International Workshop on Automatic Vectorisation of Historical Maps - 13 March 2020 - ELTE, Budapest
https://doi.org/10.21862/avhm2020.03

Chenjing Jiao[1], Magnus Heitzler[1], Lorenz Hurni[1]

Extracting Wetlands from Swiss Historical Maps with Convolutional Neural Networks

Keywords: Siegfried map, map feature extraction, Fully Convolutional Neural Networks, wetland reconstruction

Summary: Historical maps can serve as valuable resources for various kinds of researches such as ecology, land reclamation (Ngo et al, 2015), toponymy, history, etc. (Chiang, 2016). Specifically, extracting wetlands from historical maps facilitates researchers to investigate the spatio-temporal dynamics of the hydrological and ecological situation. Deep learning methods, especially Fully Convolutional Neural Networks (FCNN), provide an efficient and effective way to extract features from raster maps. To extract wetland areas from the Swiss Siegfried maps, we trained a U-Net based architecture and applied the learnt model to 573 map sheets across Switzerland. The pixel-wise prediction results were converted to polygons through a vectorization and generalization step. The vector wetland layers will be used for studies on land-cover change.

Introduction

The accessibility of the large numbers of historical maps indicates the high demand for successful feature extraction methods (Leyk & Boesch, 2009), because a wealth of cartographic information is still locked in the existing historical map series, such as building footprints, hydrography, map labels, etc. It is difficult to assess and compare wetland's historical and the current coverage and condition. Thus, historical maps are valuable sources for wetland change analysis as they retain the most reliable spatial data. Wetland areas in historical maps are often shown as composite elements that consist of clustered small symbols whose spatial distribution and patterns can define higher-level spatial objects, such as forests, grasslands, etc. (Leyk & Boesch, 2009). In *Figure 1*, the wetlands are denoted with composite horizontal strokes.

Conventional extraction approaches such as color image segmentation (CIS) have demonstrated their effectiveness in extracting map features, like roads, buildings, text, etc. Nevertheless, the conventional approaches are not perfectly suitable for historical maps, because the maps often suffer from poor graphical quality due to bleaching of the original paper maps and archiving practices (Chiang, et al., 2014). Based on the color feature of image pixels, CIS assumes that homogeneous colors in the image correspond to the same classes. The homogeneous regions that can be classified by CIS must show spatial contiguity. Thus, CIS may fail in extracting a composite object as a whole because of the trivial symbols and their complicated cluster. As we can see from *Figure 1*, one single wetland area is clustered by some trivial strokes, but the symbols are not connected with each other, although they form a contiguous texture.

[1] Institute of Cartography and Geoinformation, ETH Zürich [cjiao@ethz.ch]

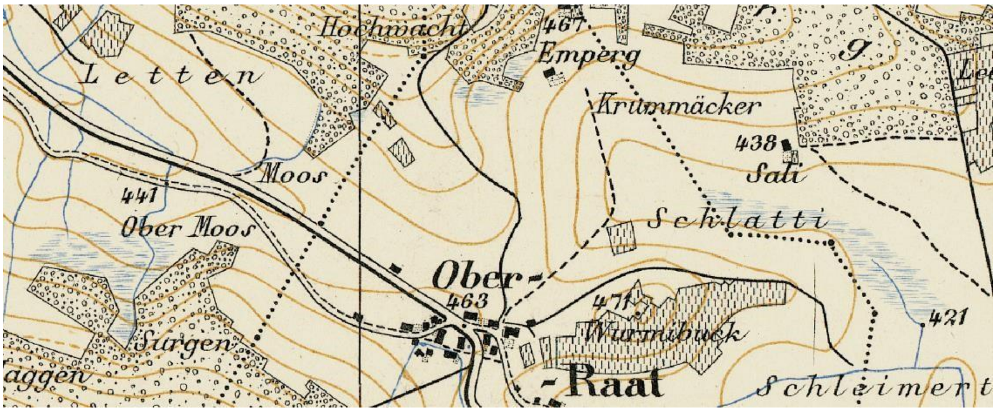

Figure 1. Wetland areas shown as composite graphic elements consisting of clustered horizontal strokes. Geodata © Swisstopo

In recent years, a group of supervised machine learning algorithms, known as Fully Convolutional Neural Networks (FCNN), has been exerting their superiority in automated feature extraction from images, especially historical maps (Saeedimoghaddam and Stepinski, 2019). An FCNN architecture, consisting of several convolutional, ReLU, and pooling layers, gradually extracts complex features from an input image, and then predicts the possibility of a feature correlating to a specific class. Yet, one of the challenges of FCNNs is that they inevitably require a large amount of training data to perform sufficiently well. However, producing training data manually might not be an option due to a lack of resources, which motivates the search for alternative sources of training data.

