DENSE MATCHING IN HIGH RESOLUTION OBLIQUE AIRBORNE IMAGES

M. Gerke

International Institute for Geo-Information Science and Earth Observation – ITC, Department of Earth Observation Science, Hengelosestraat 99, P.O. Box 6, 7500AA Enschede, The Netherlands, gerke@itc.nl

KEY WORDS: Adjustment, Bundle, Calibration, Matching, Point Cloud, Rectification

ABSTRACT:

An increasing number of airborne image acquisition systems being equipped with multiple small- or medium size frame cameras are operational. The cameras normally cover different viewing directions. In contrast to vertical images, those oblique images have some specific properties, like a significantly varying image scale, and more occlusion through high raising objects, like buildings. However, the faces of buildings and other vertically extended objects are well visible and this is why oblique images are used for instance for visualization purposes.
This paper shows results from applying the sophisticated Semi-Global-Matching technique to a set of oblique airborne images. The images were acquired by two systems, namely FLI-MAP 400 (Fugro Aerial Mapping B.V.) and Pictometry (BLOM Aerofilms) over the same area. After the joint adjustment of the images, dense matching and forward ray intersection was performed in several image combinations. The disparity maps were evaluated through the comparison with a reference map derived from LIDAR which was acquired in parallel with the FLI-MAP system. Moreover, the 3D point clouds were analyzed visually and also compared to the reference point cloud. Around 60 to 70 percent of all matches were within a range of $\pm$ 3pix to the reference. Since the images were acquired in different flight configurations, the impact of different intersection angles and baselines to the triangulation is quite obvious. In general, the overall structures on the building faces are well represented, but the noise reduction needs further attention.

1 INTRODUCTION

An increasing number of airborne image acquisition systems are operational (Petrie and Walker, 2007). Because of the availability of low-cost digital cameras with small or medium sized sensors, some of those systems carry multiple cameras covering different viewing directions. For instance from Pictometry[1] image are available already for a number of cities and they are accessible in the category "birds eye view" in Microsoft Bing Maps[2] (formerly known as Virtual Earth).

The use of oblique images for topographic mapping purposes was shown in quite some papers. In (Höhle, 2008) height determination from single oblique images is demonstrated. The verification of vector data using oblique imagery is shown in (Mishra et al., 2008). Due to the fact that building façades are well visible in oblique images, some researchers concentrate on how to automatically extract façade textures (Früh et al., 2004, Wang et al., 2008). Besides, the oblique images are interesting for cadastre applications, because the building outline as defined at the vertical wall is directly visible (Lemmen et al., 2007). Compared to vertical airborne images, oblique images have some specific properties. Depending on the tilt angle, the scale within the imaged scene varies considerably. Moreover, vertical structures of raised objects like buildings or trees are imaged, but the (self)occlusion by those objects is much more significant compared to the vertical image case.

Another interesting application and research domain concerns the derivation of high dense point information through image matching techniques. The benchmark results from the Middlebury[3] testsets show that high quality state-of-the-art techniques to dense matching are available. If it is possible to apply those techniques to oblique airborne images, interesting new applications would arise, or support existing ones, like the ones listed above. In general, point clouds as derived from dense matching in oblique images can be a complementary data source to airborne laser scanning, as those devices normally do not capture dense points on vertical structures. Of course, the traditional use of this kind of data to produce digital surface or terrain models is another possible application.

In (Besnerais et al., 2008) an approach to dense matching in oblique airborne images is presented. The authors develop a pixel wise similarity criterion which accounts for the special viewing geometry of oblique images. A dense depth map is obtained through global regularization. The approach was tested on a number of test images and showed good results. However, the ground sampling distance of the used images was not smaller than 1.4m, mostly it was even larger, up to 20m.

This paper evaluates the application of the Semi-Global-Matching technique (SGM, see (Hirschmüller, 2008)) to a set of high resolution FLI-MAP[4] and Pictometry images. One particular façade of a building is normally only visible in images taken from one viewing direction, resulting in a relatively bad intersection angle in object space. Thus, the main objective of this paper is to evaluate the overall accuracy of the derived 3D point cloud as derived from a forward intersection of matched points. Although other – may be better performing– algorithms for dense matching exist (Seitz et al., 2006) we chose SGM, because it demonstrated already its fitness for the photogrammetric production process, c.f. (Hirschmüller et al., 2005).

