eCognition - Classification and Object Features

Pedro Afonso
PARIS LODRON UNIVERSITY OF SALZBURG

December 21, 2025

All the roads lead to... OBIA.
-This semester's motto

Abstract

eCognition has a plethora of tools aimed at Object-Based Image Analysis (OBIA). In this report, I aimed to explore some of the segmentation and classification tools.

I began by comparing the Random Forest Classification and the Support Vector Machine Classification of a photo of the banks of the Salzach river taken by Planet. Afterwards, I explored several Segmentation Objects Features in an old image of the south of Salzburg, most notably sub-class creation and super-objects creation.

Keywords: eCognition, Segmentation, Classification, OBIA

1 Introduction

At this point, I am all too familiarized with the concept of image segmentation and classification. There are several ways to execute the segmentation (from the simple chessboard segmentation to more complex ones like the multi-resolution segmentation). And, likewise, the classification can be accomplished via a purely mathematical approach (as it is often the case with the over usage of NDVI), or via machine learning methods.

For this assignment, I began by exploring different machine learning methods available in eCognition for image classification, and, on a second stage, I turned my focus to the segmentation part to explore the tools provided by eCognition to analyze the results.

2 Machine Learning Classification

The first step was to segment the image with the "multiresolution segmentation" tool. I tweaked the default settings a bit and, after running it a few times, I settled on a "scale parameter" of 50, a "shape" weight of 0,3 and a "compactness" weight of 0,5, while leaving all the layer weights at 1.

Afterwards, it was time to create the possible classes which would be assigned to my segments. I simply followed the instructions and created 5 classes (as can be seen in Figure 1): Arable land, Artificial surfaces, Forest, Pastures, and Water.

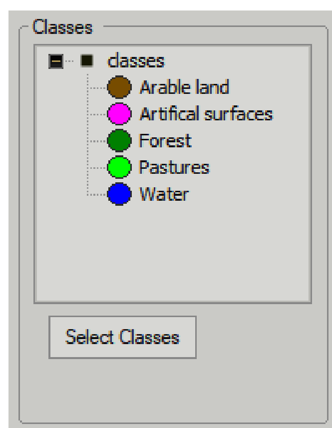

Figure 1: The five classes I created and respective assigned colors.

And, then, I assigned each class to a few segments in order to use them as samples during my model training. Figure 2 shows the several samples painted with the respective class color overlayed on top of the image.

After this, I used the "Feature Space Optimization" tool to assess which would be the best features to include in the training of my model. I picked the area of the segments, the overall brightness, the mean value and

Figure 2: Picked samples and overall photo. The Salzach river, in the middle, clearly separates and greener area (on the right) from a more arid one (on the left).

standard deviation of each layer, and the gray-level co-occurrence matrix (GLCM), which acted as a proxy for the texture (the result can be seen in Figure 8, in the Appendix, where it is clear the best number of features is around 7).

Despite the results, I was quite liberal with my choice of features afterwards. First of all, even though the "area" was included as one of the decisive choices for both 7 and 6 features, I decided I would not include it because it seemed too random and I did not want to overfit the model; it would probably only make sense if I had a larger scale parameter. Then, I decided that I couldn't simply ignore the mean value of Layer 1 (even if is, in a way, already represented in the brightness). With this in mind, I decided to simply swap the area for Layer 1's mean value. I know that, theoretically, this means that this new set of features yields a lower separation distance in my samples, but I am still confident that it is a wise choice, backed by a solid reasoning.

I then trained a model using a Random Forest Classification and another using Support Vector Machine Classification (the process tree can seen on Figure 9, in the Appendix). The results can be seen in Figures 3 and 4, respectively.

Despite using the exact same samples, the Support Vector Machine Classification resulted in less artifact generation. With the Random Forest Classification, there are several segments in the edges of fields that were mistakenly identified as artificial surfaces and, to a lesser degree, forests seem to be mistaken with these edges too. Furthermore, the Random Forest Classification appears to misidentify water much more often.

