International Cartographic Association
Association Cartographique Internationale

Deep Generation of Synthetic Training Data for the Automated Extraction of Semantic Knowledge from Historical Maps

Lukas Arzoumanidis [*], James Ormond Fethers [a], Sethmiya Herath Mudiyanselage [a], Youness Dehbi [a]

[a] *HafenCity University, Computational Methods Lab, Hamburg, Germany, {lukas.arzoumanidis, youness.dehbi}@hcu-hamburg.de*

* Corresponding author

Keywords: historical maps, generative adversarial networks, training data generation, cartographic heritage

Abstract:

As a prerequisite for backtracking and monitoring environmental changes, historical maps are invaluable as they precisely capture the spatial reality of the past (Chiang et al., 2019). To avoid further deterioration of useful information that could unlock the possibility for spatio-temporal analysis, often a *digitization* is performed, which results in a large number of scanned historical documents stored in digital archives (Wu et al., 2023). However, in this case scanned historical documents, e.g., maps or plans, are not semantically enriched which makes the underlying spatial information unreachable (Chiang et al., 2019). In this context, it is desired to digitize heterogeneous historical map corpora in a way that their semantic knowledge can be automatically extracted by machines, recently demonstrated through semantic segmentation using different Deep Learning techniques (Arzoumanidis et al., 2023, Wu et al., 2023, Petitpierre et al., 2021). After analyzing, vectorization allows to store extracted semantic knowledge in a machine readable format which sets the basis for an effective and intuitive comparison with current map features, e.g., in a spatio-temporal analysis (Wu et al., 2023).

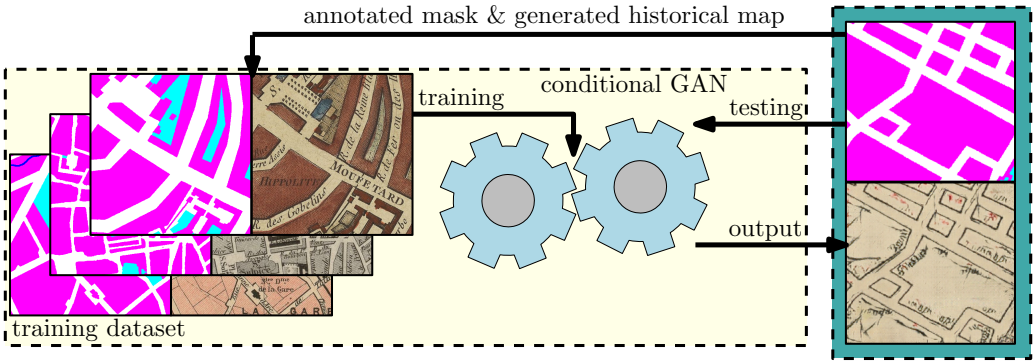

Figure 1. Our approach for synthetically generating training data.

Most of the recently applied machine learning methods rely, however, on pre-labelled training data. In particular, Deep Learning based approaches need an extensive amount of annotated maps for learning. As yet, such data has been mostly prepared with a high overhead manually. This paper paves the way to provide large datasets for the automatic generation of synthetic historical maps in order to overcome the time consuming and tedious labelling process. In this context, we proceed two-folds: Based on an heterogeneous corpus of historical maps from Paris, (1) on the one hand, we learned the style, e.g., coloring or deterioration, of the underlying maps and transfer it to a set of current and thematically different maps as can be shown in Figure 2. Herewith, we merely capture the style of the source maps from the corpus allowing for the generation and sampling of arbitrary and huge amount of new historically-styled maps. On the other hand (2), based on an annotated corpus consisting of image pairs, we were able to learn a mapping between the underlying semantically segmented objects in historical maps, e.g., streets and buildings, and the implicit style of the map accordingly. Figure 3 shows one example of image pairs (a-b) which have been used as training data. The generated historical maps are depicted on the most right. Such a result, allows for the automatic generation of synthetic historical maps based on Volunteer Geographic Information (VGI), e.g., OpenStreetMaps (OSM) or even make use of on the-fly sketched annotation masks. In a recent work, Li (2019) applied also a conditional GAN to generate historical maps from existing geographic datasets like OSM. He focused, however, on generating data for the training of text detection systems from historical maps. With our approach, we gave more attention to pave the way for generating a large amount of data, consisting of mask and image pairs, for the sake of semantic segmentation.