As no sufficient calibration and orientation information was available, the whole block first needed to be adjusted. The method for bundle block adjustment, including self-calibration of multiple devices and employing scene constraints to enhance the scene geometry was introduced and tested in (Gerke and Nyaruhuma, 2009). The dense matching algorithm then was applied to several combinations of stereo images, and results were evaluated through LIDAR data which was acquired from the FLI-MAP system.

[1] http://www.pictometry.com
[2] http://www.bing.com/maps
[3] http://vision.middlebury.edu/stereo/ (accessed 15 March 2009)
[4] http://www.flimap.nl

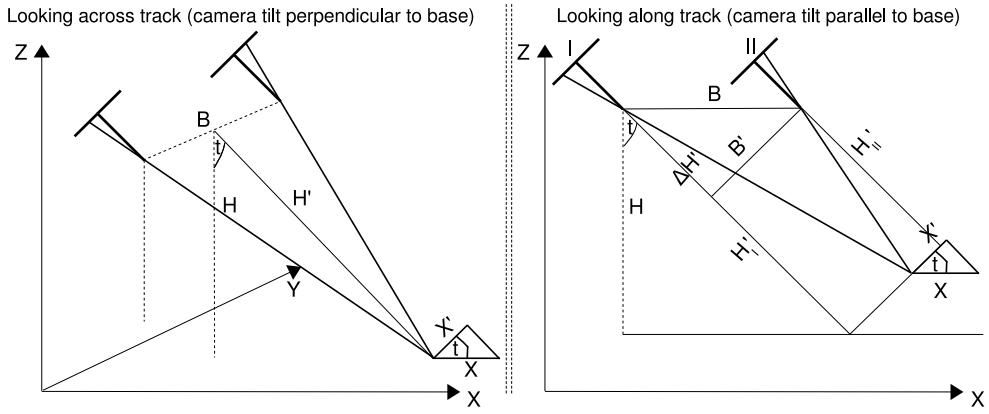

Figure 1: Geometry of across track (left) and along track (right) baseline in oblique stereo images

2 STEREO GEOMETRY AND ACCURACY IN OBLIQUE IMAGES

The dense matching results in point correspondences in a stereo pair[5]. The image rays as resulting from those matching image pairs are forward intersected in order to compute 3D points in object space. In this section a brief theoretic approximatoin of the expected accuracy from the given intersection geometry in oblique airborne images is derived. Two cases can be distinguished: A) the two cameras used for the stereo intersection are oriented across track, i.e. they inclose a right angle with the base, and B) the cameras are looking in flight direction. The intersection geometry of the first case can be derived from the standard normal case, but the varying scale and the tilt of the camera coordinate system wrt the actual coordinate system need consideration. The second case can be compared with the longitudinal tilt-setup (Albertz and Kreiling, 1980), i.e. the normal case with different camera heights. Only here the whole base is tilted. In Fig. 1 both camera geometries are sketched.

The scale within an oblique image depends on the flying height H, the focal length c the tilt angle t, and the angle between the viewing ray to a target and the vertical β (symbols according to (Höhle, 2008)):

$$m = \frac{H \cdot \cos(\beta - t)}{c \cdot \cos\beta}, \tag{1}$$

where m: scale at target point. At the principal point β equals t, whereas at the fore- and the background β is determined from t and the half field of view α: $\beta_{fore} = t - \alpha$ and $\beta_{back} = t + \alpha$.

A: tilt across track (side-looking) In the vertical image case, the accuracy for triangulated points in height (s'_H), and in X-Y plane ($s'_{X,Y}$) can be estimated by:

$$s'_H \approx \frac{H'}{B} \cdot m \cdot s_{px}, \tag{2}$$

$$s'_{X,Y} \approx s_x \cdot m \approx s_y \cdot m, \tag{3}$$

where $s_x \approx s_y \approx 0.5 \cdot s_{px}$ are the standard deviations for image coordinate and parallax measurements; the errors in the orientation components are neglected. In the case of the tilted camera system, these formulas are applicable to the tilted system, so the varying scale needs to be considered, according to equation 1, also H' needs to be adopted accordingly:

$$H' = m \cdot c. \tag{4}$$

[5]For the combination of multiple views see the experiment section

Finally, the respective error components need to be projected from the tilted system to the actual coordinate system:

$$s_H \approx \sqrt{(s'_H \cdot \cos t)^2 + (s'_{X,Y} \cdot \sin t)^2}, \tag{5}$$

$$s_{X,Y} \approx \sqrt{(s'_H \cdot \sin t)^2 + (s'_{X,Y} \cdot \cos t)^2}, \tag{6}$$

thus for a tilt angle of 45° both components will be identical.

