

The Landscape of Dense Image Matching Algorithms

NORBERT HAALA, Stuttgart

ABSTRACT

Both improvements in camera technology and the rise of new matching approaches triggered the development of suitable software tools for image based 3D reconstruction by research groups and vendors of photogrammetric software. Based on dense pixel-wise matching, the photogrammetric generation of dense 3D point clouds and Digital Surface Models from highly overlapping aerial images has become feasible. In order to evaluate the quality of these matching algorithms in terms of accuracy and reliability, the European Spatial Data Research Organisation (EuroSDR) started a benchmark on image based DSM generation in February 2013. This test is based on two representative image blocks, which were processed by different groups with different software systems. The results provided from the different groups give a profound insight to the landscape of dense matching algorithms and are used within the paper to evaluate the potential of image based photogrammetric data collection.

1. INTRODUCTION

Recent innovations in matching algorithms considerably improved the quality of elevation data, generated automatically from aerial images. Traditional stereo-matching, originally introduced more than two decades ago, usually applies feature based algorithms. These algorithms first extract feature points and then search the corresponding features in the overlapping images (Heipke 1993). The restriction to matches of selected points usually provides correspondences at high certainty. However, feature based matching was also introduced to avoid problems due to limited computational resources. In contrast, recent stereo algorithms aim on dense, pixel-wise matches. By these means 3D point clouds and Digital Surface Models (DSM) are generated at a resolution, which corresponds to the ground sampling distance GSD of the original images. To compute pixel matches even for regions with very limited texture, additional constraints are required. Local or window based algorithms like correlation use an implicit assumption of surface smoothness since they compute a constant parallax for a window with a certain number of pixels. In contrast, so-called global algorithms use an explicit formulation of this smoothness assumption, which is then solved a global optimization problem (Szeliski, 2010). One example is scanline optimization, which can be solved very efficiently by recursive algorithms. A very popular and well performing example is semi-global matching (Hirschmüller, 2008), which evaluates a cumulative cost function from multiple scanline directions. Especially when combined with sophisticated aggregation strategies it can produce accurate results very efficiently (Szeliski, 2010).

Software tools for image based 3D point cloud generation are currently developed by a number of research institutes and photogrammetric software vendors. As a result, the landscape of photogrammetric data processing is changing considerably. In order to document such rapid progress benchmarks have proven to be extremely useful. Well known examples which measure the performance of state-of-the-art matching algorithms are the Middlebury Stereo Vision Page (Scharstein, & Szeliski 2002) or the benchmark on multi-view stereo reconstruction (Seitz et.al. 2006). These benchmarks provide general purpose datasets through a platform, where results can be uploaded to automatically yield quality metrics of the respective approach. While these projects emerged from the Computer Vision community, the test on the performance of photogrammetric digital airborne camera systems (Cramer, 2010) was organized by the German society of

Photogrammetry, Remote Sensing and Geoinformation (DGPF). Within this project, the potential of photogrammetric 3D data capture from automatic image matching was already demonstrated (Haala et.al. 2010). Also in view of the rapid progress in software technology, the European Spatial Data Research Organisation (EuroSDR) started a follow-up initiative to evaluate the ongoing developments in image based DSM generation (EuroSDR, 2013). In February 2013 a project team for implementation and organization of a benchmark on image matching was established. Basic scope was the evaluation of 3D point clouds and DSM produced by the participants with their available software systems. This should allow the respective groups to demonstrate their state-of-the-art in high quality image based DSM generation. Furthermore, the benchmark should provide a platform for software developers and initiate an exchange between vendors of photogrammetric matching software and DSM producers like national mapping authorities.

The results of this benchmark give a profound insight to the potential of image matching approaches for DSM computation. They were discussed during the 2nd EuroSDR workshop on 'High Density Image Matching for DSM Computation' held from 13th to 14th June 2013 (Fritsch et. al., 2013). This workshop brought together software developers, distributors and users of dense matching software and thus provided a suitable platform for a first review the respective outcomes. The benchmark results presented by the participants gives an overview of current state of dense matching algorithms and demonstrates the current developments of image based DSM generation. As the head of the project, the author uses these results for preparation of the presented paper. The following section briefly describes the relevant test data sets in the first part and then introduces the participants and their software systems as well as the used hardware environment. Section 3 then compares and discusses the results from the different groups.

