

BUILDING GENERALIZATION USING DEEP LEARNING

M. Sester*, Y. Feng, F. Thiemann

Institute of Cartography and Geoinformatics (ikg), Leibniz Universität Hannover, Germany
(monika.sester, yu.feng, frank.thiemann)@ikg.uni-hannover.de

Commission IV, WG IV/4

KEY WORDS: Generalization, Cartography, Deep Learning

ABSTRACT:

Cartographic generalization is a problem, which poses interesting challenges to automation. Whereas plenty of algorithms have been developed for the different sub-problems of generalization (e.g. simplification, displacement, aggregation), there are still cases, which are not generalized adequately or in a satisfactory way. The main problem is the interplay between different operators. In those cases the benchmark is the human operator, who is able to design an aesthetic and correct representation of the physical reality.

Deep Learning methods have shown tremendous success for interpretation problems for which algorithmic methods have deficits. A prominent example is the classification and interpretation of images, where deep learning approaches outperform the traditional computer vision methods. In both domains - computer vision and cartography - humans are able to produce a solution; a prerequisite for this is, that there is the possibility to generate many training examples for the different cases. Thus, the idea in this paper is to employ Deep Learning for cartographic generalizations tasks, especially for the task of building generalization. An advantage of this task is the fact that many training data sets are available from given map series. The approach is a first attempt using an existing network.

In the paper, the details of the implementation will be reported, together with an in depth analysis of the results. An outlook on future work will be given.

1. INTRODUCTION

Cartographic generalization is the process of generating smaller scale representations from large scale spatial data. This process has been conducted by cartographers ever since; they are applying different operators, such as selection, simplification or displacement, which - in sum - lead to a simpler and a more clear representation of the spatial scene at a smaller scale. There are many interesting solutions to solve this problem and provide automatic processes. The main challenge today lies in the problem of the interplay between different operators. Here, the benchmark is still the human operator, who is able to design an aesthetic and correct representation of the physical reality.

Deep Learning methods have shown impressive success for interpretation problems for which algorithmic methods are difficult. A prominent example is the classification and interpretation of images, where Deep Learning approaches outperform the traditional computer vision methods. In both domains - computer vision and cartography - humans are able to produce a solution; this also implies, that there is the possibility to generate positive examples for the different cases. Thus, the idea in this paper is to employ deep learning for cartographic generalizations tasks, especially for the task of building generalization. Most Deep Learning approaches are based on supervised learning, i.e. they require example data with given input-output pairs. Fortunately, in generalization many training data sets are available from the existing map series.

To automate the generalization process, various methods are being applied, namely optimization approaches, rule based approaches, or agent based approaches. The operations are conducted mostly in vector space, but there are also approaches,

which apply image processing algorithms to this problem. Thus, the idea is straightforward to utilize the new powerful Deep Learning approaches for the generalization problem. This paper is exploring models from so-called semantic segmentation (Badrinarayanan et al. (2015)). These methods do not provide one classification for a given image patch, but lead to a classification of each pixel. Thus, the output is again an image, however with a label for each pixel. In this paper we adopt this idea. We have already used a similar approach for the generalization of lines from traffic trajectories (Thiemann et al. (2018)), which was inspired by a work on sketch simplification Simo-Serra et al. (2016).

In this paper, the problem of building generalization will be addressed. The traditional approaches to solve it are based on the application of a set of rules in order to eliminate too small building parts. In this case, the map generalization software CHANGE was used Powitz (1992). The paper presents initial experiences with a straightforward implementation of an existing network. The details of the implementation will be reported, together with an in depth analysis of the results. An outlook on future work will be given.

2. STATE OF THE ART

The automation of generalization has been studied by many researchers in recent years - mainly coming from cartography, geoinformatics, and computer science, especially computational geometry (see e.g. Mackaness et al. (2011)). The methods are often based on geometric considerations; the integration of operators is relying on optimization (e.g. Sester (2005)), rule sets, or agent based methods (e.g. Renard et al. (2011)).