B: tilt along track (forward-looking) To derive the accuracy in the tilted system H, X', Y', first the necessary parameters for the case of longitudinal tilt need to be computed: Base B' in the tilted system and the heights of the cameras I and II:

$$B' = B \cdot \cos t, \tag{7}$$

$$\Delta H' = B \cdot \sin t, \tag{8}$$

$$H'_I = m \cdot c, \quad \text{and} \quad H'_{II} = H'_I - \Delta H'. \tag{9}$$

Applying partial derivation wrt the image and parallax measurements to the formulas given in (Albertz and Kreiling, 1980), the accuracies for the coordinate components in the tilted system can be derived:

$$s'_{H_I} \approx s'_{H_{II}} \approx \sqrt{\left(\tfrac{H_{II}'}{p_x}\right)^2 \cdot s_{px}^2 + \left(\tfrac{B' \cdot \sin t}{p_x}\right)^2 \cdot s_x^2}, \tag{10}$$

$$s'_{X,Y} \approx \tfrac{H_I'}{c} \cdot s_x. \tag{11}$$

Note that the actual parallax needs to be computed for the estimation. In the approximations for the given data, see below, a mean parallax according to a mean depth in fore- and background was assumed. For the planar accuracy the more pessimistic estimation, assuming the smaller image scale, is given here. Finally, the planar and height components in the actual coordinate system are computed according to equations 5 and 6.

3 METHODS ADOPTED

3.1 Block adjustment for multiple platforms

In (Gerke and Nyaruhuma, 2009) a method to incorporate scene constraints into the bundle block adjustment is described and tested. The bundle block adjustment algorithm uses horizontal and vertical line features, as well as right angles to support the stability of block geometry. Those features can be identified at building façades, as visible in oblique images. In addition, the approach is able to perform self-calibration on all devices which are incorporated in the block. This is an important issue in the case of oblique images, as those are often acquired by non-metric cameras. The extension to the setup used for this paper where images from different platforms are involved is done without any change to the core approach.

Parameter	FLI-MAP vertical	FLI-MAP oblique	Pictometry oblique
flying height [m]	275	275	920
baseline [m]	50	50	400
tilt angle [°]	0	45	50
number of images,viewing direction	7	8xSW	2xW, 2xS, 2xE, 1xN
focal length [mm]	35	105	85
pixel size [μm]	9	9	9
sensor size [mm x mm]	36x24	36x24	36x24
GSD and theoretic accuracies (for oblique: from fore- to background)			
ground sampling distance [cm]	7	2.8 – 4	10 – 16
$s_{X,Y}$, vertical [cm]	4	NA	NA
s_Z, vertical [cm]	40	NA	NA
$s_{X,Y}$, across track base [cm]	NA	NA	22 – 44
s_Z, across track base [cm]	NA	NA	18 – 37
$s_{X,Y}$, along track base [cm]	NA	60 – 92	22 – 42(*)
s_Z, along track base [cm]	NA	60 – 92	19 – 35(*)

(*): along track base images from Pictometry were not used.

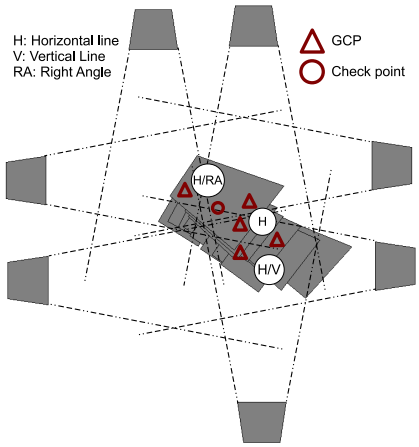

Figure 2: Image parameters and layout of sample block

3.2 Rectification and dense matching

The approach to dense stereo matching as applied in the current implementation is the Semi-Global-Matching algorithm (Hirschmüller, 2008). The basic idea behind this technique is to aggregate local matching costs by a global energy function, which is approximated by an efficient pathwise 1-dimensional optimization.

To simplify the matching, the images are rectified beforehand. For this purpose the approach proposed in (Oram, 2001) is applied. A homography is estimated which is compatible to the fundamental matrix (Hartley and Zisserman, 2004, chap. 13). The aim is to minimize distortions due to perspective effects, and thus also to reduce the disparity search space. A consequence from the particular algorithm is that the epipolar lines are coincident, but not necessarily parallel. Hence, in the subsequent rectification the images need resampling to obtain parallel epipolar lines. One disadvantage of this procedure is that straight lines are not preserved, however, this does not influence the matching. To compute the optimal homography point correspondences are required, like for instance in the case at hand the adjusted tie points. If images are taken from approximately the same viewing direction, it is also possible to extract further matches through scale invariant point descriptors like SIFT (Lowe, 2004). Outliers in the correspondences are identified through RANSAC within the estimation of the compatible homography. The inliers are here also used to estimate the disparity search range for the dense matching.