2. THE EUROSDR PROJECT ON DENSE IMAGE MATCHING

In order to limit the effort required by the potential participants for data processing the test was restricted to subsets of aerial image flights. Thus, two representative data sets were prepared for the benchmark. These data sets consisting of two aerial image sub-blocks provided different landuse and block geometry.

2.1. Data sets and deliverables

The first data set, Vaihingen/Enz was selected as an example for data usually collected during state-wide DSM generation at areas with varying landuse. As it is visible from the ortho image and the DSM of the test area in figure 1 it covers a semi-rural area at undulating terrain. Elevation differences of 200m occur between the river valley and the upper area. The data are a subset from a flight collected during the project on Digital Airborne Camera Evaluation of the German Society of Photogrammetry, Remote Sensing and Geoinformation (DGPF) (Cramer, 2009). Both ground sampling distance and image overlap of the image block are rather moderate. This situation is typical for most flights captured for national mapping agencies. The aerial images were collected by an UltraCam-X at height above ground of 2900m and a ground sampling distance (GSD) of 20cm. For the image matching benchmark PAN images were used. These images were made available as Tiled Tiff uncompressed 8 bit/pix with 9420x14430pix at 136Mpix/image. The sub-block selected for the benchmark is depicted in figure 2. It consists of 3 strips with 12 images each. Figure 2 additionally visualizes the corresponding 36 camera stations, which are represented by the blue points. The respective image footprints are depicted by rectangles in the same color. During the test a DSM had to be generated for the central part of the block. This part is highlighted by the rectangle with light blue outlines. This area corresponds to the ortho image and DSM already shown in figure 1. It has a




Figure 1: Ortho image and DSM of test area Vaihingen/Enz.

size of 7.5kmx3.0km. Participants of the benchmark were asked to generate a DSM raster at a grid with of 0.2m. For the test area this resulted in a DSM raster of 37500x15000pix.

One of the main factors which influence the quality of image based surface reconstruction is the amount of image overlap. In figure 2 the number of images available for each terrain point is represented by the color coded map. For the Vaihingen/Enz sub-block, the available overlap of 63% in flight and 62% cross flight results in variations from one-folded areas to nine-folded overlap. In figure 3 these areas are represented in red and dark green, respectively. As it is also visible, usually four to nine images were available for the central part of the block, where the DSM had to be generated.




Figure 2: Vaihingen/Enz: Image overlap (maximum nine-folded) with image footprints and camera stations.




Figure 3: München: Color-coded image overlap with image footprints and camera stations(left) and DSM (right).



Figure 4: Visualisation of München.

In contrast to the first test area Vaihingen/Enz, which mainly features semi-rural landscape with different classes of landuse, the second test data set is more typical for applications in densely built-up urban areas. For such applications, images are usually captured at a higher overlap and resolution. This urban data set München covers the central part of the city and was captured by DMC II 230 at a GSD of 10cm. Each image has a size of 220 Mpixel/image (15552x14144 pix) at 16 bit/pix. The image sub-block to be processed consists of 3 image strips with 5 images each. In the left part of Figure 3, the camera stations are represented by dark blue triangles. Similarly to Figure 2, the corresponding image footprints are shown as dark blue rectangles. The area to be processed is again highlighted in light blue. It has a size of 1.5x1.7km. The image block was captured with 80% in flight and 80% cross flight overlap. This results in in up to fifteen-folded object points, which of course provide a considerable redundancy. In Figure 3 (left) this maximum

overlap is depicted in dark green. However, this large overlap is only available in the central part of the test area. It reduces to five images per object pixel at the border of the area to be processed. The aspired DSM grid width was 10cm. As it is visible for the grey coded representation of the complete area in Figure 3 (right) and the 3D visualization for a part of the area in Figure 4, the terrain is rather flat. However, the area is densely covered by buildings with heights of up to 50m. These buildings result in occlusions especially for surface parts close to the facades. Thus, visibility can be limited for such regions, which will potentially aggravate the matching processes during DSM generation.