*Corresponding author

The approaches are being mainly applied in vector space, however, there are also methods using rasterized representations of the spatial scene. Examples for this kind of operations are aggregation, typification (Müller and Wang (1992)), displacement Jäger (1991), and also building generalization (Damen et al. (2008); Li et al. (2004)). Building generalization involves different elements, such as selection (according to size, type or usage), aggregation (in order to close small gaps) and mainly simplification of the building outline. During this process, the typical shape of a building has to be preserved or even enhanced. This often involves rectifying nearly right angles and enforcing parallel lines in the building footprint, which can be achieved using a set of rules (e.g. Sester (2005); Powitz (1992)). The factors controlling the generalization depend on the scale, and are the area of the building, as well as small structures (facade lengths). When moving to even smaller scales, these parameters would lead to an elimination of most of the buildings, therefore, then typification is applied, i.e. the replacement of the buildings by a building template (mostly square or rectangle) while preserving their spatial arrangement. This operation is applied at scales 1:40.000 and smaller.

There have been early attempts to use Machine Learning to extract cartographic rules from given examples. One goal was to learn suitable parameters of operations (e.g. Weibel et al. (1995)). In their work, the authors observed a human cartographer in order to learn his actions. In a similar way, Mustière (1998) aimed at identifying optimal sequences of operators using Machine Learning. Sester (2000) tried to extract spatial knowledge from given spatial data. While these approaches were very interesting, they mainly remained proofs of concepts.

Deep Learning as a new paradigm has re-emerged in recent years, triggered by the now available computational power (especially exploiting GPUs) – allowing to design also very deep (many layers) and complex networks – and large quantities of available training data. The success in image interpretation was much influenced by the development of new modelling schemes like Convolutional Neural Networks (CNN). They constrain the number of connections in the network to local environments, thus mimicking the human visual system with its perceptual fields, but also constraining the relations of neurons based on neighborhood principles. There are many architectures of Neural Networks, such as nets to classify images - like the popular Krizhevsky et al. (2012), which is able to classify images into 1000 classes. Other networks - LSTM - can learn sequences such as texts or recurring patterns (for an application to derive behaviour of traffic participants, see Cheng and Sester (2018)), identify objects within an image (e.g. Redmon et al. (2016)), or assign a class label to each pixel - which is called semantic segmentation. Such methods are being applied in very relevant tasks such as autonomous driving, speech recognition, health-care and finance, as well as in topographic mapping (Marmanis et al. (2018); Chen et al. (2018)).

In the generalization domain, Machine Learning has been proposed for different applications. Xu et al. (2017) use a deep autoencoder network to assess the quality of building footprint data by learning the characteristics of quality from OSM and authoritative data. Zhou and Li (2017) compare different Machine Learning approaches concerning their capability to select important road links for road network generalization. Lee et al. (2017) use different Machine Learning methods to classify buildings as a prior step for their generalization. An approach for generalization of lines from traffic trajectories using Deep Learning has been presented by Thiemann et al. (2018).

3. APPROACH

Building generalization is composed of different sub-processes: small buildings are eliminated (selection), small parts of the building outline are eliminated, the outline is simplified (simplification), neighboring objects can be merged (aggregation), too small buildings can be enlarged (enhancement), buildings are displaced (displacement) and finally, groups of buildings may be replaced by another group, however, with less objects (typification). The approach presented in this paper aims at an end-to-end training scheme, where a given input and a target output is given, and the system has to learn the "black box" in between. Thus, the expectation is that the system learns all these generalization operations in a holistic way.

In order to train the network, examples have to be provided in terms of corresponding input and output data. Those data sets are available from existing maps, where situations before and after generalization are depicted. In our approach, we rely on an image based approach, which means that the data is prepared in terms of images. One straightforward option would be to use image patches around individual buildings as training data; however, we decided to use the whole map as such, which is cut into regular image patches of given size $b \times b$. When selecting an appropriate size for b , this approach ensures, that adequate context around each building is implicitly used.