For both approaches, a conditional GAN implementing a generator-discriminator architecture has been applied. For the purpose of style transfer, we used *CycleGAN* which performs an unpaired image-to-image translation. Further, we applied

https://doi.org/10.5194/ica-abs-7-7-2024

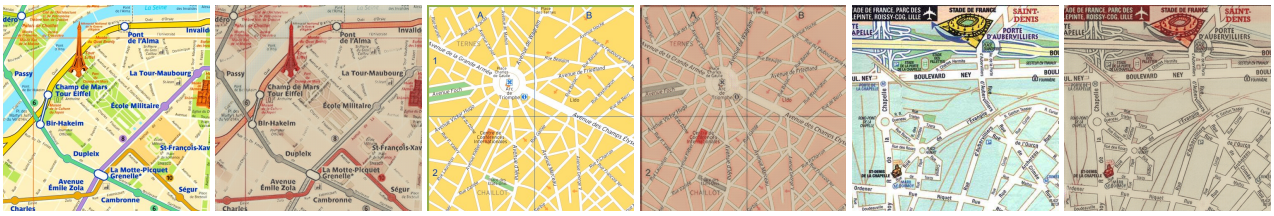

(a) present map (b) style transfer (c) present map (d) style transfer (e) present map (f) style transfer

Figure 2. Historical maps generated from current maps (`https://ontheworldmap.com/france/city/paris/`).

pix2pix for the synthetic training data generation using a paired image-to-image translation (Zhu et al., 2017, Isola et al., 2017). For both experiments, we choose a learning rate of 0.0002 and a batch size of 1. The style transfer learner has been trained for 25 epochs using 940 of historical and present maps, whereas the synthetic training data generator has been trained for 200 epochs using 940 pairs of historical maps and their corresponding annotation masks from the mentioned Paris dataset.

Experimental results demonstrate the ability to transfer historical map style into current and thematically different maps of Paris using *CycleGAN* (cf. Figure 2). While the topological and geometric structure of the underlying maps have been preserved, the networks is able to capture both coloring information and the level of visual deterioration from the learning examples. Beyond this style information, the deep synthetic data generator is able to mimic the geometry of streets, streams, building footprints from the annotated masks to the according generated historical maps. Even text labels have been placed in expected regions, they are however meaningless. It can be also stated that the semantically segmented annotated masks influences both the appearance, e.g., hatching, and the color of different objects such as buildings and streets (cf. Figure 3). As mentioned, the resulting maps could serve as learning examples for the automatic digitization of historical maps. In this context, the evaluation of our system will be subject of an upcoming research article.

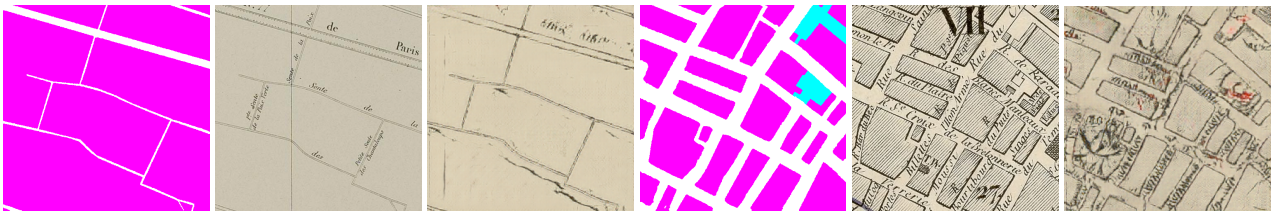

(a) annotated mask (b) ground truth (c) generated map (d) annotated Mask (e) ground truth (f) generated map

Figure 3. Historical maps (c,f) generated from pairs of annotated masks (a,d) and historical maps (b,e) respectively.

Acknowledgements

We gratefully acknowledge the open dataset of annotated historical maps from Paris provided by Petitpierre (2021).

References

Arzoumanidis, L., Knechtel, J., Haunert, J.-H. and Dehbi, Y., 2023. Self-constructing graph convolutional networks for semantic segmentation of historical maps. *Abstracts of the ICA. 31st International Cartographic Conference* 6, pp. 11.

Chiang, Y.-Y., Duan, W., Leyk, S., Uhl, J. H. and Knoblock, C. A., 2019. *Using Historical Maps in Scientific Studies: Applications, Challenges, and Best Practices*. Springer Publishing Company, Incorporated.

Isola, P., Zhu, J.-Y., Zhou, T. and Efros, A. A., 2017. Image-to-image translation with conditional adversarial networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5967–5976.

Li, Z., 2019. Generating historical maps from online maps. In: *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, SIGSPATIAL '19, Association for Computing Machinery, New York, NY, USA, p. 610–611.

Petitpierre, R., 2021. Historical city maps semantic segmentation dataset.

Petitpierre, R., Kaplan, F. and di Lenardo, I., 2021. Generic semantic segmentation of historical maps. *Computational Humanities Research Conference* pp. 228–248.

Wu, S., Chen, Y., Schindler, K. and Hurni, L., 2023. Cross-attention spatio-temporal context transformer for semantic segmentation of historical maps. In: *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*, SIGSPATIAL '23, Association for Computing Machinery, New York, NY, USA.

Zhu, J.-Y., Park, T., Isola, P. and Efros, A. A., 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2242–2251.

https://doi.org/10.5194/ica-abs-7-7-2024