2.2. Test participants, investigated software systems and used hardware environment

To address a suitable number of participants, the test procedure for the benchmark was kept relatively simple. Since the main focus is on dense image matching all participants had to use the orientation parameters made available for the image blocks without modification. No 3D point clouds were taken into account, the evaluation was limited to DSM raster in predefined size and resolution. Participants were asked to send their results to the project team for comparison and accuracy analysis. The quality of a software solution also depends on the computational efficiency, which considerably influences the processing time required. This is not only influenced by the implemented algorithm, but of course also results from the used hardware environment. This information was retrieved from the test participants by a questionnaire.

The participants of the test can be divided into users and vendors of commercial photogrammetric software systems and research institutes using software systems developed in house for own production and projects. However, frequently also software systems from research institutes are made available for the public under different license models. Almost all participants provided further information on their processing environment and strategy during the workshop on high density image matching for DSM computation (Fritsch et. al., 2013). These presentations are summarized in the following.

2.2.1. Users and vendors of commercial photogrammetric software system

- SocetSet 5.6 (NGATE) from BAE Systems was presented by C. Ginzler from Swiss Federal Institute for Forest, Snow and Landscape Research (WSL). He used a Intel Xeon CPU X5570 2.93 GHz and 24 GB memory, the required time for Vaihingen/Enz was 36h and 25h for München.
- B. Brunner from Forest Mapping and Management (FMM) Salzburg presented results from Microsoft UltraMap V3.1. This software can only be used from imagery captured by the UltraCam product family. Therefore, investigations were restricted to the Vaihingen/Enz data set. Furthermore, processing has to pass through all software modules. Since this includes Automatic Aerial Triangulation and color adjustment, the UltraMap results are based on an individual bundle block adjustment. Furthermore instead of the 8bit images provided to the participants the originally 16bit imagery was used. The IT environment consisted of 32 Xeon E5-2630/i7 cores and 5 GPU's (1 Tesla K10, 4 Tesla M2090). For the 36 images of the Vaihingen/Enz block this comparable powerful environment resulted in a processing time of about 35min for data ingest and AAT while the DSM was computed in 27min.
- Match-T DSM 5.5 from Trimble/inpho was presented by C. Ressler from the Department of Geodesy and Geoinformation (GEO) at TU Wien. He used an Intel Core i7 CPU, 3GHz with 4 cores and 8GB memory. Match-T itself matches only one image pair for each XY-location. To exploit the high image overlap all possible image pairs were matched during his investigations. The resulting DSMs were then fused using the in-house software systems Opals.

Processing time for Vaihingen/Enz resulted in 23h for matching with Match-T and 38h for data import, gridding and fusion with Opals. For the München data set the numbers were 19h for matching with Match-T and 37h for processing with Opals.

- The ImageStation ISAE-Ext software was demonstrated by R. Schneider from GEOSYSTEMS GmbH, Germany. For the München test area processing was limited to 6 stereo models. The hardware environment of 2 Intel Xeon Quadcore Processors with 16GB RAM a computation time of 40min per stereo model and 4h for the complete block was achieved. For Vaihingen/Enz 33 stereo models were processed at an average computation time per stereo model of 20min. This resulted in 11h for stereo matching in total, followed by 21h for grid interpolation of the 580 million points to the required 20cm raster.
- P. Nonin from Astrium GEO-Information Services presented results for the Pixel Factory software. For his investigations a hardware configuration of 2 Xeon E5-2640, 2.5GHz with 2x6 cores was used. Results for the München area are based on the processing of 12 stereo pairs. Due to expected problems from moving shadows no cross-track matching was applied for this scene. Total time for DSM generation was 2h 12min. For Vaihingen/Enz 57 stereo pairs were processed in a total time of 3h 9min.