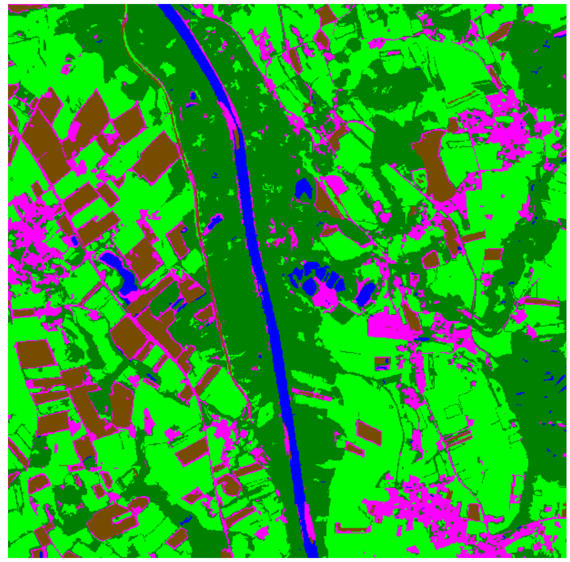

Figure 3: Classification using a Random Forest method.

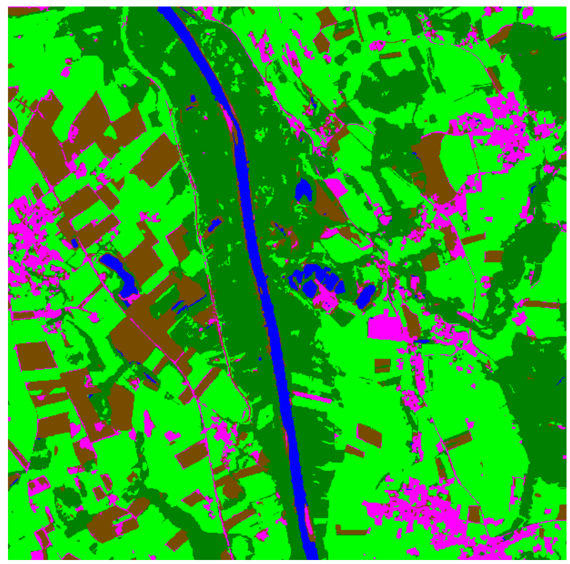

Figure 4: Classification using a Support Vector Machine method.

3 Object Features

This time, I started with the basic chessboard segmentation, dividing the image in small squares. Then, I began my exploration of Object Features, including the mean value of every layer, the maximum value of the blue layer (Layer 1), some geometry-related features and NDVI (manually created by me).

I then compared a water segment with a boat segment (both signaled in Figure 5).

Figure 5: Detailed view of the chessboard segmentation, in black, and selection of two segments, in red.

As shown in Figures 10 and 11 of the Appendix, the NDVI is unsurprisingly negative for the water segment, while being slightly positive on the boat; every layer has higher values on the boat, which is why it looks much brighter and whiter; but the most obvious thing is that both segments have an area of 100 pixels and a shape index of 1, because every segment in the chessboard segmentation is the same, making the geometric statistics quite redundant.

This is no longer true for a multi-resolution segmentation, where the water segment is two orders of magnitude larger than the boat segment (as seen in Figures 12 and 13 of the Appendix).

Sticking with the multi-resolution segmentation, I then used the NDVI feature I had created to classify my segments (using the "assign class" algorithm):

- as vegetation (green), if $NDVI \geq 0.25$
- as water (blue), if $NDVI < -0.15$
- as a "marine vessel" (red), if it is unclassified and completely surrounded by water segments.

The result follows in Figure 6, where three segments were identified as water and 57117,6 m^2 of land was identified as vegetation (as seen in Figure 14).

To make this more interesting and use my Air Quality Layer (which had gone forgotten until now), I divided the vegetation class in two subclasses: High Air Quality Vegetation (light green) and Low Air Quality Vegetation (dark green). The final composite is illustrated in Figure 7.

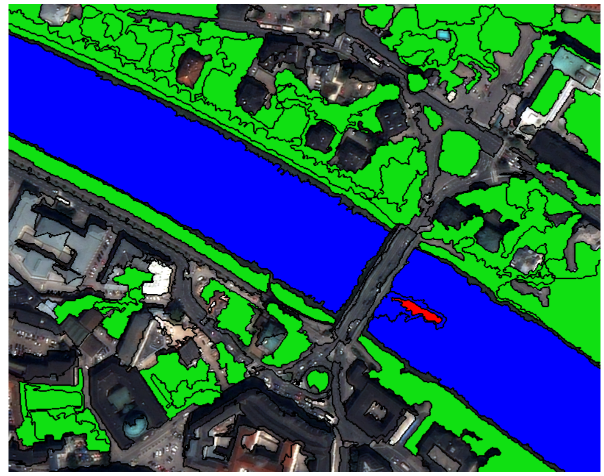

Figure 6: Multiresolution segmentation with segments colored according to their class. Unclassified segments were left unfilled.

