile Mapping has the unique capability to map large areas with an unprecedented degree of detail at relatively low cost. As a complementary technique to aerial photogrammetry, it enables an array of intriguing applications in computer vision, autonomous driving, and robotics. With high-precision instruments to locate the platform absolutely and relatively in space, its positioning accuracy is competing with other surveying-grade technologies.

These merits are opposing implications on the positioning capabilities of typical acquisition areas of MM, such as urban areas. Unlike aerial surveys, MM campaigns are directly affected by GNSS multipath and non-line-of-sight effects, and thus unable to correct for IMU drifts. Traditionally, ground control points (GCPs) are observed in acquired data products and introduced into adjustment solutions to fix the data products and/or the platform's trajectory (Cavegn et al., 2016). Alternative approaches may utilise digital maps (Gruyer et al. (2014), Roh et al. (2016), Schindler (2013)) or other tertiary data to increase the localisation accuracy (Gu et al. (2016), Groves (2011)).

A cost-efficient as well as accurate approach is enabling the adjustment of MM data by introducing observations by aerial images tying the data sets together. Particularly, MM imaging data and aerial imagery complement each other's coverage by a diverging perspective on the scene. This property may lead to quasi-complete representations of an area, yet it is complicating the registration of the data. By employing a multitude of strategies, this problem has been solved for aerial nadir imagery in conjunction with MM panoramic images (Jende et al., 2017). A remaining shortcoming

is, however, the establishing of correspondences in texture-, and feature-less areas. Relying on an aerial nadir view only, confines the potential overlapping area with MM data to roads, pavements, or other ground-related entities.

We therefore propose the extension of this approach to aerial oblique imagery. A successful integration will not just increase the number of observations but will also stabilise the geometry within a bundle adjustment considerably by enabling the integration of observations along the vertical direction, and by an inherently better intersection geometry.

2. RELATED WORK

Generally, this work intersects two major research problems, namely localisation and positioning as well as image registration. As mentioned in the previous paragraph, there is a number of alternative strategies to cope with the localisation problem. It is noteworthy, however, that approaches either aim for the correction of acquired data (as in our approach) or the platform's location or rather its trajectory. The latter is mainly required for real-time capable applications, such as Simultaneous Localisation and Mapping (SLAM) or Visual Odometry (VO) (Badino et al. (2013), Gupta et al. (2016), Kümmerle et al. (2011), Zhang and Singh (2015)). These approaches regard the absolute accuracy of collected data as a means to an end, thus it is secondary.

If the goal is to correct the data product (e.g. imagery), offline or post-processing becomes feasible. This allows for the design of algorithms aiming for high data accuracy through the integration of highly accurate tertiary data products, such as aerial images. For instance, Javanmardi et al. (2017) proposes an approach to

---

* Corresponding author

correct Mobile Laser Scanning (MLS) data with support from aerial nadir images by determining areas unlikely feasible for matching. The authors report an accuracy in the sub-decimetre range. Ji et al. (2015) correct MM images in urban areas by employing aerial nadir images in conjunction with a particle filter estimation. The results ascertain the feasibility of the approach, although the error modelling within the estimation pipeline has to act on many assumptions.

An extension towards a registration between MM and oblique images requires a more sophisticated approach. Morel and Yu (2009) proposed a registration pipeline based on synthesised views. The images are warped with an affine transformation to tackle perspective differences. A more flexible approach with respect to wide baseline matching problems has been proposed by Roth et al. (2017), which utilises synthesised views based on projective transformations. These registration techniques, however, are entirely unguided – in the case of Roth et al. (2017), only geometric verification is used. If the orientation of the image pair is coarsely known, the registration task can be facilitated considerably. Wu et al. (2018) suggest registering terrestrial and aerial oblique images with the support of plane priors extracted from 3D meshes. Their aim is to combine terrestrial and aerial point clouds to achieve a complementary coverage as well as to increase the resolution of the resulting 3D model.

