Robustness of Visual Localization Methods

LOCATE - Visual Localization in Natural Environments

Jan Brejcha and Martin Čadík

Faculty of Information Technology, Brno University of Technology

Project LOCATE deals with localization in natural environments. It proposes the investigation of image to model registration techniques allowing for building accurate visual localization systems that are more robust to qualities of input photos. Particular attention is devoted to finding applications of visual localization systems in challenging outdoor scenarios.

Main aim of our actual work is the comparison of our edges-to-silhouettes matching method [1] with state-of-the-art approach based on horizon curve fitting [2]. According to our results, horizon-based method [2] has significant disadvantages mainly on places, where the horizon curve is ill-defined, e.g. the horizon watched from high mountain peak is almost flat and hence not descriptive enough.

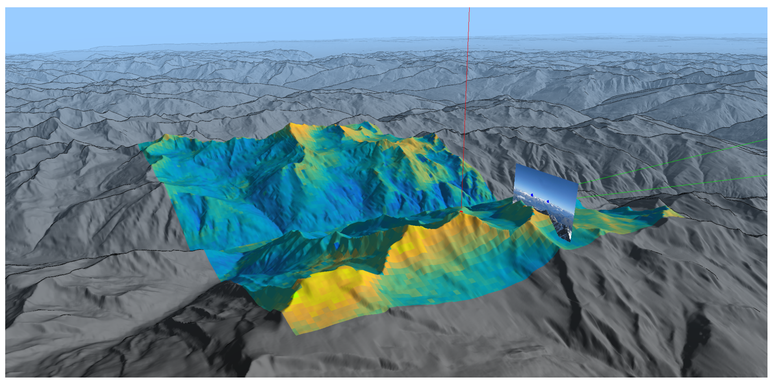

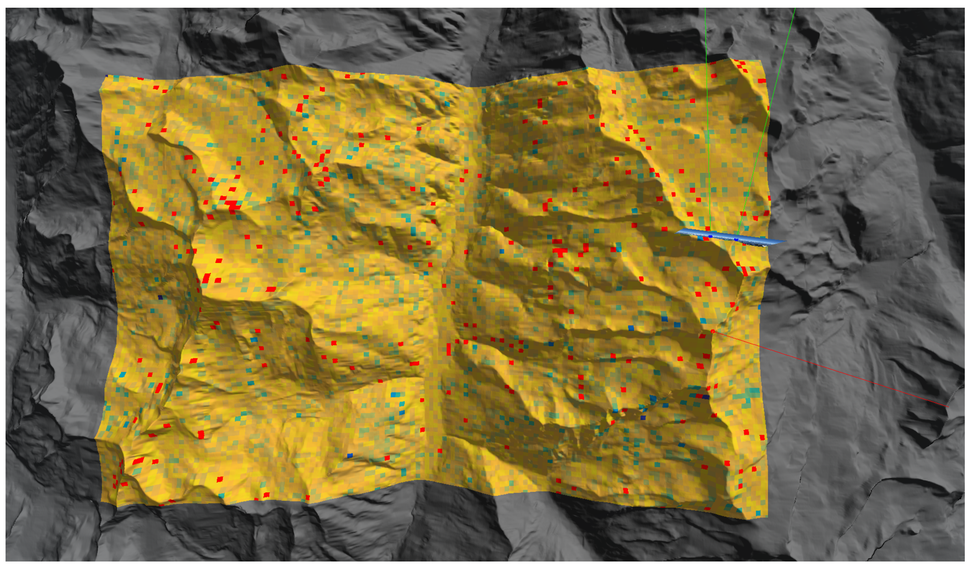

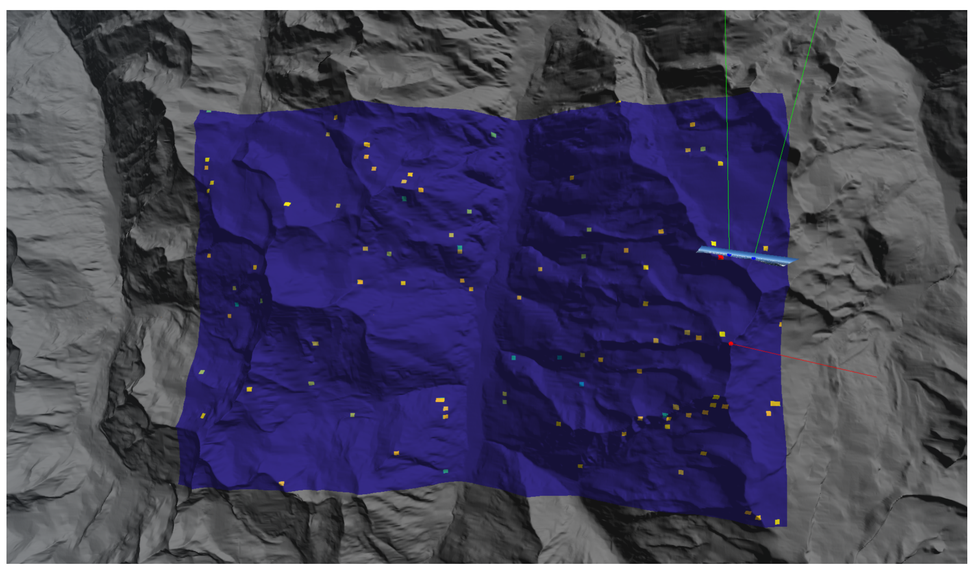

Horizon curve fitting method [2] was reimplemented and studied. Our test dataset spans the area of 15 x 10km in Swiss Alps. The search space was uniformly sampled by 0.0015° in west-east, and 0.001° in north-south direction respectively totalling 12693 places. In each place a panorama is rendered 1.8m above the surface. For [2] the horizon curves from each panorama were extracted to obtain local features from which the Bag-Of-Words database has been created.

The visualisations show that the tested methods deliver different results. Method [2] finds a lot of candidates that are dispersed randomly throughout the whole search space. This is good result only in case that enough correct candidates are chosen for the geometric check. Then the burden of chosing the right result is only on the geometric verifier (ICP). The bright side of [2] is that during our experiments often at least some good matches at the first positions were found near the ground truth.

For [2] only direction and location voting in a Bag-Of-Words database was used. The left image shows the results for reimplementation of [2] without geometric check, the results with geometric check are shown in the right image. The geometric check is run on only 100 of best candidates.

The results are visualised using a heatmap. More the red color, the better score. The pallete for the heatmap can be seen below. Ground truth is marked by the red line with red ball above the surface.