In this paper, we used the data from another study on wetlands in Switzerland as training data and trained an FCNN model to detect wetland pixels, which is followed by a vectorization and generalization step that generates wetland geometries.

Data and Methods

The Siegfried map is a comprehensive Swiss national map series that was published between 1872 and 1949. In a recent real-world use case, wetlands for whole Switzerland for the year 1880 are to be extracted for a Non-governmental Organization (NGO) from the Siegfried Map Series to carry out a study on land-cover change. For training an FCNN model, we need the input map sheet and the corresponding labeled wetland data. However, due to the unavailability of the labeled data for the year 1880, data of the year 1900 from the study on the historical development of wetlands in Switzerland has been used for training (Stuber & Bürgi, 2018). One Siegfried map sheet is shown in *Figure 2(a),* with its corresponding labeled training data presented in *Figure 2(b)*, whose pixel holds a one if it belongs to a wetland, and zero otherwise. Naturally, these two datasets for 1880 and 1900 come from somewhat different distributions. The symbolization may have changed slightly between the two periods, and the map sheets of the training dataset have been georeferenced in a different way from the sheets of 1880.

https://doi.org/10.21862/avhm2020.03

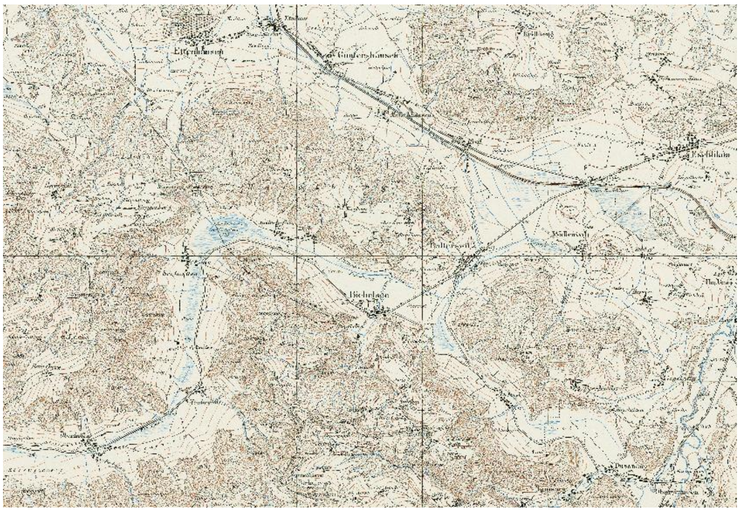

(a)

(b)

Figure 2. Data example: (a) One map sheet; (b) the labeled wetland data corresponding to (a). Geodata © Swisstopo

We developed a segmentation model that is based on the U-Net architecture (Ronneberger & Brox, 2015), which was adapted to the problem at hand. According to Heitzler and Hurni (2020), the U-Net is comprised of two diametrical network paths, the contracting path and the upsampling path, which is presented in Figure 3. The contracting path consisting of convolution layers and max-pooling layers generates increasingly higher abstract representations of the input image. The max-pooling layers used in conjunction with the convolution layers allow for reducing the extent of the input matrix. The up-sampling path infers the class of a pixel based on its position and its surroundings by gradually combining these features on different granularities. Deconvolution layers are used in the up-sampling path to quadruple the extents of input layers. Each layer is followed by a corresponding elu activation layer except for the max-pooling layers and the final convolution layer, which has a sigmoid activation layer. These activation layers delimit the output probabilities to the range [0, 1]. In the used implementation, the convolution layers have 32 filter kernels.