2.2.2. Research institutes with own software developments

- M. Idrissa from the Royal Military Academy (RMA), Brussels showed results for the RMA DSM Tool. For processing a Linux Cluster with more than 90 Fedora CPUs (Intel – 2.4 GHz) and more than 30 nodes with 272GB RAM (total) and a 1000Mbit/s network was used. Processing times for München was approximately 18min per stereo couple and about 5h for the complete block, for Vaihingen/Enz the numbers were 8min per stereo pair and again 5h in total.
- Results of the Remote Sensing software package from Joanneum Research, Graz, were presented by K. Gutjahr. He used a Windows PC Intel Xeon CPU E5-2650, 2.0 GHz with 16 Cores and 32GB RAM for processing. Processing time for Vaihingen/Enz for 57 stereo pairs summed up to 17h 17min for the processing steps preparation, normalization and prediction, matching, forward intersection, DSM generation and DSM finalization. Accordingly the processing time for München with 22 stereo pairs was 21h 39min.
- The software system MicMac developed at IGN France was presented by M. Pierrot-Deseilligny. Within MicMac multi image matching is realized by energy minimization either using graph cut or a modified dynamic programming/SGM approach. Vaihingen/Enz was processed within 6h and München in 12h.
- M. Rothermel presented results from the SURE software developed at the Institute for Photogrammetry (IfP), University of Stuttgart. It is based on a SGM variant with integrated coarse-to-fine strategy based on efficient implementation of SGM on dynamic data structures. He used a PC with one i7 CPU quadcore at 3.4GHz and 32GB Ram. For the München data set in total 4h 13min were required. The processing time for Vaihingen/Enz summed up to 4h 37min. For an alternative implementation using an NVIDIA 580 GTX GPU and an i3 CPU, dual-core the Vaihingen/Enz data set was processed in 2h 50min.
- H. Hirschmüller, German Aerospace Center (DLR) provided results from a FPGA implementation of his SGM algorithm. Processing was realized on a Intel Xeon E5-1620 quad core PC. However, for processing a Virtex 6 FPGA board with 8GB Ram, connected to the PC via GigE was applied. With this configuration a processing time for the München data set of 5h 51min and 19h 31min for Vaihingen/Enz was recorded.




This compilation shows that very heterogeneous hardware environments were used for processing. It spread from standard desktop PCs with single to multi core to the use of low-end to multiple

high-end graphics cards till the application of computer clusters. Since a considerable amount of data has to be processed, time required for DSM computation is additionally influenced by reading and writing operations. Thus, efficiency of data ingest also depends on the available storage system and network environment. Finally, within one software solution processing time can be influenced by following different strategies e.g. while compiling the number of stereo-pairs to be processed. This is the reason why the reported processing times should not be regarded as fix number but as a hint to the capacity of area covering data collection.

3. EVALUATION OF DSM QUALITY

For quality assessment of the DSM data sets provided from the different participants, the comparison to a reference surface is probably most suitable. This reference surface can in principle be provided from LiDAR measurements. However, elevation differences between data from image matching and airborne LiDAR cannot exclusively be put do errors of the matching process. As an example, differences as a results from plant growth and harvesting are to be expected due to the time gap of some weeks between the image and the LiDAR flights within the Vaihingen/Enz test area (Haala et.al. 2010). Height variations can additionally result from the different measurement principles. Light pulses from airborne laser scanning partially penetrate a tree canopy, while matching will most probably relate to the visible surface. Finally, pixel-wise matching reconstructs surface representations at a resolution, which is frequently not available from standard LiDAR flights. As an example, LiDAR data was available for Vaihingen/Enz with a median point density of 6.7 pts/m². For the München area the available LiDAR points have a density of 4 pts/m². In contrast, the amount of detail to evaluate the DSMs from image matching is much higher. In the test, a grid width corresponding to the GSD of the respective imagery had to be generated. Thus, the 20cm raster of Vaihingen/Enz results in 25pts/m² while the 10cm grid for München corresponds to 100pts/m². Due to these problems in providing suitable ground truth from independent measurements, an alternative approach was used. Based on the results of the 10 participants, a median DSM was generated. Of course this median DSM does not provide independent ground truth at higher order accuracy. Still, it can be used very well to illustrate differences between the respective solutions.

These differences are depicted for a part of the Vaihingen/Enz area in Figure 5. The top left image shows an ortho image of the area. A shaded representation of the computed median DSM is given in the top right. The remaining images provide a color coded representation of the DSM differences to this median, which are overlaid to a shaded relief of the generated DSM. Despite the fact that larger differences are available, color coding was always limited to differences between -2m and +2m. The legend is given in the caption of the figure. As it is visible in Figure 5, larger differences to the median surface occur for a number of DSM in the shaded area of the quarry. Differences are also available in the river area, as well as in the vicinity of patches of trees and in the area of the vineyards. In order to allow a quantitative analysis of DSM differences especially in these areas, elevation profiles were extracted for further analysis. These lines of interest are overlaid to the ortho images in the top left of Figure 5. The profile line numbered with label 0 is extracted from the respective surface models at the quarry area, which is subject to a cast shadow. Profile line 1 is located at another steep slope of the quarry, which in contrast is well illuminated. Furthermore, line 1 covers some vineyards. Apparently, this type of landuse results in larger differences between the respective solutions. Profile 2 is situated at an area with patches of trees. Profile 3 has its starting point close to a building and then crosses a river area. This area again is subject to larger differences of the DSMs depicted in Figure 5.