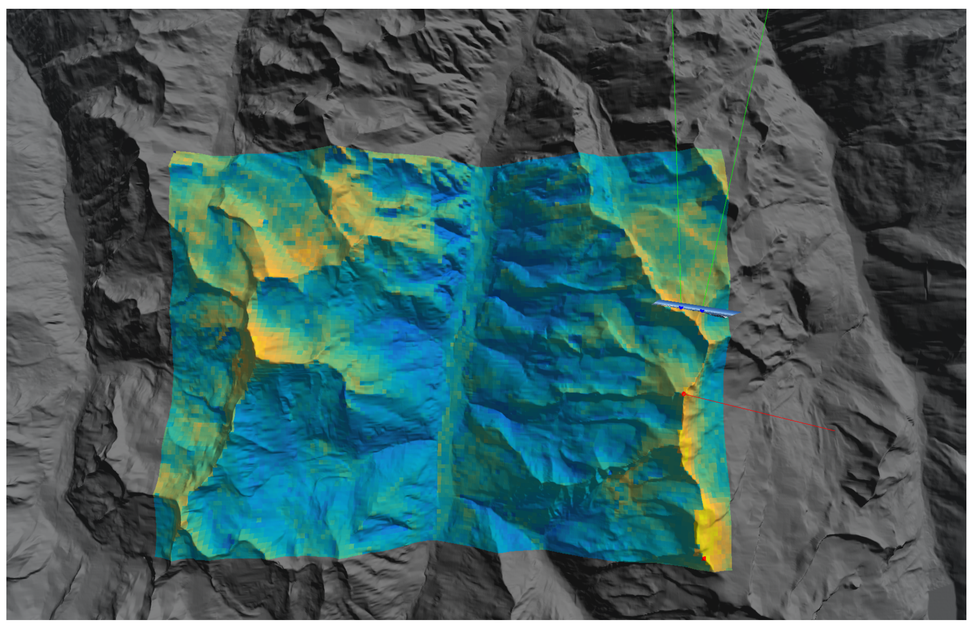

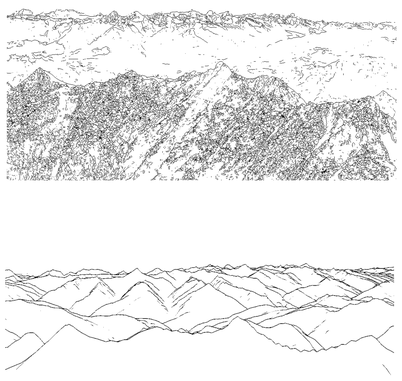

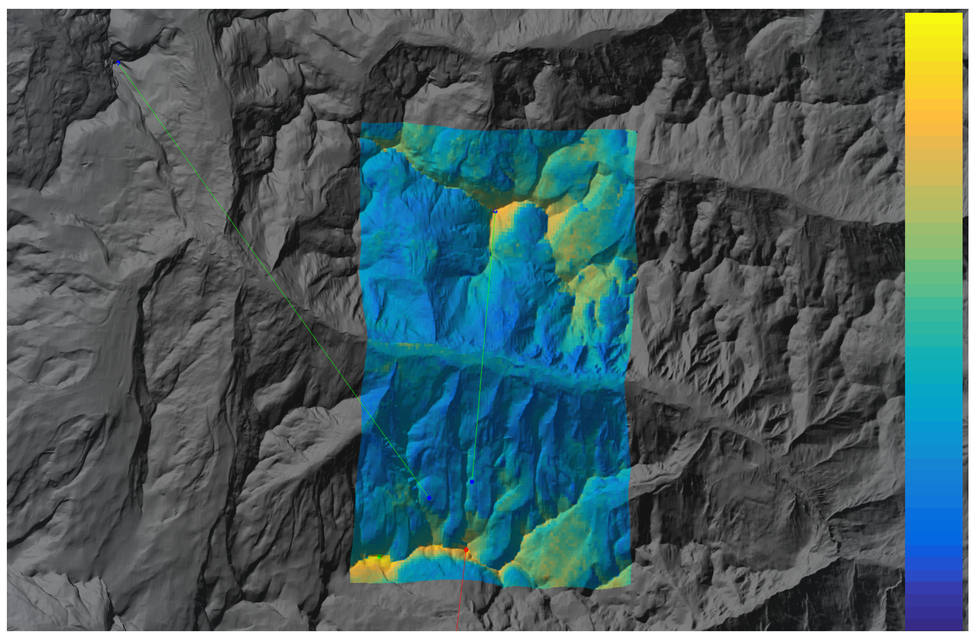

Results from the cross-correlation part of edges-to-silhouettes matching [1] give better distribution of possible candiates. Thanks to the geometric verification we could process better candidates and therefore there is a higher chance of successfull geo-localization.

In further research, the vectors from the cross-correlation will be visualised to better synthetic silhouettes over the results. The main downside of the cross-correlation is the computational complexity. Therefore faster approximations are seached.

[1] Baboud L., Čadík M., Eisemann E., and Seidel H.-P. Automatic photo-to-terrain alignment for the annotation of mountain pictures. In Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition, pages 41–48, 2011.

[2] Baatz G., Sauer O., Köser K., and Pollefeys M. Large scale visual geo-localization of images in mountainous terrain. In Proceedings of the 12th European conference on Computer Vision, pages 517–530, 2012.

The project (LOCATE 4SGA8694) is financed from the SoMoPro II programme. The research leading to this invention has acquired a financial grant from the People Programme (Marie Curie action) of the Seventh Framework Programme of EU according to the REA Grant Agreement No. 291782. The research is further co-financed by the South-Moravian Region. This work reflects only the author's view and the European Union is not liable for any use that may be made of the information contained therein.