4 EXPERIMENTS

4.1 Description of used data

Part of the data used for these experiments was acquired by the Fugro Inpark FLI-MAP 400 system in March 2007 over Enschede, The Netherlands. Besides two LIDAR devices and two video cameras, the system carries two small-frame cameras, one pointing vertical, and one oblique camera, looking in flight direction, tilted by approx. 45°. Additional Pictometry images were made available through BLOM Aerofilms. Those images were acquired only one month before the FLI-MAP data. A small block of 7 vertical and 8 oblique images from FLI-MAP as well as 7 images from Pictometry was chosen for the experiments. In Fig. 2, upper part some parameters of the images are given, the GSD and accuray estimation was done according to equations 1 to 11, while a standard deviation for image measurements of a half pixel was assumed. In the bottom of that figure the layout of the block is shown, including GCP, check points and the approximate position of defined scene constraints. The highly overlapping images in the center are from the FLI-MAP acquisition, while the 7 regularly aligned outer images are from the Pictometry-flight. Note that no along track images are chosen from Pictometry. The airplane acquired the images in N-S-direction, so the East- and West-looking images belong to one flight line (baseline approx. 400m) and the two South-looking images are from two adjacent strips, baseline approx. 350m. For the accuracy estimation the two South-looking images can be treated like across-track images.

4.2 Block adjustment results

Four full and one height GCP were used for the adjustment. Additionally, one right angle, 3 horizontal and 4 vertical line constraints were defined. It was assured that in every image at least one of the features used for the scene constraints was visible. In Table 1 the adjustment results in terms of RMSE at the control and check points, or features respectively are listed. One observation from the residuals is that the Z-component is smaller than the X/Y values for all control and check features. Also the residuals at vertical constraints are larger than the residuals at horizontal constraints, and those are also influenced by the Z-component only. One reason for this can be that the tilt of the Pictometry images is larger than 45° and thus the X,Y-component is less accurate than the Z-component, refer also the the listed theoretic accuracies in Fig. 2. One general drawback of this block-setup is that outside the overlapping areas no GCPs or scene constraints are available and applicable, respectively, so the overall block geometry at the borders is not optimal. However, since the residuals at the façades are at least for the Pictometry images less than one pixel this result can be considered satisfactory.

Assessment	RMSE value[cm]
X-Res. at GCP	2.1
Y-Res. at GCP	4.8
Z-Res. at GCP	1.3
X-Res. at Check	16.2
Y-Res. at Check	5.8
Z-Res. at Check	1.5
Res. at H-constraints	1.4
Res. at V-constraints	6.8
Res. at RA-constraints (°)	0.01

Table 1: Residuals from bundle block adjustment

4.3 Dense matching results

The dense matching was performed in several stereo image combinations. Besides the matching in images from one platform,

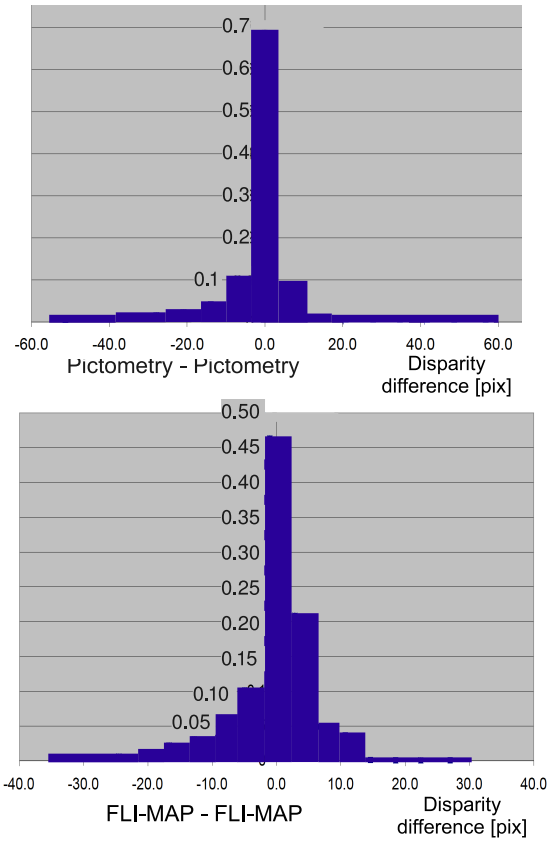

Figure 3: Two sample relative histograms of disparity differences wrt reference disparity map. Note the scale difference between both diagrams.

matching between FLI-MAP and Pictometry was tested. This is interesting, because by this means the scene can be observed from approximately the same direction through multiple views. The overlap from consecutive Pictometry images is not large enough to create 3-ray points, however, incorporating also FLI-MAP images makes this possible. Besides, this setup gives an interesting geometry for forward intersection.