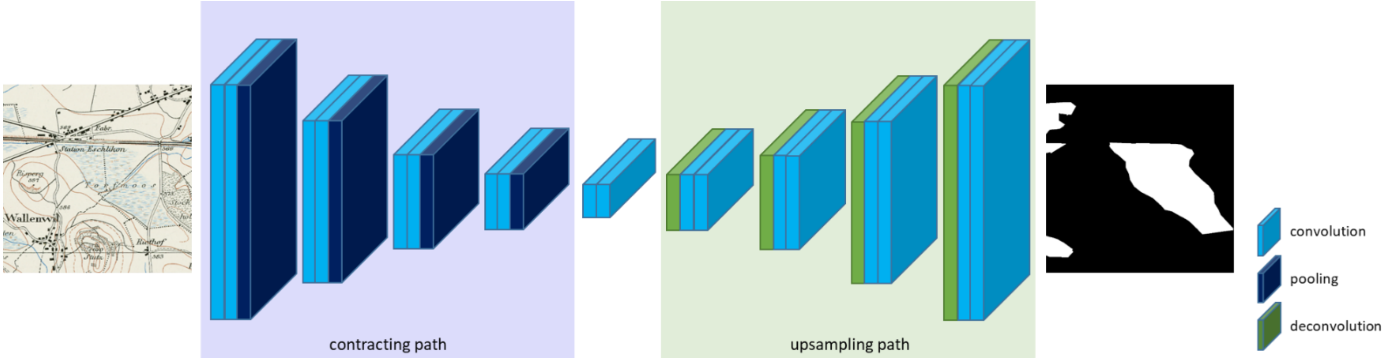

Figure 3. The FCNN architecture that is adapted based on the U-Net by Ronneberger et al. (2015) Geodata © Swisstopo

Each map sheet sized 7,000×4,800 pixels is subdivided into a grid of tiles with the size of 200x200 pixels to be processable by the U-Net model. Furthermore, to avoid misclassifying the border pixels, the extent of the input tile is expanded by 60 pixels on each side. A well-learnt model is obtained after training with 36 map sheets, namely 36×35×24 tiles. When a map sheet is fed into the trained model, it generates a pixel wise prediction result, holding the probability of each pixel belonging to a wetland. The pixels with a possibility greater than 0.5 are segmented as wetland areas.

Once the wetland segmentation results are generated, vectorization is performed using functions from the Geospatial Data Abstraction Library (GDAL) to convert the raster binary results into vector geometries. Next, generalization is implemented using the Douglas–Peucker algorithm, which iteratively traces the line segments of the original polyline and only keeps the start point, end point, and the farthest point of the current line segment (Douglas & Peucker, 1973). If required, the vectorized and generalized geometries can still be manually corrected to yield high-quality wetland vector layers.

Result

An example for a wetland reconstruction result is shown in Figure 4, where Figure 4(a) presents the map sheet. Fig 4(b) illustrates the pixel-wise prediction result of a subsection of the area, in which the luminance of the red color corresponds to the possibility that the pixel belongs to a wetland. In other words, the brighter red pixels indicate a higher probability of them belonging to a wetland. The final result of this subsection is presented in Figure 4(c), which overlays the wetland geometries on the map sheet. We can see that the wetland boundaries are accurately extracted. Besides, a wetland area denoted with composite elements can be extracted as a whole, even though it is intersected with contour lines or other composite elements. The results verify the effectiveness as well as the robustness of the reconstruction approach. Wetland areas have been extracted and vectorized from 573 Siegfried map sheets across Switzerland with this method.