3.1 Network Architecture

The structure of the network was designed as Fully Connected Network (FCN), inspired by Simo-Serra et al. (2016). It consists of a series of convolution layers, which are followed by up-convolution layers in order to generate images of the original size. The architecture of the network is shown in table 1. Due to downsampling and upsampling the network tends to have problems with preserving the exact boundaries and induces some blurring. In order to mitigate this effect, so called skip connections which link earlier layers with deeper layers of the network. The inclusion of the skip-connections is visualized in Figure 1.

Type	Kernel	Strides	Output Size
Input	-	-	$1 \times H \times W$
Convolution	3×3	1×1	$24 \times H \times W$
Convolution	3×3	1×1	$24 \times H \times W$
Down Convolution	3×3	2×2	$24 \times H/2 \times W/2$
Convolution	3×3	1×1	$64 \times H/2 \times W/2$
Convolution	3×3	1×1	$64 \times H/2 \times W/2$
Down Convolution	3×3	2×2	$64 \times H/4 \times W/4$
Convolution	3×3	1×1	$128 \times H/4 \times W/4$
Convolution	3×3	1×1	$256 \times H/4 \times W/4$
Convolution	3×3	1×1	$256 \times H/4 \times W/4$
Convolution	3×3	1×1	$256 \times H/4 \times W/4$
Convolution	3×3	1×1	$256 \times H/4 \times W/4$
Up Convolution	4×4	$1/2 \times 1/2$	$64 \times H/2 \times W/2$
Convolution	3×3	1×1	$64 \times H/2 \times W/2$
Convolution	3×3	1×1	$64 \times H/2 \times W/2$
Up Convolution	4×4	$1/2 \times 1/2$	$24 \times H \times W$
Convolution	3×3	1×1	$12 \times H \times W$
Convolution	3×3	1×1	$1 \times H \times W$

Table 1. Model architecture for our building generalization network

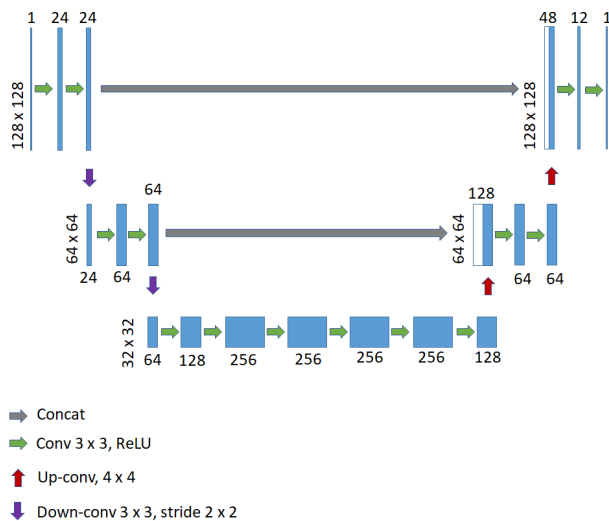


Figure 1. Model architecture for our building generalization network

3.2 Data preparation

Building polygons available at OpenStreetMap were used as input data for training our network. This data has an approx. scale of 1:5000. We generalized the buildings using the software CHANGE. Three different target scales were calculated, 1:10.000, 1:15.000 and 1:25.000. The target scale controls different parameters, such as the minimum length of a facade element (3m, 4.5m 7.5m) or the minimum area (9, 20, 56 m²) to be preserved.

After the generalization with CHANGE, buildings in the original layer and in the target layers were available. Both the input and output were rasterized in 0.5m×0.5m grids. This grid size ensures that the details of the building ground plan are preserved during the rasterization process.

3.3 Training the network

In total, 19.000 image tiles ($b=128px \times 128px$) without overlap were used for training and 1000 image tiles were used for testing. The training was scheduled for 80 epochs, with batch size 32. The Keras / Tensorflow environment was used.