Two methods were used to assess the results: one quantitative and one qualitative. For the quantitative assessment a reference disparity map was computed from the FLI-MAP LIDAR data, then the differences to the disparities from image matching were analyzed using histograms. For a more qualitative assessment 3D point clouds were computed from the matching results and then assessed visually, also in comparison to the LIDAR point cloud.

Disparity map assessment For this assessment the reference LIDAR points (density: 20 points per m^2) were projected into the image plane as defined by the respective image orientation and calibration parameters and subsequently a reference disparity map was computed. Two issues are important here: first, only first pulse LIDAR points should be considered, as also in the image only the visible surface can be matched. Second, through the oblique viewing direction as realized with the cameras one has to take into account self-occlusion through buildings; the laser scanner scans vertical and thus scans other parts of the scene, especially on the backside of buildings visible in the images. To avoid errors from that circumstance, only areas which do not show these effects were used for the evaluation.

The disparity maps were assessed by calculating the difference disparity map and computing a histogram out of that one. Only pixels showing a disparity value in both maps were considered,

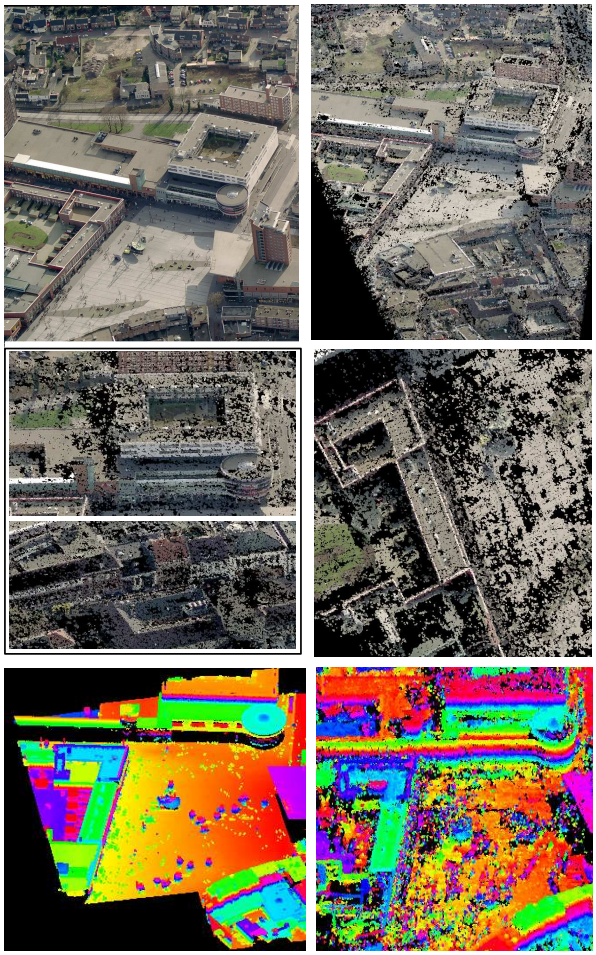

Figure 4: Results from dense matching in two overlapping South-looking Pictometry images. Top: left image and 3D cloud from matching, center row: zoom to point cloud from matching at façades (left) and top view (right), bottom row: point cloud color coded height: reference (left), from matching (right)

thus matched points at façades which were not acquired by the LIDAR device can not be assessed. Two of such relative histograms are shown in Fig. 3. The upper histogram shows the differences from the matching within two Pictometry images (see Fig. 4). For this histogram approx. $50 \cdot 10^3$ matches were considered (out of $2.2 \cdot 10^6$ in total), and around 70% of them show a difference of ± 3 pixels to the reference. The histogram at the bottom shows the analysis from the matches within two oblique images from FLI-MAP, refer to Fig. 5. For this histogram approx. $200 \cdot 10^3$ matches were considered (out of $6.4 \cdot 10^6$ in total). Because of the smaller baseline between consecutive FLI-MAP images, compared to Pictometry, the overlapping area is larger, and thus results in more matches. Approximately 60% are within the difference of ± 3 pixels. All matches outside this tolerance can be considered as blunder. A more in depth analysis revealed that most blunders were caused in shadow areas or other areas with poor texture. When assessing those histograms is should be considered that errors from the image calibration and post estimation also contribute to those residuals, thus a final conclusion on the absolute matching accuracy of the SGM implementation can not be made.