3. METHODOLOGY

3.1 Overview

The aim of our procedure is to identify reliable correspondences between MM panoramic and aerial oblique images. The registration pipeline is designed to work in a fully automatic fashion (see Figure 1), as it will be integrated into a complete position estimation and verification workflow for MM imagery. A registration between two data sets, which do not share the same viewpoint, is not a trivial task. Hence, a registration cannot be conducted directly. Since MM data is predominantly acquired in urban areas, building façades are suitable objects visible from aerial oblique images as well. However, not every façade is per se adequate for registration purposes or visible from all data sets. For instance, repeated patterns, which are certainly common with respect to façades, may hinder a successful registration. Moreover, visibility hypotheses are needed to determine if a potential registration of a certain façade is even feasible. Thus, multiple strategies have been implemented to fulfil these prerequisites prior to the actual registration.

The first step in our approach is the registration of MM images along the trajectory. Since the MM images are adjusted relatively, reliable correspondences can be determined. The correspondences are used to triangulate image observations in 3D-space. Subsequently, planes are fit into the sparse point cloud. These planes are used to determine which points lie on a plane. Afterwards, virtual object points are created around every point on the plane. These points shape patches, and are used to extract image information for the registration in a later step. By employing the exterior orientation (EO) parameters of all aerial oblique images, individual visibility hypotheses are created for each patch. If the visibility hypothesis for a specific patch is successful, the patches are used to extract image information by back-projection from MM and aerial oblique images. This technique helps to overcome image differences, such as scale, rotation, and to some extent perspective. However, differences in illumination, contrast, or lighting have to be tackled within the actual registration step. Patches created from both image sources are co-registered with a template matching approach.
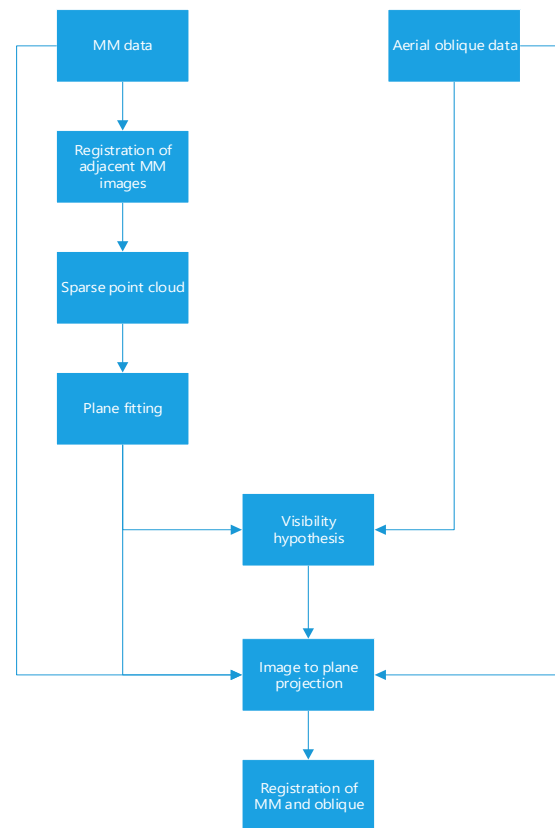

Figure 1. MM to oblique image registration pipeline

3.2 Registration of MM images along the trajectory

The MM system records a 360*180 degrees panoramic image every 5 metres along a predefined trajectory (see Figure 2). Since the data provider adjusts the MM images after acquisition, the data features a high relative accuracy. MM images are encoded in an equirectangular projection, where each pixel corresponds to an angular measurement. This entails a spherical projection leading to strong distortions if e.g. compared to a perspective image.
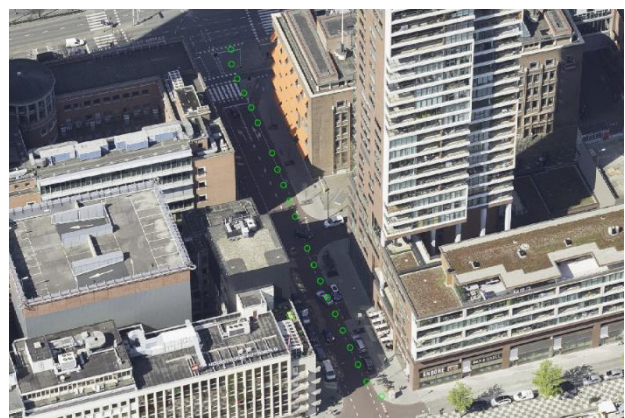