During the training process, we used binary cross entropy as the loss function and Adadelta (Zeiler (2012)) to optimize the loss function. Data augmentation with randomly given rotations was used to avoid duplicated training inputs at each epoch. The training would stop when the loss on the validation data set is not improved over 10 epochs. Only the model which performed the best on the validation set was used for further independent tests. We also investigated and tested using different network architectures and hyper parameters (including dropout rate, learning rate, batch size and tile size), based on its behavior on validation set, the current structure (as shown in Table 1) was selected.

The accuracy score is one of the most often used metrics for evaluation of classification tasks, measuring the correspondence between prediction and target output. However, for building generalization, there are only slight differences between input and output: the differences mainly occur on the boundary of the objects; areas within and outside the objects (especially the large

areas without buildings) may lead to huge amount of True Negative (TN) pixels. Thus, when we simply compare the source and target images, we can directly achieve an accuracy of over 98% (see Table 2, second column). Thus, this metric does not reflect the optimization we did for generalization. Therefore, we used not only the accuracy score but also the pixel wise IoU (Intersection Over Union) as a second metric for our evaluation. In that way, we neglect the huge areas without buildings and only consider if the predicted building are similar to the targets.

For all the three target scales, the networks were trained separately with the same parameters. We conducted an independent test on a map with an extent of 4270px×2560px. For each target scale, we compared input and prediction with respect to the target using the accuracy score and pixel-wise IoU as shown in Table 2. The figures show, that the accuracy values are very high and there is hardly an improvement when going from the input to the prediction. This changes, however, when the IoU is used as a metric for evaluation: for the smaller scales (15k and 25k) the values improve, e.g. from 0.88 to 0.92 in 1:25k. This is not the case for the target scale 1:10k, where there is a small deterioration of the quality.

	Accuracy	IoU
Input vs Target (25k)	0.9818	0.8836
Prediction vs Target (25k)	0.9878	0.9203
Input vs Target (15k)	0.9908	0.9395
Prediction vs Target (15k)	0.9937	0.9582
Input vs Target (10k)	0.9948	0.9652
Prediction vs Target (10k)	0.9936	0.9572

Table 2. Comparison of the accuracy values for different scales and the two metrics Accuracy and IoU

Another model was generated by including the skip connections and also augment the input data by rotating the input patches (see Figure 1). This generated the following accuracy measures shown in Table 3.

	Accuracy	IoU
Input vs Target (25k)	0.9818	0.8836
Prediction vs Target (25k)	0.9901	0.9369
Input vs Target (15k)	0.9908	0.9395
Prediction vs Target (15k)	0.9949	0.9671
Input vs Target (10k)	0.9948	0.9652
Prediction vs Target (10k)	0.9963	0.9761

Table 3. Comparison with skip connections and data augmentation with rotations

This model leads to improvements, i.e. a higher similarity between prediction and target map for all three different scales; the accuracy levels measured by the IoU are between 93 and 97 %.

4. EXPERIMENTS AND RESULTS

4.1 Experiments with specific buildings and orientations

In order to evaluate the behaviour of the network with respect to typical building shapes, a test series of buildings has been produced. All the buildings have offsets and extrusions of different extents. In this test, also the potential dependency on the orientation was investigated by rotating the shapes in 4 different directions. The results are shown in Figures 2 to 5. Please note that

the size of the buildings is different - ranging from approx. 80 m to 25m. The figures visualize the original image on the top and the three generalizations below.




Figure 2. L-shaped building: 1 row: original, 2 row: 1:10k, 3 row: 1:15k, 4: 1:25k




Figure 3. L-shaped building: 1 row: original, 2 row: 1:10k, 3 row: 1:15k, 4: 1:25k




Figure 4. Rectangular building with offsets: 1 row: original, 2 row: 1:10k, 3 row: 1:15k, 4: 1:25k




Figure 5. Small building with annex: 1 row: original, 2 row: 1:10k, 3 row: 1:15k, 4: 1:25k

It can be observed that the extrusions and intrusions are continuously eliminated with increasing scale. It is also clear, that the modifications of the building outlines are very subtle for the larger scales, and get more and more visible for the smaller scales. In Figure 4 the offsets vanish with smaller scales, so does the annex in Figure 5. The purpose of this experiment was mainly to check, if there is a dependency on the orientation of the buildings. This is not the case – obviously, the network has "seen" enough examples of extrusions, intrusions and offsets in different orientations, and is able to generalized them appropriately.

