Evaluation of master's thesis:

Jakob Martinussen

Three-dimensional roof surface geometry inference from using remote sensing data

Conlusion

The thesis is graded A.

Justification

The introduction has a very good motivation, description of the problem setting, and explanation of the practical and societal significance of solving the problem. The research questions are made very clear based on an enlightening explanation of their context.

The chapter on «Data and preprocessing» provides a very good introduction to the application field and the terminology and knowledge required to read and understand the following chapters. The candidate here demonstrates strong domain knowledge and very good grasp of underlying theory from the application field, such as cartography, mapping and geographical data formats.

The chapter on «Modeling» gives a very good impression of the candidate's knowledge and theoretical understanding: He gives a very good description and delineation of different machine learning task and their relations in terms of complexity and order in the processing pipeline; He is mathematically rigorous in the definition of target properties for network training and prediction; The split of the dataset into training, validation and test sets is carefully handled to ensure valid validation and testing; The motivation of design choices for preprocessing and methodology is generally very thorough. However, the final choice of using the U-Net architecture over other alternatives could have been argued better for. Selection and validation of hyperparameters such as layer depth and dimensions should also have been discussed.

The formalism under which the pseudoinverse mapping is defined and analysed is ingenious and well-founded, and provides for a strong analysis of the problem, even though it difficult to assess how original the psedoinverse approach is, and how it relates to existing literature. A statement on its originality would been expected.

The «Postprocessing» chapter shares the good qualities of the previous chapters: The overall approach to problem-solving is extremely structured, systematic and thorough. The use of mathematical definitions is very rigorous. It complements the textural descriptions nicely and makes the exposition of the material very readable and clear. The reader is given good guidance by logical ordering of information and sequencing of sections and chapters that provides excellent readability.

The experiments and results displayed in the «Experiments» chapter are solid, thorough, well designed and well presented. The candidate provides very good pseudocode which outlines practical implementation of all methods. The effort he makes to explain all details and practical issues that are important for successful implementation are appreciated. He also gives due attention to time-efficient execution of implemented routines, for instance by using the r-tree index. In the analysis of the data, the candidate demonstrates excellent analytical skills. For instance, the analysis of the Lidar residual dataset displays a skillful combination of methodology and observation to interpret data.

The chapters describing theory and methodology do not follow the archtypical structure of a thesis, which is on the whole not a problem, as the manuscript has unusual pedagogical qualities and is written with great clarity. It provides the reader with good guidance in terms of good definitions of terminology and pointers to sections where key terms will be or have been defined in more detail. However, the state-of-the-art should have been more clearly separated as an own section. The survey of previous work is brief, which is perhaps because the topic has received little attention in the literature, but this should have been pointed out and clarified. The use of references is sufficient, but not excessive, and could be extended.

Moreover, it could have been pointed out more clearly what the contributions of the candidate are. The overall impression is that the extent of the thesis work is impressive, also noting that a lot of the data preparation and preprocessing is done by the candidate as part of his previous specialisation project. The extent of this underlying work is considerable, and an appropriate presentation is given, although they are not devoted a lot of space in the thesis. It is well documented how many design choices are based on conclusions from the specialisation project. However, the candidate would benefit from stating more boldly and clearly what his contributions are, so this would not have to be assumed.

No technical errors nor major flaws have been discover. However, the performance of the U-Net in roof segmentation and surface normal prediction is so good that it is an immediate question why deep learning has not also been applied to the instance segmentation problem, for which it should also be very suitable. This remark is made more pertinent by the observation that the chosen clustering approach underperforms as an approach to instance segmentation, as judged by the good starting point provided by the masked surface normal prediction. The motivation of using the DBSCAN and kNN algorithms for instance segmentation is sound and convincing. However, the results are underwhelming. Figures 5.4, 5.8 and 5.12 indicate that the surface normal prediction provides an excellent starting point for instance segmentation, which the chosen clustering algorithms have not been able to utilise. This is not commented on in the text.

The achievements of the current work have not been compared to nor discussed in light of the listed related work. This would have been appreciated. Validation of the model complexity in terms of number of layers and filters per layer is not mentioned. This is, along with the lack of motivation for choosing the U-Net architecture, a weakness.

On the whole, the thesis documents an impressively extensive work of high quality. The presented work is strongly driven by the desire to solve a very concrete problem. It is a strong engineering performance more than a scientific endeavour, although the approach to problem-solving is very, structured, systematic and well-founded in rigorous mathematics and solid understanding of the chosen methodologies. With the exception of the instance segmentation, the performance of the pipeline is very good. The multitask learning approach is the most innovative methodological choice in the thesis, which is otherwise based on standard neural network architectures and machine learning methods. The modified U-Net used for surface normal prediction is also noteworthy.

Stian Normann Anfinsen

Massimiliano Ruocco

Tromsø / Trondheim, 20 November 2020