Figure 2. MM platform recording locations projected into aerial oblique image

In order to enable a reliable registration between adjacent MM images, each panoramic image is transformed to six perspective images (i.e. two image triplets) with yaw deviations from the driving direction as depicted in Figure 3 (left). The defined yaw angles enable enough coverage while still maintaining a manageable baseline for feature matching. Moreover, the pitch

angle for every perspective image is set to 30 degrees above the horizon with a vertical field of view of 75 degrees. These settings allow for an ideal capture of building façades if the assumption holds that buildings are parallel to the road and thus to the trajectory of the MM platform (see Figure 3 on the right).
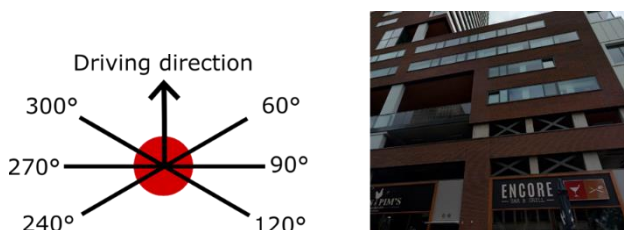

Figure 3. Left: yaw deviations from driving direction for perspective image creation. Right: Perspective image from panoramic image

Feature matching is realised using the AGAST-detector (Mair et al. (2010)) in combination with the DAISY-descriptor (Tola et al. (2010)). As identified correspondences are used for plane fitting as well as patch creation, AGAST as a corner detector proves useful for façade matching, since corners represent points of intersecting image gradients rather than distinctive areas (i.e. blobs). Consequently, image patches are e.g. created at and around window frames instead of arbitrary areas in the image, which may not be salient enough for a registration with the aerial oblique image. Additionally, the DAISY-descriptor works well for wide-baseline matching problems, and thus is able to register MM perspective images with a baseline of 5 metres.

The registration procedure is designed to work on image triplets for a more reliable outlier rejection. Since potential features on façades are parallel to the platform's trajectory, enforcing the epipolar constraint with a fundamental matrix will lead to geometric ambiguities. Surely, the same applies to the trifocal tensor setup, but the number of outliers can be strongly reduced. In practice, the trifocal tensor is constructed from three projection matrices composed from the intrinsic and extrinsic parameters of the perspective images (Hartley and Zisserman, 2004). The focal length can be derived for perspective images by:

$$f = \tan\left(90 - \frac{vfov}{2}\right) * w/2$$

where   vfov = vertical field of view
               w = image width

In order to verify if a putative correspondence in three images is valid, the trifocal point transfer is employed (see eq. 8 in Hartley (1997)). Hence, putative matches between two images and the trifocal tensor are used to compute the location of the keypoint in the third image. By comparing the distance between the computed and putative location of the third keypoint, outliers can be rejected. Subsequent to the registration of an image triplet (see

exemplary result in Figure 4), image observations are transformed back into the equirectangular projection for triangulation.

3.3 Plane fitting

In this step, planes are fit into a set of object points. To this end, a virtual recording location above the centre recording location is created. This allows for defining a normal vector, which is always perpendicular to the platform's trajectory. This vector serves as the reference vector for plane fitting. In particular, this reference vector prevents that planes are found, which are not parallel to the platform's trajectory and perpendicular to the ground (see Figure 5). In our case, the threshold has been set to 30 degrees.
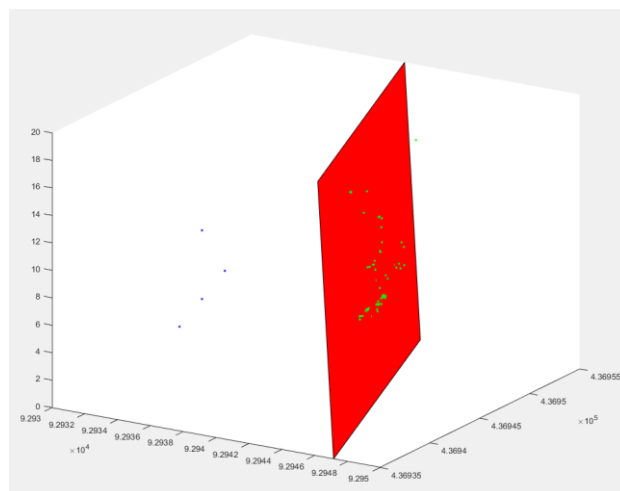