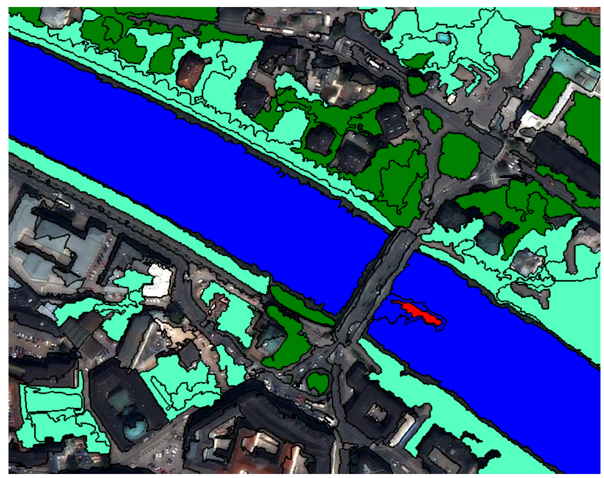

Figure 7: Multiresolution segmentation with segments colored according to their class and subclass. Unclassified segments were, once more, left unfilled.

The final Object-Feature Tool I explored was the creation of super-objects. I created a new segmentation level below my previous one, with a scale parameter of 50, this time. Then, I created super-objects for the vegetation class using a hierarchical distance of 1, of course, since I only had one other level. The result is, to no one's surprise a copycat of the previous Vegetation layer (as shown in Figure 15a, on the appendix), except each of the previous segments has been divided into multiple smaller ones.

If a hierarchical distance of 0 or any quantity larger than 1 was chosen, the result would be null, given there are no other levels with classification layers. I tried it simply to assess the result, and I present the outcome on Figure 15b, in the appendix.

4 Appendix

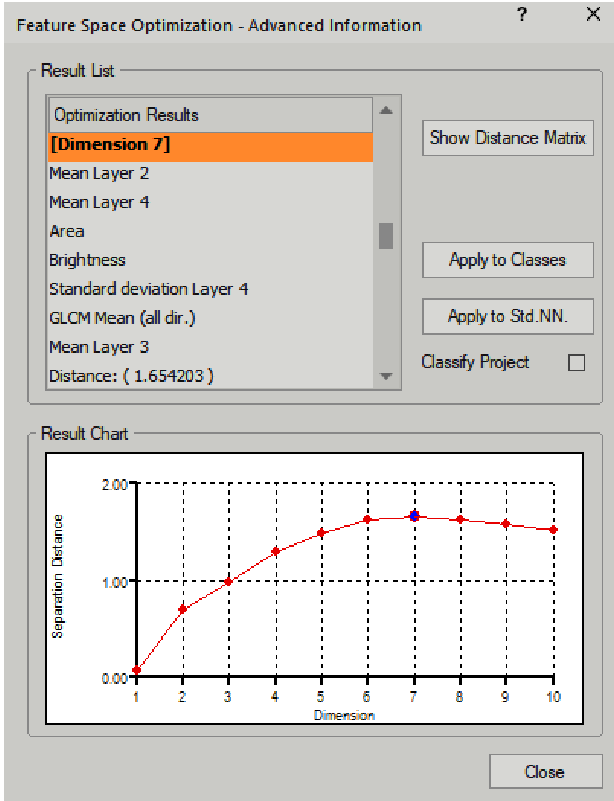

Figure 8: Trend of separation distance in respect to the best combination of different numbers of features (in the bottom); and best combination of seven features (at the top).