A general observation can be made: the outlines of the generalized buildings are not necessarily exactly straight-lined. Also, in some corners a kind of "overshoot" can be observed (see, e.g. Figure 2, last row).

4.2 Experiments with specific scales

The following tests have been conducted for target scale 1:15.000. The first experiment investigates certain building structures (Figures 6, 7 and 8). The three encircled areas in Figure 8 visualize the generalization capabilities of the model: *A* shows that the intrusions at the corners are filled; *B* shows that the F-shaped building is simplified to a rectangle - similar to the target output produced by CHANGE. The area encircled in *C* visualizes that the system is capable of simplifying very complex building outlines: the many extrusions and intrusions are replaced by a smooth outline. Still, however, the outlines are not exactly straight lines, nor are the corners exact right angles, as given in the target output. Note, however, that these visualizations are enlarged substantially in order to show the generalization effect appropriately. To illustrate the real size, Figure 9 shows input,




Figure 6. Building Polygons from OpenStreetMap

output and the prediction of our model in the dedicated target scale (here 1:15.000). It can be stated, that the effect of generalization of our prediction is clearly positive: as opposed to the original image (left) scaled down to 1:15.000, the forms are simpler, no tiny details disturb the visual impression.

Figure 10 shows the input and target, together with the prediction of our model. Clearly, the outlines are simplified. The z-shaped buildings in the center have been generalized different from the target generalization - however, their is also a valid simplification.




Figure 7. Target output by CHANGE for scale 1:15.000




Figure 8. Prediction with Deep Learning for scale 1:15.000




Figure 9. Buildings shown in original size of scale 1:15.000

In addition to this visual inspection, quantitative tests have been conducted. We compared the map before generalization, map after generalization with CHANGE and map after generalization with our proposed building generalization network (shown in Figure 6 and 10). Based on these two independent test sets, we compared accuracy score and pixelwise IoU of the input and prediction w.r.t. the target as shown in Table 4.


4.3 Map series in different scales

The next visualizations show the same spatial extend in the three different target scales (Figure 11). In the upper row, the solution with CHANGE is given, whereas in the lower row the predictions with the Deep Learning models are shown.


Figure 12 shows the target scales in their respective size.



(a) Building Polygons from OpenStreetMap



(b) Target output by CHANGE



(c) Our prediction

Figure 10. Second test on independent Data

	Accuracy	IoU
Input vs Target (Test1)	0.9764	0.8856
Prediction vs Target (Test1)	0.9906	0.9536
Input vs Target (Test2)	0.9829	0.9261
Prediction vs Target (Test2)	0.9937	0.9723

Table 4. Comparison of quality of results using accuracy and IoU

It is nicely visible, that the models are able to both simplify existing outlines, eliminate too small buildings, and also merge adjacent buildings; i.e. the models contain a combination of different individual operators.

4.4 Results with large extent of data set

Figure 13 shows the result of a larger area in the target scales 1:15.000 and 1:25.000. The extend comprises both residential buildings and industrial buildings, as well as an inner city area. It can be seen, that the Deep Learning model is able to produce appropriate simplifications for the different building types, i.e. more