Figure 5. Plane fitting example. Green: object points. Red: fit plane. Blue: Three recording locations (image triplet) plus additional fourth virtual recording location

The plane fitting technique is based on MLESAC (Torr and Zisserman, 2000). After the plane has been found, it is used as a mask to determine the object points that contributed to the plane. Based on the normal vector of the inlying points, an orthogonal basis is computed. The basis is used to discretise a patch around every inlier object point in world coordinates. The pixel spacing and size of the patch can be arbitrarily defined. Since these patches are used for extracting image information from both, the MM and aerial oblique image, it is useful to set the pixel spacing to conform to the lowest resolution involved in process. In our case, this is 10 cm, as this is the average resolution of the aerial oblique data set.

3.4 Visibility hypothesis

In order to register MM and aerial oblique images, it is required that the patches created earlier depict the same object in all
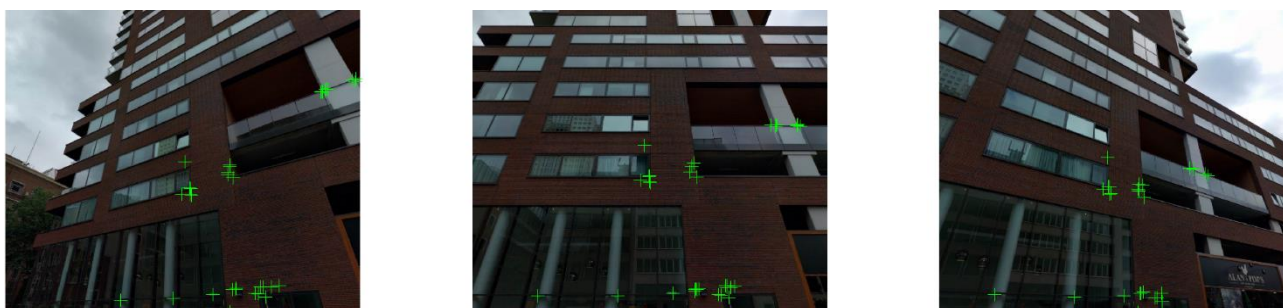

Figure 4. Image matching result of image triplet. From left to right: 60, 90, and 120 degrees perspective images

images. Especially in the case of aerial oblique images and an urban scenario, occlusions may occur and no direct line-of-sight between a certain façade and the aerial image's projection centre exists. To this end, three methods are used to create a visibility hypothesis. First, two angles φ and θ are computed, which constrain the angular offset of the oblique camera to the plane's normal vector in the vertical and horizontal dimension. In our scenario, both angles are defined not to exceed a maximum of 70 degrees. For instance, if φ, the azimuthal angle were 90 degrees, a façade would be in the nadir of the aerial oblique image, and thus not visible. Second, the scale of the plane in the aerial oblique image is obtained by the distance from the plane to the projection centre, the oblique camera's focal length, and the corresponding pixel size. This ensures that the required minimum spacing of pixels of the aerial oblique image patch can be maintained. At last, the centre point of a patch is projected into the camera to verify if the resulting image coordinates are inside the image plane. These three techniques can significantly reduce the number of aerial images considered for registration; nevertheless, there may be false positives due to occlusions by other objects (see Figure 6).

Figure 6. Back-projected patch centre coordinate. Left: False positive of visibility hypothesis. Right: True positive of visibility hypothesis of the same point

3.5 Image to plane projection and registration

Each patch is a discretised grid containing world coordinates around an object point labelled an inlier after plane fitting. Thus, each cell of the grid can be projected into the MM and aerial oblique image, as their orientation elements are known. The respective patch cell is then assigned the RGB value of the back-projected coordinate in the image. Certainly, depending on a potential position offset of the MM platform, the assigned image information may differ. However, as the plane is rigidly set in space, the image information – be it MM or aerial oblique image – is rectified. Moreover, both image projections have the same pixel spacing according to defined parameters earlier (please see Figure 7 and Figure 8).