Figure 5: DSM differences for Vaihingen/Enz, color coded between -2m  +2m.

The resulting height profiles are shown in Figure 6. The top left of Figure 6 depicts elevations of profile 0 for data from IfP SURE, Astrium, Match-T, SocetSet, MicMac, GeoSystemes, DLR, RMA, Joanneum and UltraMap. Apparently, all elevations are very similar for the area of the quarry, which is well illuminated. However, large differences are visible for the area subject to cast shadow. Profiles extracted for GeoSystemes (dashed blue line) and MicMac (solid black) apparently are higher than the median DSM. Furthermore, the profile from the SocetSet DSM (solid green) contains a considerable amount of no-data areas. The RMA profile (dashed light blue) is relatively noisy, while for Astrium (solid red) a certain smoothing can be observed. The first part of profile 1 is extracted at another steep slope of the quarry. In contrast to profile 0 this part is well illuminated. Except for a no-data area from SocetSet, apparently all profiles are relatively similar. The left diagram in the middle row of Figure 6 shows a more detailed view of the profiles 1. The corresponding part of the ortho image is shown right to that profile. As it is visible the line covers some rows of vines. These objects have a size close to the GSD of the available imagery. Still, some solutions like Astrium (solid red), MicMac (solid black), and IfP-SURE (solid blue) at least partially resolve these structures. Furthermore, almost all solutions are able to reconstruct the shape of the small hut. Apparently, at object surfaces not subject to problems of image matching, differences between almost all solutions are in the pixel level of 20cm. An additional profile situated at an area with patches of trees is shown in the bottom left of Figure 6. This profile 2 shows some differences between the respective solutions mainly at tree borders. However, the respective results are remarkably consistent. Still, the profile from DLR (dashed red) has a constant offset compared to all other solutions. This is also visible in the difference image of Figure 5. The offset did not occur for a DSM additionally provided by DLR from a CPU based solution. However, the supposed problem in the investigated FPGA solution could not be fixed before finalizing this paper.

The last profile labeled by 3 starts close to a building and then crosses the river. Since such surfaces are almost impossible to reconstruct from image matching, large differences occur. The solution from UltraMap (dashed black) is remarkably smooth.





Figure 6: DSM profiles Vaihingen/Enz.

DSM results for the test area München are presented in Figure 7. Again a subsection was selected for further analysis. The corresponding part of the ortho image is depicted on top left of Figure 7. Profiles which were subject to further analysis are overlaid, again. As already discussed, the UltraMap software is not able to process imagery from the DMC II camera. Due to this reason, results were only available from the remaining 9 software solutions. Again, a median DSM was computed as reference surface and the respective differences to this surface were color coded between the values -2m and +2m.

As it is visible in Figure 7, the solution from SocetSet considerably deviates from all other data sets for this urban area. Differences between the remaining solutions are mainly visible at small details and steep edges which occur in the vicinity of building borders. Furthermore, cast shadows seem to result in differences and increasing noise for the reconstructed surfaces. Thus, profiles were extracted at such areas for further analysis.






Figure 7: DSM differences to Median surface München, color coded between -2m  +2m.