Point clouds: Pictometry to Pictometry For the following evaluations a forward intersection of the matched points was performed. A simple blunder detection was implemented by applying a threshold to the residual for image observations. For two-ray intersections this method can filter some blunders, but be-

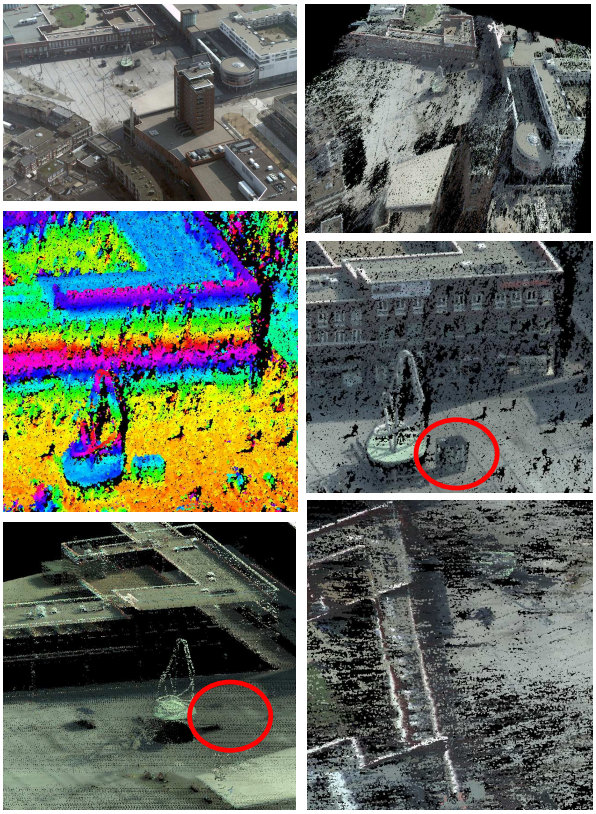

Figure 5: Results from dense matching in two overlapping FLI-MAP images. Top: part of left image and point cloud from matching, centre: 3D point cloud from matching, color coded height (left) and color picked from images, bottom row: reference point cloud and top view to matched point cloud . The circle indicates an elevator box which is visible in the point cloud from matching, but not in the laser scanning data.

cause only one redundant observation is available, quite a lot of blunders will not be detected.

The point cloud as resulted from the triangulation of matches in a Pictometry image pair are shown in Fig. 4. The top row in that figure shows part of the left image and an overview on the 3D scene defined by the matched point cloud. The center row shows a zoom in to the colored point cloud from matching, focusing on some façades and the vertical view to that scene. Finally, the bottom row shows the reference point cloud at the left hand side, where the color codes the height (one full color cycle equals 10m). The corresponding part of the matched point cloud is depicted on the right hand side. From that figure some interesting observations can be made. Although most of the flat roofs show poor texture, the corresponding part in the matched point cloud is quite dense and homogeneous. However, the height of the roof in the upper part is not correct, it is approx. 50cm lower than shown in the reference point cloud. In the vertical view on the point cloud from SGM the occluded areas are clearly visible, whereas the vertical façades are not visible in the reference point cloud. Overall, the structures are well represented, but the mismatched pixels impose a certain level of noise to the scene. Those mismatches can hardly be detected, as only stereo matches are available in this case. The detailed zoom on the vertical image shows that the accuracy in x-y direction is quite good, and apparently even better than the estimated one ($s_{X,Y}$: 22 – 44cm).

Point clouds: FLI-MAP to FLI-MAP The triangulated point cloud from the matching in two consecutive oblique images from the FLI-MAP data is shown in Fig. 5. The zoom in to the colored point cloud shows quite some details, for instance the elevator box close to the statue is clearly visible. The statue itself is well represented and even gaps in the matches where people were walking during image exposure are clearly identifiable. However, those views onto the point cloud were made from the viewing direction of the cameras, so the main error showing and effect in viewing direction is not visible. The estimated depth accuracy (in viewing direction, s'_H) of the along-track FLI-MAP data varies from 90 to 130cm, and the error in X', Y'-direction is only 2cm. To assess the overall quality, the vertical view needs to be considered: Here the uncertainty in viewing direction is quite obvious. If the result from the vertical view-zoom is compared to the one from the Pictometry data (Fig.4, center row), it can be observed that the result from the FLI-MAP data is more inaccurate. This visually achieved observation confirms the theoretically approximated accuracy, which is about four times worse.