Figure 7. Original image projections of MM [left] and aerial oblique image [right]. Please note the grey dots in the left hand image - these are projected grid cells used for the extraction of image information

Figure 8. Top: Image projections of MM and aerial oblique image onto grid in 3D-space. Both images have the same pixel spacing and share a similar perspective. Bottom: Same images after Wallis filtering

It becomes apparent that the projected images differ in contrast, illumination, and to some extent in content and perspective. To compensate for these differences, Wallis filtering (Wallis, 1976) is used to enhance the contrast of both images. Afterwards, a mesh grid around the centre pixel in the projected MM image is created. For every cell in the mesh grid including the centre pixel, a template is created, which is registered with the projected aerial oblique image. The corresponding peak in the correlation matrix is fed into a subpixel estimation process based on polynomial fitting. Subsequently, correspondences are transformed into their original image geometries, where further outlier removal techniques can be applied.

4. EXPERIMENTAL RESULTS AND DISCUSSION

This section shows and discusses first experimental results of the MM to aerial oblique matching pipeline. The data, which is used in this case, has been acquired in the city of Rotterdam, the Netherlands. The MM data has been artificially deteriorated in the horizontal dimension, as the original accuracy of the data cannot be disclosed due to proprietary knowledge of the data provider.

4.1 Registration of MM and aerial oblique images

As mentioned in the paragraphs before, a number of prerequisites have to be fulfilled in order to register aerial oblique and MM imagery successfully. Our registration pipeline is at an early stage, thus quantitative results are labelled future work. For demonstration purposes, a couple of examples are shown and discussed in this section. Figure 9 shows successful registration results between MM and aerial images using a template matching approach based on normalised cross correlation. Please note that no outlier removal step has been applied. The inter-image differences are rather strong, especially regarding perspective (top row) and content (bottom row). Although both images have been projected onto the same discretised grid in 3D-space, the original perspective affects the outcome – in particular for oblique images. Furthermore, the images differ regarding their

Figure 9. Four exemplary registration results between MM and aerial oblique images. The MM images are on the left hand side, the aerial oblique images on the right hand side

radiometry, and the matching task has to overcome sensor differences as well. As mentioned earlier, Wallis filtering has been used to counter differences in contrast, as normalised cross correlation is very sensitive towards illumination differences. Radiometric differences appear to be the most challenging part of this registration scenario alongside with diverging perspectives. Future endeavours will also focus on the utilisation of phase correlation, as it proved to be the most successful registration method for our aerial nadir to MM image matching pipeline.

4.2 Critical cases

4.2.1 Registration of MM images along the trajectory and its effects on plane fitting: Although the MM data has a high relative accuracy, and thus geometric relations between adjacent images can be exploited, the approach is not immune to outliers. There are multiple issues related to this part of the procedure, which may propagate to subsequent steps in the algorithm. For instance, if there is vegetation or people in the foreground (see Figure 10), feature matching may involve them into the registration result. If not classified as outliers, these parts of the image will be part of the resulting sparse point cloud and may affect the plane fitting process (see Figure 11).
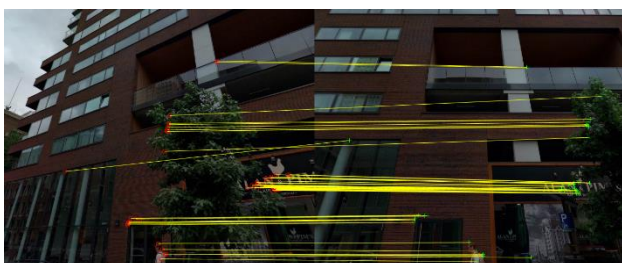