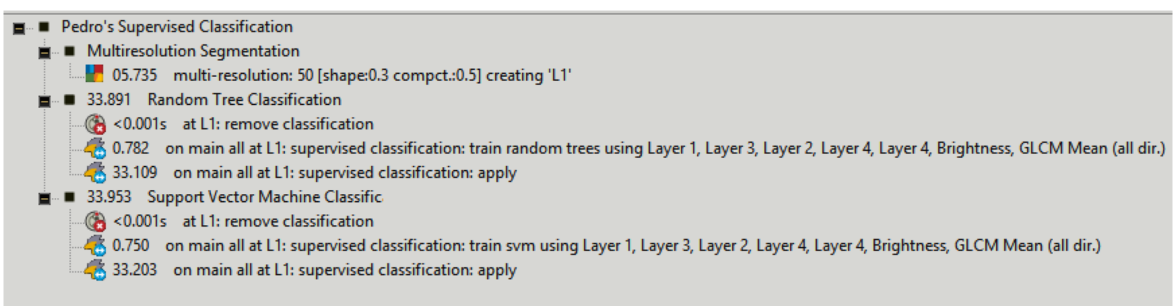

Figure 9: Full Process Tree of the First Part of the Report. It includes a "remove classification" step at the beginning of each training step to start from a blank state.

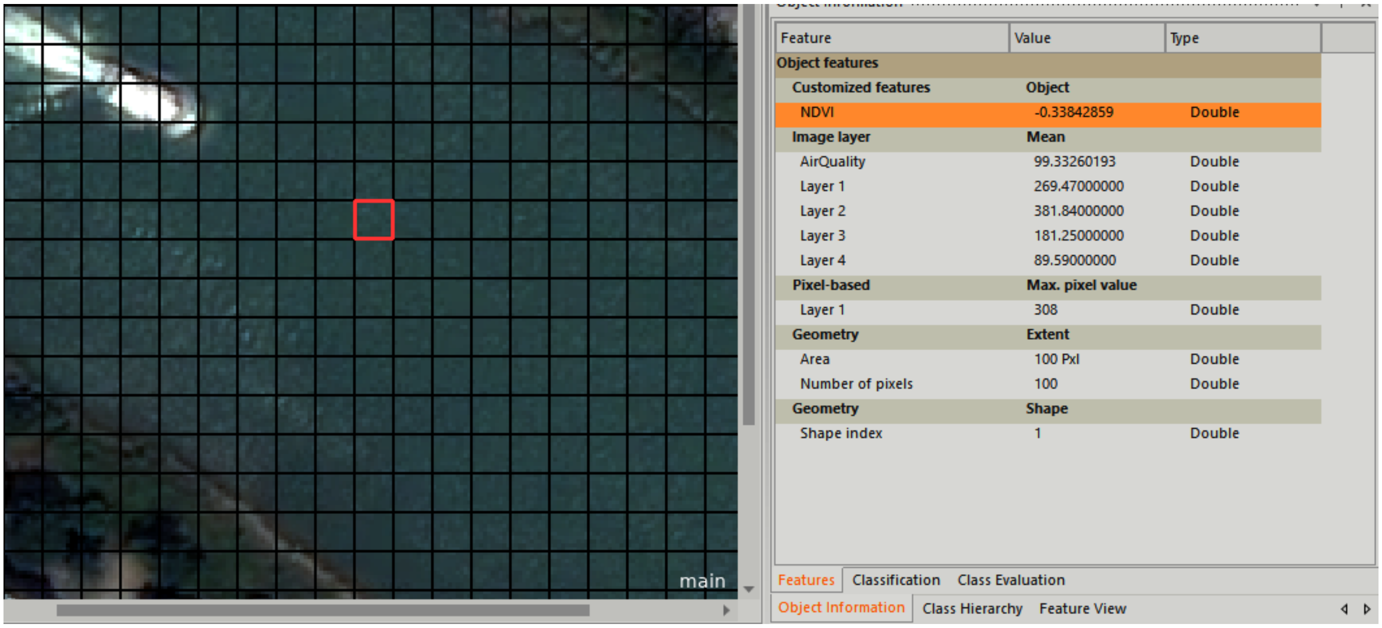

Figure 10: Object Features of the water segment created by the chessboard segmentation