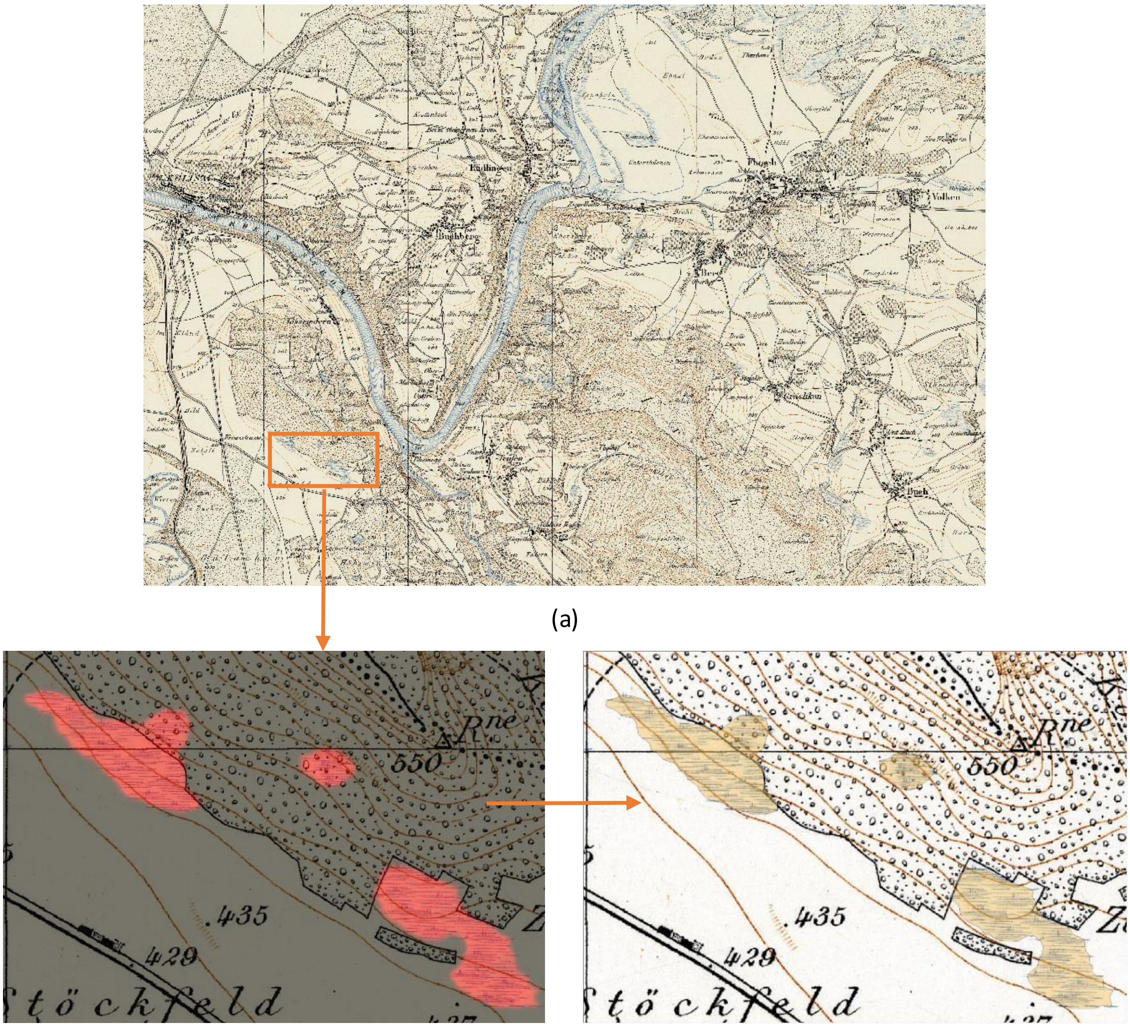

(a)

(b)

(c)

Figure 4. The wetland extraction result: (a) the map sheet; (b) the pixel-wise prediction result of a fractional area; (c) the vector geometry after vectorization and generalization. Geodata © Swisstopo

Conclusion

This paper addresses the wetland reconstruction issue by developing and training a U-Net based architecture, and applying the well-learnt model to the target map sheets. It is followed by a vectorization and generalization step, which outputs vector wetland layers that can be analyzed by Geographic Information Systems (GIS). The wetland reconstruction method performs well for recognizing and extracting the composite wetland elements. Future work will not only focus on the wetland extraction, but also on other hydrological features on the historical maps.

References

Chiang, Y., Leyk, S., Knoblock, C.A., 2014. A survey of digital map processing techniques. *ACM Computing Surveys*, *47*, 1:1-1:44. https://doi.org/10.1145/2557423

Chiang. Y., 2016. Unlocking textual content from historical maps - potentials and applications, trends, and outlooks. *International Conference on Recent Trends in Image Processing and Pattern Recognition*, 111-124, Bidar, India. https://doi.org/10.1007/978-981-10-4859-3_11

Douglas, D., Peucker, T., 1973. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *The Canadian Cartographer* 10: 112-122, https://doi.org/10.3138/FM57-6770-U75U-7727

Heitzler, M., Hurni, L., 2020. Cartographic reconstruction of building footprints from historical maps: a study on the Swiss Siegfried map. *Transactions in GIS*. https://doi.org/10.1111/tgis.12610

Ngo, V., Swift, J., Chiang, Y., 2015. Visualizing land reclamation in Hong Kong: a web application. In: *International Cartographic Conference*, Rio de Janeiro, Brazil.

Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computerassisted intervention*, 234-241, Munich, Germany. https://doi.org/10.1007/978-3-319-24574-4_28

Saeedimoghaddam, M., Stepinski, T.F., 2019. Automatic extraction of road intersection points from USGS historical map series using deep convolutional neural networks. *International Journal of Geographical Information Science*, 1-22. https://doi.org/10.1080/13658816.2019.1696968

Leyk, S., Boesch R., 2009. Extracting composite cartographic area features in low-quality maps, *Cartography and Geographic Information Science*, 36(1): 71-79. https://doi.org/10.1559/152304009787340115

Stuber, M., Bürgi, M., 2018. Vom "eroberten Land" zum Renaturierungsprojekt. Geschichte der Feuchtgebiete in der Schweiz seit 1700. Bristol-Schriftenreihe: Vol. 59. Bern: Haupt Verlag.