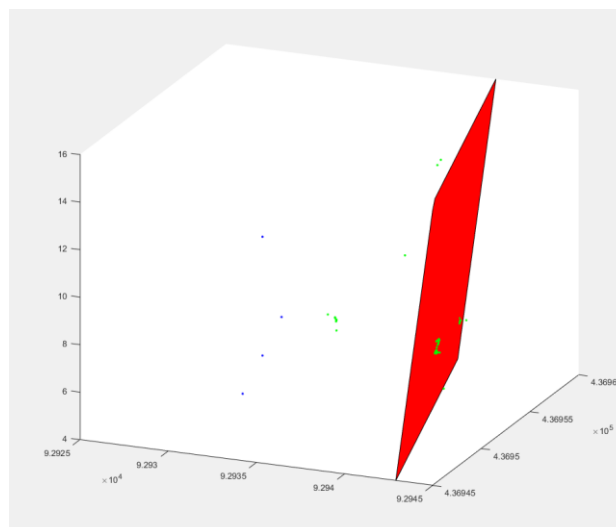Figure 10. Vegetation and pedestrian are part of the registration result

Figure 11. Sparse point cloud and fit plane of registration result in Figure 10; the pedestrian walked close to the wall and became part of the object points contributing to the plane. The tree in the foreground can be clearly distinguished from the façade points

In this case, the plane has been correctly fit into the point cloud. However, the fitting process led to a slight angle, which may affect subsequent image projections, and as a consequence the registration of MM and aerial oblique image. To constrain the fitting process to rigidly vertical planes, however, will not account for façades that are truly slanted. Another issue with the registration occurs if no suitable planar objects in the scene are present (see Figure 12), or the MM and aerial data differ too much due to e.g. a time difference between the data acquisitions. Although the registration result in Figure 12 is correct, all further steps render useless in this scenario.
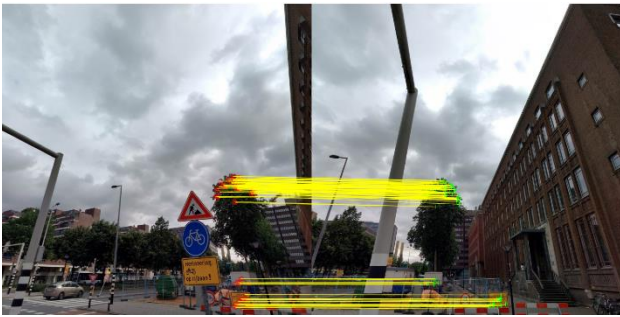
Figure 12. No planar objects in the scene prevent a continuation of the workflow

Certainly, mismatches may occur during this process. However, wrong correspondences lead to wrong triangulation results, and thus it depends on individual correct correspondences that allow for correct plane fitting (see Figure 13). Repeated patterns remain an issue for the MM registration task. In Figure 14, however, wrong correspondences did not affect the plane fitting process, as the triangulated object points are too scattered, and thus not considered by the MLESAC plane fitting process.
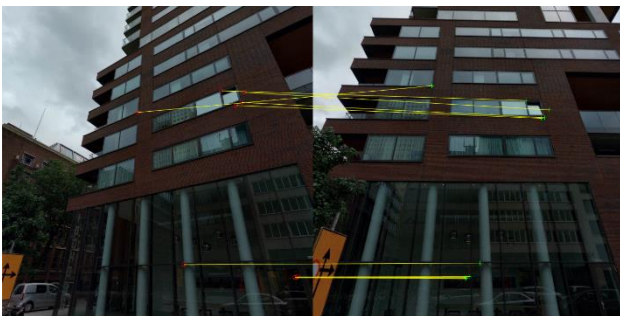

Figure 13. Mismatches due to repeated patterns on the same epipolar line
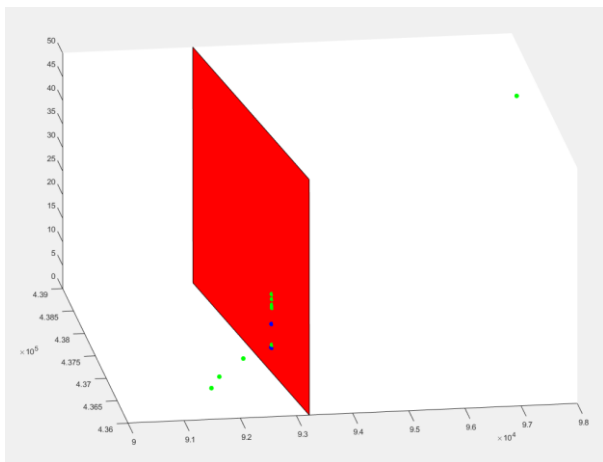

Figure 14. Plane fitting in a sparse point cloud with outliers. Please note: the triangulation of mismatches scaled the scene accordingly. The blue points depicting the MM recording locations are not on the plane.

