České vysoké učení v Praze

Professor Jiri Matas

Center for Machine Perception
Department of Cybernetics
Faculty of Electrical Engineering
Czech Technical University,
Karlovo namesti 13,
121 35 Prague 2, Czech Republic
phone: +420 603 140 180

web: http://cmp.felk.cvut.cz/~matas/
email: matas@cmp.felk.cvut.cz

August 10, 2021

Review of the PhD Thesis of Jan Brejcha

The title of the doctoral thesis is "Visual Localization in Natural Environments". The term "natural environments" does not capture the content of the thesis well, since it solely focusses on localisation in mountainous terrains. The localisation is developed and evaluated for a particular class of photos, which could be described, from the examples in the thesis, as "broad view, good acquisition conditions, daytime, large area depicted". The exact scope of the thesis could have been specified more precisely, in the abstract and the introduction.

The problem of visual localisation in mountainous environments, even in the slightly restricted conditions mentioned above, is a challenging open problem. The thesis provides a thorough analysis of the state of the art, including information about relevant data that are essential for both development and evaluation of novel methods.

The contributions of the thesis cover diverse aspects of the localisation problem, presenting a new dataset and novel ideas and algorithms, with excellent results. On the annotated dataset side, the contribution is impressive both in quality of the annotations provide and the size of the dataset. On the algorithmic part, the thesis contributes significantly to an instance of cross-modal matching, mainly focusing on the use of Digital Elevation Models (DEM), but also proposing a method that exploits orthophoto maps. The thesis provides a detailed and readable presentation of the proposed methods.

The main contributions of the thesis were published in papers presented at quality conferences and good journals, including ECCV a top-tier, highly selective conference. Acceptance of the papers that roughly correspond to chapters of the thesis is a clear indicator of the contribution of the thesis to the state-of-the-art.

I will provide the candidate with a copy of the PhD thesis in .pdf format with a large number of comments, typically referring to typos or small technical issues that I do not repeat in the review. The comments do not affect my highly positive assessment of the thesis.

Possible question for the defence:

- Could you give a technical specification of photos that are likely to be localised correctly by the methods presented in the thesis? Such information is very useful for a potential user.

- Please discuss the characteristic of the photos in the GeoPose3K dataset. They seem to be restricted to good weather conditions and views that include many ridgelines. This could be quite a small subset of photos acquired in mountainous terrain.
- A large number of static or pan-tilt-zoom cameras are installed in natural environments, their GPS location is known, the data is publicly accessible. Have you considered testing the method on such data? Often, the cameras operate 24/7 and as such are highly suitable for testing robustness of a localization algorithm w.r.t. the lighting conditions, time of day, season, and the weather
- How sensitive is the localisation method that relies on semantic segmentation to the season in which the query photo was taken? Have you conducted any quantitative experiment to assess this?
- The matching score for edges, formula (4.1), seems arbitrary, with parameters that might not be easy to set. Please explain why this score function was chosen.
- The edge-based matching against DEM ridgelines might work at very low levels of light (late afternoon, early morning, full moon). Have you conducted any experiments to check this?
- Fig. 6.2 . First, would discriminant correlation filter work better? Second, the section suggests that the change of a pattern to x' = const – x matters, and a combination of a positive and negative patterns is superior to std. correlation. If this is the case, why is this not a standard, well-known procedure?
- Chapter 7. Have you considered mapping the whole image to the other domain, rather than only the descriptors?

In conclusion: the thesis contributes to the state-of-the-art. The ideas are presented in well-structured manner; the claims are supported by experimental evidence.

I recommend the thesis be accepted.

Prof. Jiri Matas