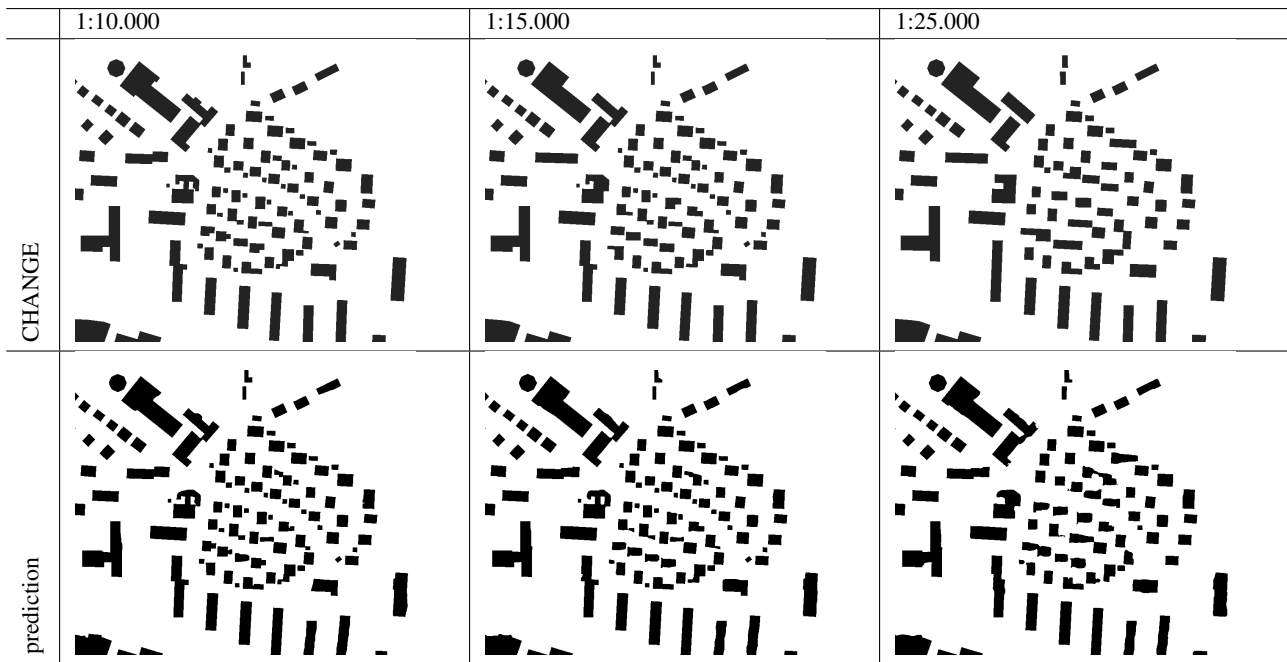


Figure 11. comparison of target scales (top) and predicted presentations (bottom).

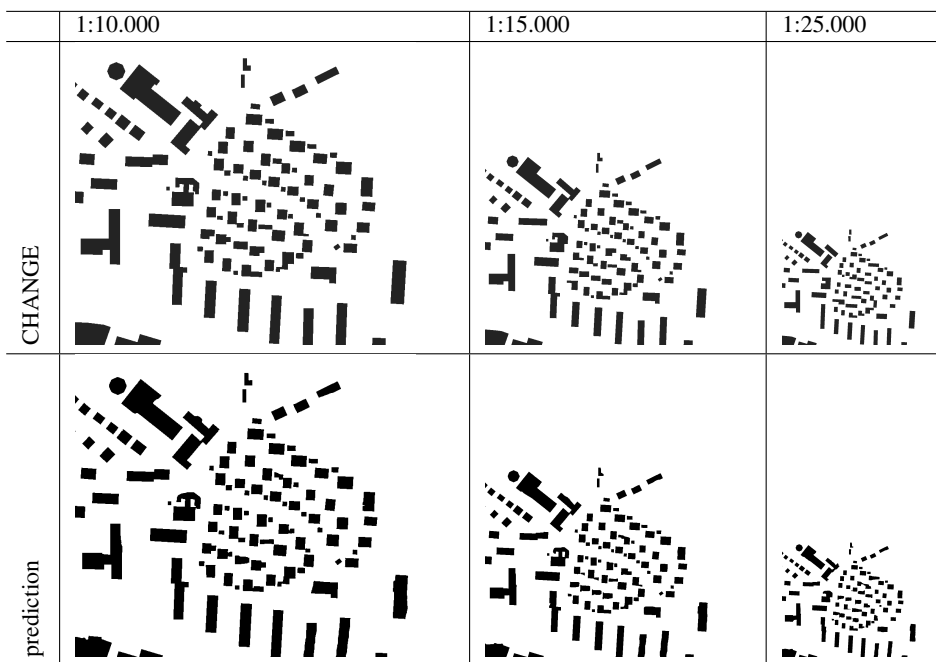


Figure 12. comparison of target scales (top) and predicted presentations (bottom): visualization original target scale size.

rectangular shaped residential buildings, more irregular shaped buildings in the city center, where different individual buildings are merged. It is also visible that small buildings have been eliminated in the smaller scales.