To summarise, the MM registration process could be further advanced by e.g. integrating feature tracking. Currently, features are identified in image triplets but are not propagated. The plane fitting process is based on MLESAC, which finds a solution on maximum likelihood rather than the number of inliers (compared to RANSAC). However, the spatial distribution of correspondences and thus object points could be potentially used

to stabilise the plane estimation. To this end, more weight could be assigned to points in the same vertical dimension with greater spatial distribution rather than cluttered points. Alternatively, least square fitting could be used for the plane estimation, thus no weighting is needed.

4.2.2   Visibility hypothesis and its limitations: As discussed in the previous paragraphs, a visibility hypothesis is based on three independent methods: two angular constraints, scale, and back-projection into the image plane. This approach proved to be quite robust but cannot obviate situations of occlusions or severe image differences (see Figure 6). For example, Figure 15 shows a building whose façade has been properly detected in the earlier steps. Due to the vegetation along the avenue, a successful registration is not feasible. This shortcoming has an unfortunate effect on the processing speed, as this problem becomes apparent not earlier than in the registration phase itself.
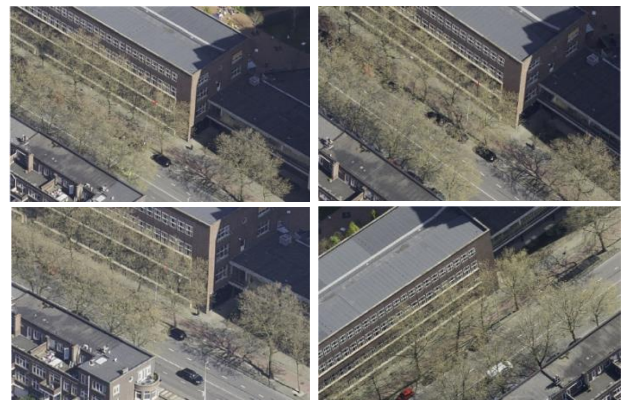

Figure 15. Façade is occluded due to vegetation

An approach to overcome that is to select only façades for further processing if the sparse point cloud is not scattered or no object points are between the plane and the recording locations, similar to the approach of Nyaruhuma et al. (2012). Moreover, the scale information obtained during this process could be used to have a resolution-adaptive image projection approach. This could potentially increase the accuracy of the registration and avoid aliasing effects (visible in Figure 8), as the scale in oblique images varies strongly.

5.   CONCLUSION

This paper presented an approach to register images with an entirely different geometry, radiometry, and content by exploiting common entities visible in both images. It has been designed for the registration problem between mobile mapping panoramic images and aerial oblique images. The approach is based on finding façades in sparse point clouds created from terrestrial image correspondences. These façades are employed as projection surfaces for both image data sets. The actual registration is conducted by a template matching approach. It could be shown that a registration of MM and aerial oblique images is feasible if the prerequisites are met. Even though a case study with quantitative results is labelled future work, the success rate – without outlier removal and with the data used in this paper – is about 80%. Many limitations can be counteracted by an array of techniques involving an understanding of the scene. In particular, a visibility hypothesis, plane estimation, and image projection are valuable tools to approach this non-standard registration problem. In our future work, we will focus on the robustness of the registration procedure with respect to overall image differences, and the generalisation of the method towards other data sets.

ACKNOWLEDGEMENTS

REFERENCES

Badino, H., A. Yamamoto and T. Kanade, 2013. Visual odometry by multi-frame feature integration. *Proceedings of the IEEE International Conference on Computer Vision Workshops*

Cavegn, S., S. Nebiker and N. Haala, 2016. A systematic comparison of direct and image-based georeferencing in challenging urban areas. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.,* XLI-B1, pp. 529-536.

Groves, P.D., 2011. Shadow matching: A new GNSS positioning technique for urban canyons. *Journal of Navigation* 64, pp. 417-430.