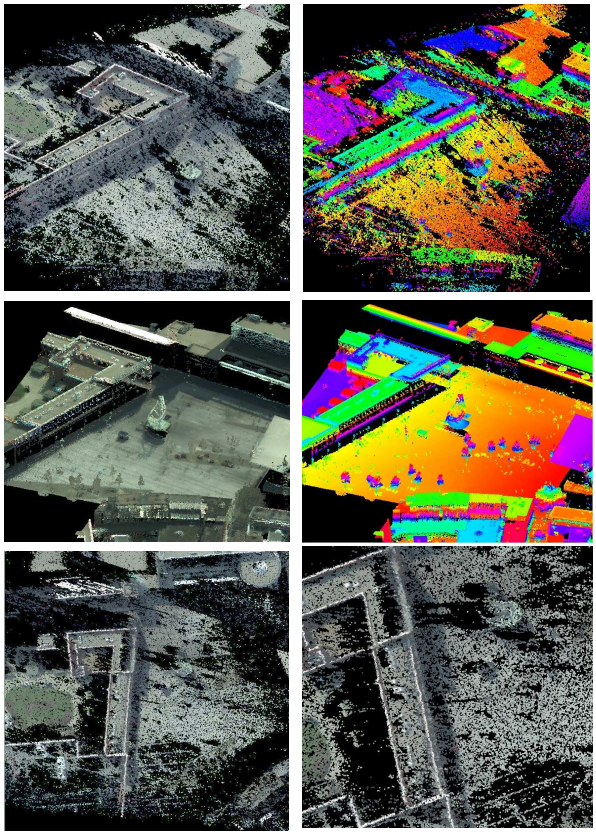

Figure 6: Results from multiple view triangulation. Top: matched point cloud, color and height, center: reference point cloud, bottom: top view

Point clouds: Multiple view For this experiment it was desired to exclude the wrong matches from the triangulation. To achieve this goal, the dense matching results from several matches were combined in the following manner: Dense matching was performed in the three combinations: ① FLI-MAP1 ↔ FLI-MAP2; ② FLI-MAP2 ↔ Pictometry1; ③ FLI-MAP1 ↔ Pictometry1. The matches from ① and ② are linked in a way that the matching points from the right image of ① are searched in the left image of ② and by this means corresponding matches ③' FLI-MAP1 ↔ Pictometry1 are created. In the subsequent triangulation only those matches were considered which coincide with the original matches from ③. Thus it can be expected that through this double check some blunders were removed. For more details on this method see (Gerke, 2008). In Fig. 6 some details on the results are shown. The point cloud contains now less points ($1.6 \cdot 10^6$ points from the FLIMAP-only point cloud vs. $190 \cdot 10^3$

matches here), but the overall accuracy seems to be better, see e.g. the height of the buildings. Also the detailed look from vertical shows less noise, compared to the two-fold matches before.

5 CONCLUSIONS AND OUTLOOK

This paper reports about the utilization of the dense image matching technique Semi-Global-Matching to a set of high resolution oblique airborne images. The images were acquired from different platforms and thus in different configurations. The comparison of the depth maps from matching with a reference computed from LIDAR data showed that roughly 70% of all matches are within an error range of $\pm$ 3pixel, however, also the residual errors from camera calibration and orientation have an impact on this evaluation. The remaining matches can be considered as blunders. How can those blunders be removed and the noise level be reduced? If multiple overlaps are available, sophisticated error analysis prior to triangulation is feasible (Hirschmüller, 2008). Also the method as applied here shows good results, namely to eliminate wrong matches through linking matches of adjacent images and applying a double check through a direct match. Other authors use the much stronger trinocular stereo geometry for matching (Heinrichs et al., 2007), or apply a similarity criterion for multiple views directly (Besnerais et al., 2008). If only two-fold overlap is available – as mostly in facade observations from oblique images – one method could be to incorporate reliable SIFT features within the SGM approach directly: The disparities as defined by them can be used to reduce the respective matching cost.