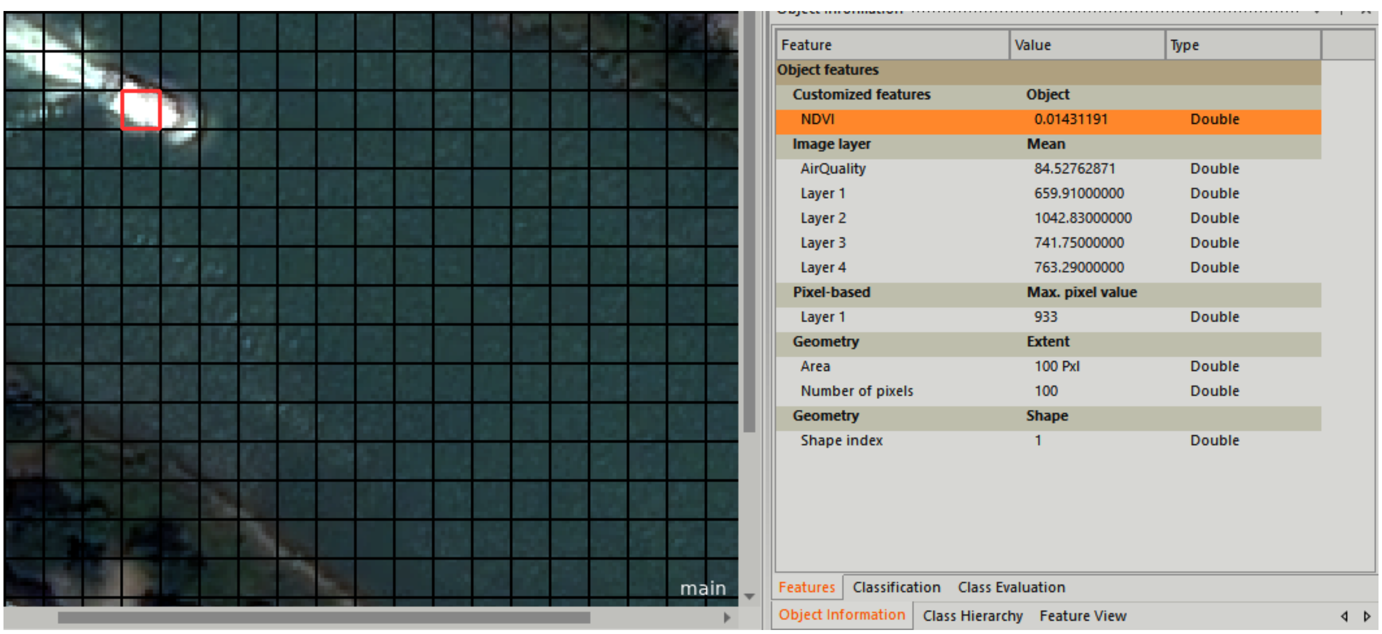

Figure 11: Object Features of the boat segment created by the chessboard segmentation