Gruyer, D., R. Belaroussi and M. Revilloud, 2014. Map-aided localization with lateral perception. *Intelligent Vehicles Symposium Proceedings, 2014 IEEE*

Gu, Y. L., L. T. Hsu and S. Kamijo, 2016. GNSS/Onboard Inertial Sensor Integration With the Aid of 3-D Building Map for Lane-Level Vehicle Self-Localization in Urban Canyon. *IEEE Transactions on Vehicular Technology* 65(6), pp. 4274-4287.

Gupta, A., Chang, H., Yilmaz, A., 2016. Gps-Denied Geo-Localisation Using Visual Odometry. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information*, Vol. III-3, pp. 263-270.

Hartley, R., Zisserman, A., 2004. Multiple View Geometry in Computer Vision, Cambridge University Press

Hartley, R.I., 1997. Lines and points in three views and the trifocal tensor. *International Journal of Computer Vision* 22, pp. 125-140.

Javanmardi, M., Javanmardi, E., Gu, Y., Kamijo, S., 2017. Towards High-Definition 3D Urban Mapping: Road Feature-Based Registration of Mobile Mapping Systems and Aerial Imagery. *Remote Sensing* 9, p. 975.

Jende, P., Nex, F., Gerke, M., Vosselman, G., 2017. Fully Automatic Feature-Based Registration Of Mobile Mapping And Aerial Nadir Images For Enabling The Adjustment Of Mobile Platform Locations In GNSS-Denied Urban Environments. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci*. XLII-1/W1, pp. 317-323

Ji, S., Shi, Y., Shan, J., Shao, X., Shi, Z., Yuan, X., Yang, P., Wu, W., Tang, H., Shibasaki, R., 2015. Particle filtering methods for georeferencing panoramic image sequence in complex urban scenes. *ISPRS Journal of Photogrammetry and Remote Sensing* Vol. 105, pp. 1-12.

Kümmerle, R., Steder, B., Dornhege, C., Kleiner, A., Grisetti, G., Burgard, W., 2011. Large scale graph-based SLAM using aerial images as prior information. *Autonomous Robots* 30, pp. 25-39.

Mair, E., Hager, G.D., Burschka, D., Suppa, M., Hirzinger, G., 2010. Adaptive and generic corner detection based on the accelerated segment test. *European conference on Computer vision* 2010, pp. 183-196.

Morel, J.-M., Yu, G., 2009. ASIFT: A New Framework for Fully Affine Invariant Image Comparison. *SIAM Journal on Imaging Sciences* 2.2, pp. 438-469.

Nyaruhuma, A.P., Gerke, M., Vosselman, G., Mtalo, E.G., 2012. Verification of 2D building outlines using oblique airborne images. *ISPRS Journal of Photogrammetry and Remote Sensing* Vol. 71, pp. 62-75.

Roh, H., J. Jeong, Y. Cho and A. Kim, 2016. Accurate Mobile Urban Mapping via Digital Map-Based SLAM. *Sensors* 16(8): p. 1315.

Roth, L., Kuhn, A., Mayer, H., 2017. Wide-Baseline Image Matching with Projective View Synthesis and Calibrated Geometric Verification. *PFG – Journal of Photogrammetry, Remote Sensing and Geoinformation Science* Vol. 85, pp. 85-95.

Schindler, A., 2013. Vehicle self-localization with high-precision digital maps. *2013 IEEE Intelligent Vehicles Symposium (IV)*, pp. 141-146.

Tola, E., Lepetit, V., Fua, P., 2010. Daisy: An efficient dense descriptor applied to wide-baseline stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, pp. 815-830.

Torr, P.H., Zisserman, A., 2000. MLESAC: A new robust estimator with application to estimating image geometry. *Computer vision and image understanding* 78, pp. 138-156.

Wallis, R.H., 1976. An Approach to the Space Variant Restoration and Enhancement of Images, *Proceedings of the Symposium on Current Mathematical Problems in Image Science*, Montery.

Wu, B., Xie, L., Hu, H., Zhu, Q., Yau, E., 2018. Integration of aerial oblique imagery and terrestrial imagery for optimized 3D modeling in urban areas. *ISPRS Journal of Photogrammetry and Remote Sensing* Vol. 139, pp. 119-132.

Zhang, J., Singh, S., 2015. Visual-lidar odometry and mapping: low-drift, robust, and fast. *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2174-2181.