5. DISCUSSION AND OUTLOOK ON FUTURE WORK

The quality of the results can be evaluated both quantitatively and qualitatively. As described above, the general measure of a pixel-wise comparison of the correct predictions is not able to capture the generalization effect. Instead, using the IoU leads to meaningful accuracy values.

A visual inspection of the results clearly indicates that the net-

works have learned the simplification of the buildings for the respective scales: small buildings are eliminated in smaller scales; close, neighboring buildings are aggregated, and outlines are simplified by eliminating small extrusions and indentations. This can nicely be observed in Figure 8, where the buildings with the complicated ground plans in the right part of the image have been simplified to L-shaped outlines. The detailed analysis, however, also showed that characteristic features of most of the individual buildings, namely rectangularity and parallelism are not necessarily preserved or enhanced. Also, it can be observed that some of the corners are rounded out, or small peaks are introduced. A remedy could be to clean up such small irregularities in a post-processing step, using e.g. an approach proposed by Sester and




Figure 13. Scales 1:15.000 (left) and 1:25.000 (right) – residential and industrial buildings, inner city

Neidhart (2008), which explicitly takes these constraints into account and optimises it in an adjustment process.

The visualizations showing the original sizes of the target scales (e.g. Figures 9 and 13) reveal that the result of the prediction can be used for simple presentation graphics: the generalization - simplification of the outline, reducing the content, as well as preserving the overall structure of the buildings and building structure is clearly achieved with the approach.

This work was intended to be a proof-of-concept whether applying or slightly modifying existing Deep Learning networks could achieve satisfying generalization results. The results indicate, that this was successful.

However, there are several avenues to go in the near future: Obviously, the simple loss function included does not enforce that characteristic building shapes are preserved (parallelism, rectangularity). Future work will be devoted to investigate which other loss functions could be applied. Also, other generalization function will be tested, such as typification or displacement. Finally, we want to explore the use of learning in vector space by creating a representation of the building outline, which can be simplified using approaches such as LSTMs. A possible representation could be a vocabulary, which was proposed for a streaming generalization approach (Sester and Brenner (2005)). Another potential approach is to use GANs (Generalized Adversarial Networks) in order to synthesize a new representation based on a given input (see e.g. Isola et al. (2017); Peters and Brenner (2018)).