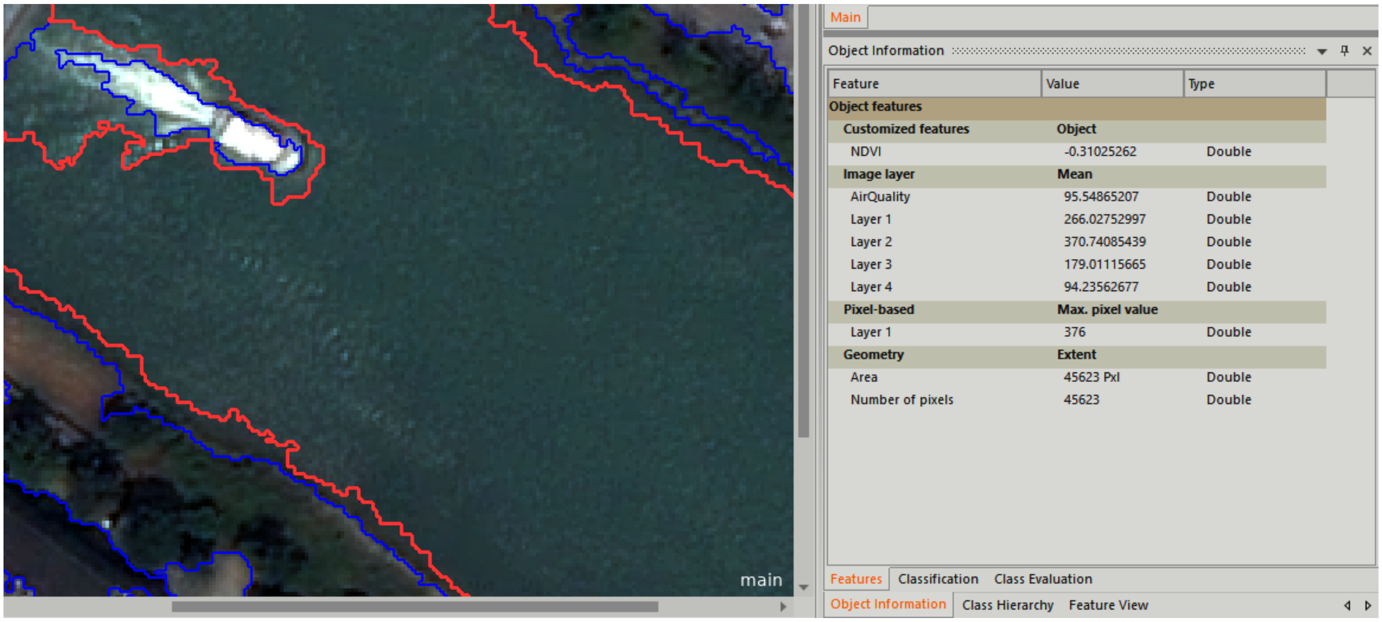

Figure 12: Object Features of the water segment created by the multi-resolution segmentation

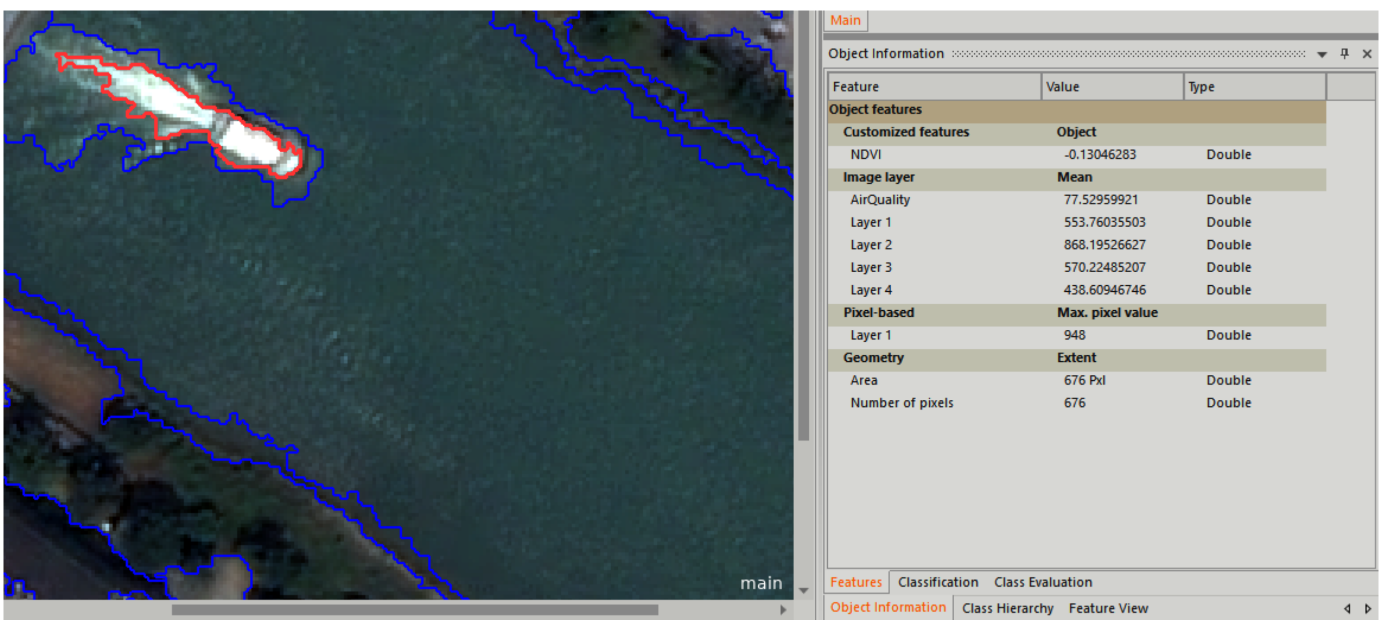

Figure 13: Object Features of the boat segment created by the multi-resolution segmentation