The point cloud as resulted from the triangulation of the respective matches revealed the sensitivity to the ray intersection angle and base length of cameras. For instance in the case of consecutive FLI-MAP images the theoretic standard deviation of triangulated points in viewing (depth) direction is – due to the small effective baseline – around 1m, but perpendicular to that – due to the low flying height – around 2cm only. In the shown examples these theoretic measures were confirmed. In the tested images from Pictometry the intersection geometry is better because of the longer baseline. In general, the overall structures on the building faces are well represented, but the noise reduction needs further attention.

In the current research the focus is put on the automatic detection and extraction of buildings in oblique images. Here, the point cloud as derived from the matching can give valuable cues. Another issue concerns the derivation of a more complete coverage by merging the point clouds as derived from different viewing directions. At least for the roof areas this can be done in a similar manner as shown above, namely through linking matches, since the majority of roof areas are visible from multiple directions.

ACKNOWLEDGEMENTS

I want to thank Daniel Oram for providing the code for general rectification on his homepage[6]. Further I like to thank Matthias Heinrichs, TU Berlin, for providing me with his code for Semi-Global-Matching. I also want to thank BLOM Aerofilms for providing the Pictometry dataset. Finally I would like to thank the anonymous reviewer for their comments.

REFERENCES

Albertz, J. and Kreiling, W., 1980. Photogrammetrisches Taschenbuch. 3. edn, Herbert Wichmann, Karlsruhe.

Besnerais, G. L., Sanfourche, M. and Champagnat, F., 2008. Dense height map estimation from oblique aerial image sequences. Computer Vision and Image Understanding 109(2), pp. 204 – 225.

Früh, C., Sammon, R. and Zakhor, A., 2004. Automated texture mapping of 3d city models with oblique aerial imagery. In: 3D Data Processing Visualization and Transmission, International Symposium on, pp. 396–403.

Gerke, M., 2008. Dense image matching in airborne video sequences. In: ISPRS: XXI congress : Silk road for information from imagery, Vol. XXXVII-B3b, International Archives of Photogrammetry and Remote Sensing, Beijing, pp. 639–644.

Gerke, M. and Nyaruhuma, A., 2009. Incorporating scene constraints into the triangulation of airborne oblique images. In: High-Resolution Earth Imaging for Geospatial Information, Vol. XXXVIII, International Archives of Photogrammetry and Remote Sensing, Hannover.

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

Heinrichs, M., Hellwich, O. and Rodehorst, V., 2007. Efficient Semi-Global Matching for trinocular stereo. In: IAPRS, Vol. 36, pp. 185–190. Part 3-W49A (PIA conference, Munich, Germany).

Hirschmüller, H., 2008. Stereo processing by Semi-Global Matching and Mutual Information. IEEE Transactions on Pattern Analysis and Machine Intelligence 30(2), pp. 328–341.

Hirschmüller, H., Scholten, F. and Hirzinger, G., 2005. Stereo vision based reconstruction of huge urban areas from an airborne pushbroom camera (HRSC). In: Lecture Notes in Computer Science: Pattern Recognition, Proceedings of the 27th DAGM Symposium, Vol. 3663, Springer, Berlin, pp. 58–66.

Höhle, J., 2008. Photogrammetric measurements in oblique aerial images. Photogrammetrie Fernerkundung Geoinformation 1, pp. 7–14.

Lemmen, M., Lemmen, C. and Wubbe, M., 2007. Pictometry : potentials for land administration. In: Proceedings of the 6th FIG regional conference, International Federation of Surveyors (FIG).

Lowe, D. G., 2004. Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision 60(2), pp. 91–110.

Mishra, P., Ofek, E. and Kimchi, G., 2008. Validation of vector data using oblique images. In: Proceedings of the 16th ACM SIGSPATIAL International conference on advances in Geographic Information Systems, ACM, Irvine, California.

Oram, D., 2001. Rectification for any epipolar geometry. In: Proceedings of the 12th British Machine Vision Conference (BMVC).

Petrie, G. and Walker, A. S., 2007. Airborne digital imaging technology: a new overview. The Photogrammetric Record 22(119), pp. 203–225.

Seitz, S. M., Curless, B., Diebel, J., Scharstein, D. and Szeliski, R., 2006. A comparison and evaluation of multi-view stereo reconstruction algorithms. CVPR 1, pp. 519–526.

Wang, M., Bai, H. and Hu, F., 2008. Automatic texture acquisition for 3d model using oblique aerial images. In: Intelligent Networks and Intelligent Systems, International Workshop on, pp. 495–498.

[6] http://www.aspa29.dsl.pipex.com (accessed 15 March 2009)