Spatial data science: from aerial data to discrete spatio-temporal event analysis

Luca Colomba

Summary

The high availability of connected sensors and aerial monitoring systems enable the constant collection of geospatial data which, coupled together with data mining and machine learning techniques, enables a wide range of applications. Examples of these are traffic monitoring, energy production estimation, predictive maintenance and automatic fault detection systems, as well as Earth Observation applications. In the context of this thesis, we tackle different aspects of geospatial data, with the main focus being the remote sensing and multispectral domain. More specifically, we introduce two novel datasets in the context of emergency management, investigating the post-wildfire effects on flora and fauna. The two datasets, which are localized in Europe and California respectively, gather information from Sentinel missions and require the machine learning algorithm to predict whether an area (i.e., a pixel) has been previously affected by a forest fire. As such, the proposed task for both datasets is an image segmentation task. The European dataset consists of a multiclass segmentation problem, with five different severity levels as published by the Copernicus Emergency Management Service, whereas the Californian one introduces the binary problem. The two datasets were collected starting from information publicly released by public entities, such as the aforementioned Copernicus EMS and the California Department of Forest Fire and Protection.

In this context, we leveraged the introduced datasets to validate existing state-of-the-art architectures on the multispectral domain, demonstrating their adaptability to the remote sensing field. Furthermore, we introduce two novel frameworks, namely the Double-Step and Magnifier frameworks, which achieved better performances compared to baseline models. More specifically, we evaluated UNet and SegFormer models against the proposed methodologies, investigating the effects of the newly introduced components

in the two frameworks. Additionally, we evaluated the importance of different channel groups in the context of automatic burned area detection by means of occlusion-based strategies and the attention mechanism, comparing the results against domain expertise evaluated through index-based analyses. Finally, we assessed the performance of a novel type of architecture based on Graph Neural Networks on the multispectral vision domain in the context of land cover classification. Such existing architecture, namely ViG, demonstrated superior performance in both mid-sized and large-scale datasets, surpassing both transformer-based and CNN-based models.

To conclude, this thesis introduces a novel type of spatio-temporal pattern, namely the spatio-temporal invariant pattern, with its associated distributed mining algorithm. Such data mining algorithm analyzes a dataset of spatio-temporal events, such as critical events in a bike sharing system or traffic events, and investigates a restricted set of events of interest. Given such set of events, the spatial and temporal neighborhood of each relevant event is analyzed and a spatio-temporal projection step is performed to obtain a sequence of events. Given the resulting sequential database, we extract frequent sequences of events. We compared the results obtained by the proposed methodology against tree-based patterns known in literature, demonstrating a higher informative content carried by the spatio-temporal invariant pattern. Moreover, scalability tests were performed to evaluate the complexity of the methodology with respect to the number of input events.