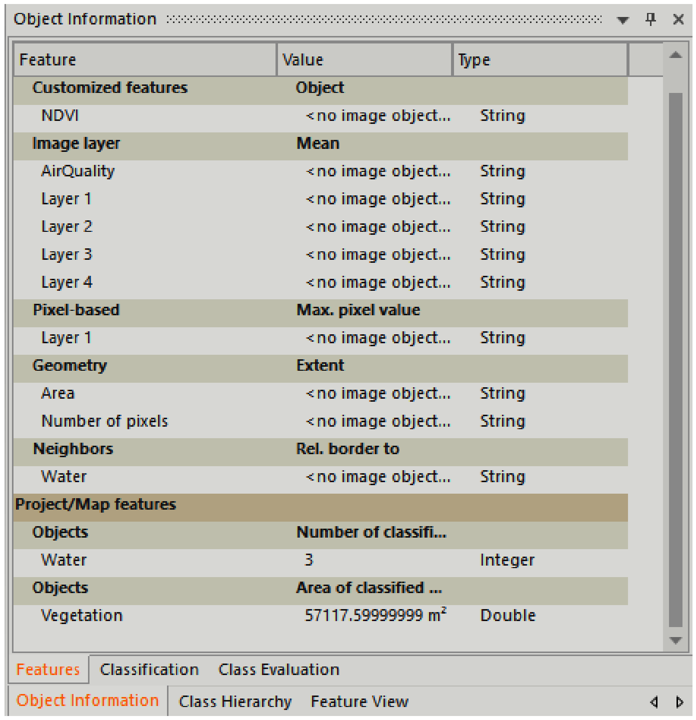

Figure 14: Selected features of the classified image.

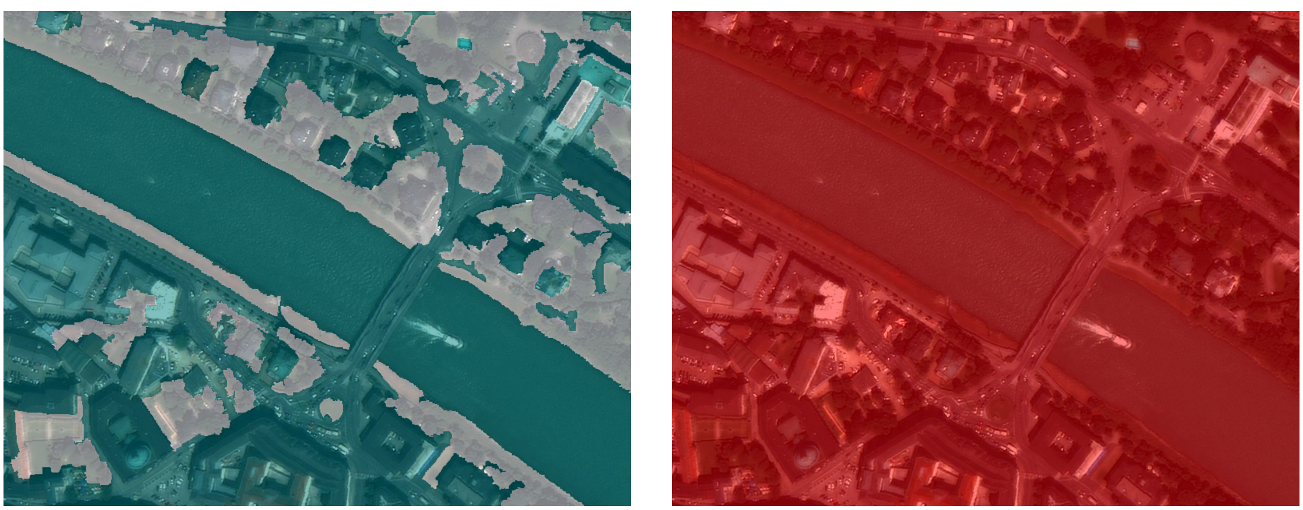

(a) Distance equal to 1.

(b) Distance different than 1.

Figure 15: Segments' Classification as a result of the super-objects creation. As counter-intuitive as it may be, the vegetation is colored grey, while non-vegetation colored in what appears to be teal. I didn't pick the colors this time!