Learning to look at LiDAR: combining CNN-based object detection and GIS for archaeological prospection in remotely-sensed data

Verschoof-van der Vaart, W.B.

Summary

The manual analysis of remotely-sensed data, i.e, information about the earth obtained by terrestrial, aerial, and spaceborne sensors, is a widespread practice in local and regional scale archaeological research, as well as heritage management. However, the amount of available high-quality, remotely-sensed data is continuously growing at a staggering rate, which creates new challenges to effectively and efficiently analyze these data and find and document the seemingly overwhelming number of potential archaeological objects. Therefore, computer-aided methods for the automated detection of archaeological objects are needed. In this thesis the development and application of automated detection methods, based on Deep Convolutional Neural Networks, for the detection of multiple classes of archaeological objects in LiDAR data is investigated. Furthermore, the implementation of these methods into archaeological practice and the opportunities of knowledge discovery—on both a quantitative and qualitative level—for landscape or spatial archaeology are explored.

The introduction, **Chapter 1**, gives an overview of archaeological spatial analysis and remote sensing, the Big Data problem, and automated approaches for large-scale archaeological prospection. Finally, the general concept of Deep (Region-based) Convolutional Neural Networks is introduced.

In **Chapter 2**, a multiclass archaeological object detection workflow, called WODAN1.0, is presented. The WODAN workflow consists of three steps: 1) A pre-processing step that converts LIDAR data into the required format; 2) an object detection step using an adjusted Faster R-CNN architecture; and 3) a post-processing step that transforms the results into geospatial data. Initial experiments show that WODAN1.0 is able to detect and categorize barrows and Celtic fields in LiDAR data from the *Veluwe* area in the Netherlands.

In the subsequent **Chapter 3**, the WODAN workflow is further improved by increasing the number of detected archaeological classes and by applying different modifications to the Faster R-CNN architecture and training scheme, resulting in the WODAN2.0 workflow. An additional post-processing step (Location-Based Ranking) is developed and incorporated to reduce False Positives caused by specific regions in the research area, based on archaeological domain knowledge. The practical value of WODAN2.0 is investigated by testing it on a test dataset, which better represents the real-world situation of scarce archaeological objects in different types of complex

terrain. Experiments show that WODAN2.0 is able to detect barrows, Celtic fields, and charcoal kilns and outperforms WODAN1.0. However, the workflow does not reach or exceeds general human performance, when compared to the results of a Citizen Science project conducted in the same research area. Following this research, additional changes were made to address specific issues, resulting in the WODAN2.5 workflow (**Chapter 3.8**).

In **Chapter 4** the application within archaeological practice and the transferability of WODAN2.5 is investigated by using it to detect barrows and Celtic fields in the Dutch *Midden Limburg* region, which differs in archaeology, terrain, and land-use from the *Veluwe* area in which WODAN was developed. It is shown that WODAN is able to detect (previously unknown) archaeological objects, provide information about the structuring of the landscape in the past, and can highlight biases in the current archaeological record.

In the second part of this thesis, CarcasonNet, an approach to automatically detect and trace hollow roads (i.e., more complex, large-scale landscape patterns) is presented (**Chapter 5**). The output of CarcassonNet consists of both geospatial polygons and lines, to efficiently study the roads themselves and the resulting route network. Experiments show that CarcassonNet is able to effectively detect hollow roads in LiDAR data from the Veluwe. Subsequently, in **Chapter 6** the transferability of CarcassonNet, trained on LiDAR data from the Netherlands, is investigated by testing it on data from Germany and Slovenia. The results show that CarcassonNet is able to detect roads in these regions, although the quality of the LiDAR data is of influence on the performance.

In the **last Chapter (7)** the possibilities and challenges of archaeological automated detection in remotely-sensed data and the progress made towards the aims of this research are discussed. Subsequently, the evaluation and transferability of the approaches is reviewed. It is argued that the incorporation of automated detection in archaeological practice is dependent on the role of these methods in the broader research framework. This thesis proposes Human—Computer strategies, in which automated detection precedes or is used in conjunction with manual analysis, to highlight areas of archaeological interest that require (field) verification. These strategies, with various levels of human involvement depending on the task, resolve many of the current issues of archaeological automated detection methods, while also offering insight in the biases in both manual analysis and automated detection.