The extracted profiles for test area München are given in Figure 8. Due to the large differences of the solution from SocetSet already visible in Figure 7, this data set was not included in these diagrams. By these means the visibility of the remaining results could be improved. The top left of Figure 8 shows results for profile 0. It was extracted at a building with a distinct façade structure in the front. The profile then continues at the top of the roof along some dormers. The first part of the profile is on the ground which is covered by a cast shadow. Especially in this area the solutions from MicMac (solid black) and RMA (dashed light blue) appear relatively noisy. At the step edge defined by the façade Match-T (solid light blue) deviates from the other results. A close up view for profile 0 which covers two dormers is shown in the top right of Figure 8. As it is visible almost all solutions seem to reconstruct these structures correctly. The profile with label 1 depicted in the bottom left of Figure 8 was extracted at a fountain. The fine structures of this object should be suitable for investigation of object parts at a size close to the GSD of 10cm for the processed imagery. Apparently, this structure is captured by almost all software solutions. Furthermore, for the well illuminated terrain surface the differences between all solutions are at the 1 pixel level except for the relatively noisy DSM provided by MicMac. The final example is Profile 2, depicted bottom right in Figure 8. As it is also visible in the difference DSMs in Figure 7, problems seem to occur for the steep roof surface, which is subject to a cast shadow. Apparently, the availability of

shadow increases the noise level of this roof surfaces for a number of software systems. However, the profiles show larger deviations only for the DSM from GeoSystems (dashed blue).




Figure 8: DSM profiles München.

4. CONCLUSIONS

The investigations presented in the paper clearly show, that a growing number of software tools for detailed, reliable and accurate image based DSM generation from airborne imagery are available. Processing not only benefits from improved algorithms for efficient stereo image matching but takes advantage from large image overlaps in order to efficiently eliminate erroneous matches. This provides a considerable reliability of DSM at vertical accuracies close to the sub-pixel level. The results of the EuroSDR benchmark additionally show acceptable run-times for a number of software systems even if a standard hardware environment is used. Data sets like the München area with 15 images at 220MPixel/Image for generating 10cm DSM grids or the Vaihingen/Enz area with 36 images and 136MPixel/image for 20cm DSM grids can be processed without problems.

The interpretation of the benchmark results identified some scenarios which still can cause some problems during image based surface reconstruction. Some solutions showed decreasing accuracies at cast shadows. Differences between the respective results also increased at fine object structures close to the resolution of the available images. However, a clear ranking for the respective results was avoided – also due to the fact that current software development for image matching is subject

to a considerable momentum. Thus further improvements can be expected both for matching accuracy and computational performance.

5. ACKNOWLEDGEMENTS

Sincere thanks go to my co-members of the project team, Wolfgang Stoessel, State Agency of Surveying and Geoinformation Bavaria, Dr. Michael Gruber, Microsoft Photogrammetry, Prof. Dr. Norbert Pfeifer, Vienna University of Technology, and Michael Franzen, Federal Office for Metrology and Surveying, Vienna. The most important prerequisite for realization of the benchmark on high density image matching for DSM computation is a sufficient number of participants volunteering for data processing. This work is especially acknowledged – without their valuable contribution such an initiative cannot be implemented.

6. REFERENCES

- Cramer, M. (2010): The DGPF-Test on Digital Airborne Camera Evaluation – Overview and Test Design. *Photogrammetrie – Fernerkundung – Geoinformation (PFG)*, Heft 2 (2010), pp. 75-84.
- EuroSDR (2013): <http://www.ifp.uni-stuttgart.de/euroedr/ImageMatching/> web-page of the EuroSDR project Benchmark on Image Matching, last accessed July, 2013.
- Fritsch, D., Pfeifer, N. & Franzen, M. (Eds.) (2013): Proceedings of the 2nd EuroSDR workshop on High Density Image Matching for DSM Computation. EuroSDR Publication Series, No. 63.
- Haala, N., Hastedt, H., Wolff, K., Ressel, C. & Baltrusch, S. (2010): Digital Photogrammetric Camera Evaluation – Generation of Digital Elevation Models. *Photogrammetrie – Fernerkundung – Geoinformation (PFG)*, Heft 2, pp. 99-115.
- Heipke, C. (1993): Performance and state-of-the-art of digital stereo processing. Proceedings of the 44th Photogrammetric Week, Stuttgart, pp. 173-183.
- 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.
- Scharstein, D. & Szeliski, R. (2002): A Taxonomy and Evaluation of Dense Two-Frame Stereo Correspondence Algorithms. *International Journal of Computer Vision*, Volume 47, Issue 1-3, April-June 2002, pp. 7-42.
- Seitz, S., Curless, B., Diebel, J., Scharstein, D. & Szeliski, R. (2006): A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'2006)*.
- Szeliski, R. (2010): *Computer Vision: Algorithms and Applications*. Springer, Chap. 11, Stereo Correspondence.