Master thesis evaluation
"Unsupervised Learning of Semantic Landmarks for Visual Navigation over Extended Periods of Time"

Author: Jan Blaha Opponent: Keerthy Kusumam

The thesis focuses on unsupervised learning of semantic landmarks for the long-term deployment of a visual teach-and-repeat navigation system. The student lays out the problems encountered by teach and repeat systems when deployed for extended periods of time, where the changes in the environment causes local feature-based registration methods to fail. The student proposes three different approaches for semantic landmark discovery.

The assignment was challenging because the topic is still investigated in scientific literature. Automated generation of annotations for unsupervised landmark detection methods is difficult by itself, which is exacerbated by the fact that the images are captured from a mobile platform in realistic outdoor conditions. The thesis is excellently motivated and the state of the art is very thorough and well written. Mainly due to the extensive literature survey, the student was able to identify and conceptualise three methods based on different principles. The thesis shows that the student worked in a very systematic manner and every step towards the thesis goals was well motivated. The thesis structure and content is of high technical quality. The description of the methods demonstrates a clear understanding of the problem and it's obvious that the student has a solid theoretical background and he can apply it accordingly. Mathematical formalism is clear and contributes to the quality of the thesis. However, the language of some sections (e.g. the section on temporal properties) becomes obfuscated and hard to follow. This makes comprehension of the advanced concepts rather difficult without re-reading the relevant sections. However, this is alleviated by the excellent command of the English lanquage.

Since the student performed an extensive survey and designed, implemented and evaluated three different landmark discovery pipelines, and tested them on a real robot, the goals of the thesis were not only only achieved, but surpassed.

Therefore, I propose to classify the thesis as

A - excellent.

Cambridge UK,
June 8, 2022

Keerthy Kusumam
University of Nottingham, UK

Questions

1. I would like to ask about details on the clustering methods used in Section 3.1.3. In particular, how does figure 10 contribute to understanding of the 'clever' clustering mentioned at the end of Section 3.1.3?

2. The overall functionality of the navigation systems based on the proposed methods is based on the so-called 'convergence theorem'. How is the convergence of the navigation methods typically tested in real experiments and how does it differ from the experiment performed in the thesis?