Thomas Zwölfer

Implementation and Evaluation of the DAISY Descriptor for a Semi-Global Matching framework

Betreuer: Dipl.-Ing. Mathias Rothermel

Prüfer: Prof. Dr.-Ing. Norbert Haala

Dauer der Arbeit: 3 Monate
Abschluss: September 2012

Introduction

In the past years dense image matching has encountered a big growth in popularity in Photogrammetry. Numerous investigations have shown its potential as an addition to more established methods of 3D-Data generation, like airborne or terrestrial laser scanning.
A popular algorithm is the Semi-Global Matching approach, which is also used in this investigation.
In order to obtain initial information about pixel correspondences, one needs so-called matching cost functions.
For the case of a short-baseline and highly overlapping images, matching cost functions like Census and Mutual Information have been developed and proven to be highly effective. In contrast to that, wide-baselines have been shown to be more challenging. A larger amount of perspective distortion and bigger occluded areas in the used images contribute to the difficulties. Nevertheless, using wide-baseline image-pairs is worth the consideration, as the triangulation into object space is more accurate in Z-direction, due to the improved intersection angle.

For this purpose Engin Tola, of Ecole Polytechnique Federale de Lausanne (EPFL), introduced the DAISY descriptor, which is supposed to combine robustness of descriptors like Lowe's SIFT with computational efficiency. By integrating a cost function based on the DAISY descriptor into a Semi-Global Matching framework, this investigation intends to evaluate achievable results.

Figure 1: Image taken from EPFL Fountain100 dataset with accompanying depth map.

Investigation

The existing SGM algorithm was extended with a cost function based on the DAISY descriptor.
It was implemented using the open source DAISY library, provided by EPFL.
This way matching strategies relying on different matching cost functions could be applied to input imagery.
Evaluation of the image matching algorithms is done by comparing depth maps, generated by the image matching pipeline, and LiDAR reference depth maps. The evaluation criteria were accuracy of the depth data and completeness of the reconstructed scene.
Different matching scenarios were used in order to evaluate the performance based on growing image baselines.

Figure 2: Depth maps generated by dense image matching. The left image shows the result of the DAISY based matching strategy. The right image is produced with a combination of Mutual Information and Census matching functions.

Results

Four different matching strategies were applied, some of which utilize a combination of matching cost functions on different pyramid levels, taking advantage of the modular implementation of the SGM algorithm.
Matching strategies, that include DAISY, show a big improvement in completeness. In some cases DAISY based matching strategies were able to reconstruct 50% more points than the other matching strategies.
Regarding accuracy an increasing accuracy of the matches with a growing intersection angle is apparent. This was expected because of the improving geometrical configuration. The matching strategies based on DAISY yield an accuracy that is slightly worse than the other matching strategies.

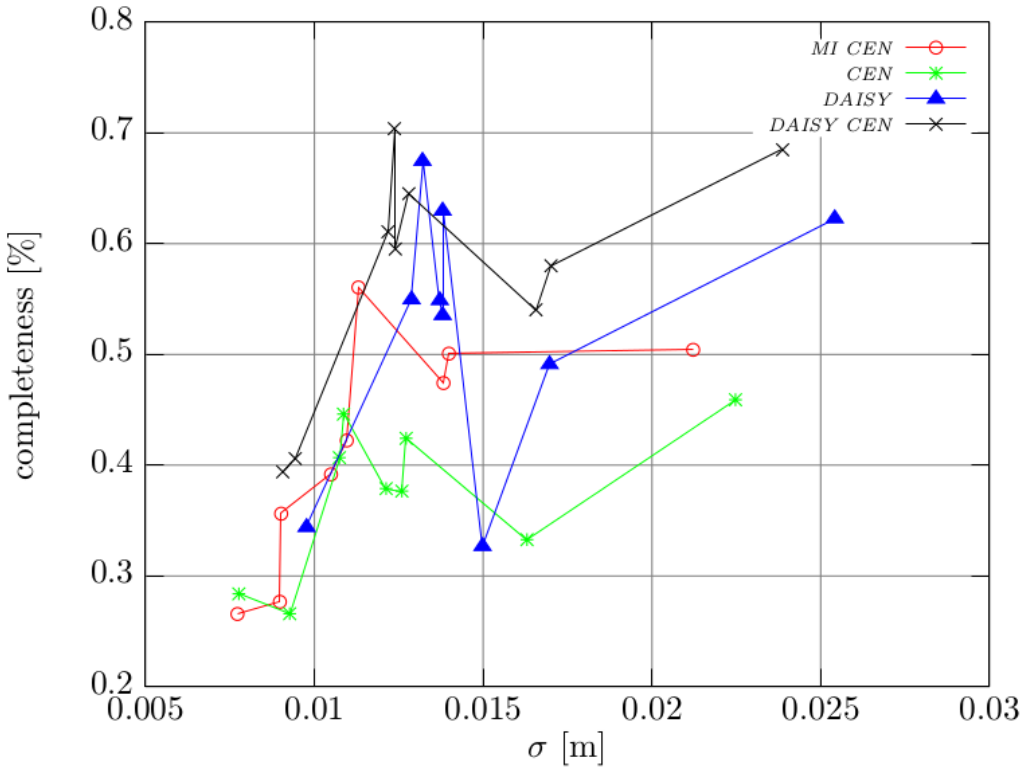

Figure 3: The standard deviation σ can be seen on the x-axis, the completeness in % of possible matches on the y-axis.

In order to judge the effectiveness of the matching costs, the values for completeness in relation to accuracy were investigated (Figure 3). Using DAISY, one is able to produce rather dense pointclouds of scenarios with wide baselines, which allow for higher triangulation accuracy.

Vice versa, Mutual Information and Census also reach this level of accuracy in rather wide baseline scenarios, but at a lower rate of completeness.

Conclusion

It was shown that concerning completeness and therefore density of the resulting point clouds

DAISY outperforms the other matching costs by up to 50% for some wide baseline scenarios. Especially by using a combination of an initialization on the lowest pyramid level by Census

with the good performance of DAISY on the higher pyramid levels, very dense results can be obtained.

Regarding accuracy, the behavior of DAISY is slightly worse than the other investigated matching costs. Accuracy increases with a growing baseline,but it does not completely reach the accuracy of Mutual Information and Census.

Conclusively, DAISY proves to be a very promising way of calculating a pixel-wise matching cost. With only little adjustment to the actual Semi-Global Matching algorithm, remarkable results regarding density of pointclouds have been obtained.