References

- Badrinarayanan, V., Kendall, A. and Cipolla, R., 2015. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *arXiv preprint arXiv:1511.00561*.
- Chen, K., Fu, K., Yan, M., Gao, X., Sun, X. and Wei, X., 2018. Semantic segmentation of aerial images with shuffling convolutional neural networks. *IEEE Geoscience and Remote Sensing Letters* 15(2), pp. 173–177.
- Cheng, H. and Sester, M., 2018. Mixed traffic trajectory prediction using lstm-based models in shared space. In: *The Annual International Conference on Geographic Information Science*, Springer, pp. 309–325.
- Damen, J., van Kreveld, M. and Spaan, B., 2008. High quality building generalization by extending the morphological operators. In: *11th ICA Workshop on Generalization and Multiple Representation*, Montpellier, France.
- Isola, P., Zhu, J.-Y., Zhou, T. and Efros, A. A., 2017. Image-to-image translation with conditional adversarial networks. *arXiv preprint, CVPR2017*.
- Jäger, E., 1991. Investigations on automated feature displacement for small scale maps in raster format. In: *15th Conference of the ICA*, pp. 245–256.
- Krizhevsky, A., Sutskever, I. and Hinton, G. E., 2012. Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*, pp. 1097–1105.
- Lee, J., Jang, H., Yang, J. and Yu, K., 2017. Machine learning classification of buildings for map generalization. *ISPRS International Journal of Geo-Information* 6(10), pp. 309.
- Li, Z., Yan, H., Ai, T. and Chen, J., 2004. Automated building generalization based on urban morphology and gestalt theory. *International Journal of Geographical Information Science* 18(5), pp. 513–534.
- Mackaness, W. A., Ruas, A. and Sarjakoski, L. T., 2011. *Generalisation of geographic information: cartographic modelling and applications*. Elsevier.
- Marmanis, D., Schindler, K., Wegner, J. D., Galliani, S., Datcu, M. and Stilla, U., 2018. Classification with an edge: Improving semantic image segmentation with boundary detection. *ISPRS Journal of Photogrammetry and Remote Sensing* 135, pp. 158–172.
- Müller, J.-C. and Wang, Z., 1992. Area-patch generalization: A competitive approach. *AThe Cartographic Journal*.
- Mustière, S., 1998. GALBE: Adaptive Generalisation. The need for an Adaptive Process for Automated Generalisation, an Example on Roads. In: *Geographic Information Systems PlaNet*, Lisbon, Portugal.
- Peters, T. and Brenner, C., 2018. Conditional adversarial networks for multimodal photo-realistic point cloud rendering. In: *Martin Raubal and Shaowen Wang and Mengyu Guo and David Jonietz and Peter Kiefer, editors. Spatial Big Data and Machine Learning in GIScience, Workshop at GIScience 2018, Melbourne, Australia, 2018.*, pp. 48–53.

Powitz, B., 1992. Computer-assisted generalization – an important software tool for gis. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. 30/4, ISPRS Commission IV Symposium, Washington, USA, pp. 664–673.

Redmon, J., Divvala, S., Girshick, R. and Farhadi, A., 2016. You only look once: Unified, real-time object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788.

Renard, J., Gaffuri, J., Duchêne, C. and Touya, G., 2011. Automated generalisation results using the agent-based platform cartagen. In: *Proc. 25th International Cartographic Conference (ICC11)*.

Sester, M., 2000. Knowledge acquisition for the automatic interpretation of spatial data. *International Journal of Geographical Information Science* 14(1), pp. 1–24.

Sester, M., 2005. Optimization approaches for generalization and data abstraction. *International Journal of Geographical Information Science* 19(8-9), pp. 871–897.

Sester, M. and Brenner, C., 2005. Continuous generalization for visualization on small mobile devices. In: *Developments in spatial data handling*, Springer, pp. 355–368.

Sester, M. and Neidhart, H., 2008. Reconstruction of building ground plans from laser scanner data. *AGILE08* p. 111.

Simo-Serra, E., Iizuka, S., Sasaki, K. and Ishikawa, H., 2016. Learning to Simplify: Fully Convolutional Networks for Rough Sketch Cleanup. *ACM Transactions on Graphics (SIGGRAPH)*.

Thiemann, F., Cakir, S. and Politz, F., 2018. Generalisierung mittels Deep Learning (CNN) am Beispiel der Straßenextraktion aus GPS-Trajektorien. In: *38. Wissenschaftlich-Technische Jahrestagung der DGPF und PFGK18 Tagung, Muenchen Publikationen der DGPF, Band 27*.

Weibel, R., Keller, S. F. and Reichenbacher, T., 1995. Overcoming the knowledge acquisition bottleneck in map generalization: The role of interactive systems and computational intelligence. In: *COSIT*, pp. 139–156.

Xu, Y., Chen, Z., Xie, Z. and Wu, L., 2017. Quality assessment of building footprint data using a deep autoencoder network. *International Journal of Geographical Information Science* 31(10), pp. 1929–1951.

Zeiler, M. D., 2012. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.

Zhou, Q. and Li, Z., 2017. A comparative study of various supervised learning approaches to selective omission in a road network. *The Cartographic Journal* 54(3), pp. 254–264.