



HAL
open science

Is deep learning the new agent for map generalization?

Guillaume Touya, Xiang Zhang, Imran Lokhat

► **To cite this version:**

Guillaume Touya, Xiang Zhang, Imran Lokhat. Is deep learning the new agent for map generalization?. International Journal of Cartography, 2019, 5 (2-3), pp.142-157. 10.1080/23729333.2019.1613071 . hal-02124904

HAL Id: hal-02124904

<https://hal.science/hal-02124904>

Submitted on 1 Jul 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Is Deep Learning the New *Agent* for Map Generalization?

Guillaume Touya^{a,*}, Xiang Zhang^b, Imran Lokhat^a

^a Univ. Gustave Eiffel, LASTIG GEOVIS, IGN, ENSG, F-94160 Saint-Mande, France, firstname.lastname@ign.fr

^b School of Resource and Environment Sciences, Wuhan University, Wuhan, Hubei, China, xiang.zhang@whu.edu.cn

Abstract: The automation of map generalization has been keeping researchers in cartography busy for years. Particularly great progress was made in the late 90's with the use of the multi-agent paradigm. Although the current use of automatic processes in some national mapping agencies is a great achievement, there are still many unsolved issues and research seems to stagnate in the recent years. With the success of deep learning in many fields of science, including geographic information science, this paper poses the controversial question of the title: is deep learning the new agent, i.e. the technique that will make generalization research bridge the gap to fully automated generalization processes? The paper neither responds a clear yes nor a clear no but discusses what issues could be tackled with deep learning and what the promising perspectives. Some preliminary experiments with building generalization or data enrichments are presented to support the discussion.

Keywords: map generalization, machine learning, deep learning

1. Introduction

The day when automatically generalized maps are as good as manually edited maps is coming closer thanks to the effort of many researchers in the past decades. In particular, agent-based generalization (Ruas & Duchêne, 2007; Duchêne et al., 2018) was a successful paradigm-shift that gave very promising results, even in production environments (Duchêne et al., 2014). But maybe it reached its limits as little new breakthrough research appeared in recent years, at least nothing as high-impacting as the first knowledge acquisition projects, constraint-based modeling, or agent or optimization-based models. One reason is that research generalization processes became so complex that it is very difficult to tune their parameters to achieve good results for a given map (Taillandier et al. 2011). And this parameter problem is not specific to agent-based generalization but applies to other approaches (Harrie, 2003; Zhou & Li, 2016). Moreover, the scope of map generalization has broadened from the needs of National Mapping Agencies to the needs of new mapmakers and map users (Touya, 2017), making a new paradigm shift more urgent. Innovative generalization techniques should be more flexible, easier to parametrize, and adapted more to optimize user experience rather than traditional mapping rules.

When looking at the next paradigm shift that will bring the state-of-the-art one step forward, deep learning seems to be a rightful but tricky candidate. It is a rightful

candidate because deep learning has been successfully used in many applications: image, video and speech recognition, natural language understanding and translation, or analyzing particle accelerator data (LeCun et al., 2015). So why not map generalization? But it is also a tricky candidate because its reputation among information technology scientists clearly overestimates its actual utility in many cases, compared to other artificial intelligence techniques. In order to help the map generalization community to sort this out, several questions are asked in this paper:

Is map generalization a good application for deep learning? Can we go beyond an effective copy of existing techniques? Can deep learning foster creative and on-demand generalization? What can we learn to improve map generalization automation? How can we learn map generalization? We try to answer these questions with discussions and reports of preliminary experiments.

This paper is structured as follows. Section 2 surveys past attempts to use machine learning for map generalization. Then Section 3 explains why map generalization is a good use case to apply deep learning techniques. Section 4 describes the components of a map generalization process that would benefit from machine learning. In Section 5, different strategies to use deep learning are presented and some experiments are reported. Finally, Section 6 discusses the current limitations of this approach and Section 7 draws some conclusions.

2. Machine Learning in Map Generalization

(Weibel et al., 1995) envisioned great potential of machine learning for knowledge acquisition, in order to monitor the map generalization process. This potential was regularly explored in the following to automate different parts of the generalization process:

- enrich data with implicit structures and relations (Plazanet et al., 1998; Sester, 2000; Steiniger et al., 2008; Touya & Dumont, 2017)
- acquire procedural knowledge to orchestrate algorithms (Plazanet et al., 1998; Burghardt & Neun, 2006; (Karsznia & Weibel, 2018)
- acquire procedural knowledge to parameterize algorithms (Plazanet et al., 1998; Cheng et al., 2013; Zhou & Li, 2017)
- evaluate generalized maps (Harrie et al., 2015)

Knowledge was extracted from different types of sources:

- expert interviews (Kilpelainen, 2000; Plazanet et al., 1998)
- traces and logs of interactive generalization from experts (Weibel et al., 1995; Taillandier et al., 2011)
- analysis of a generalized map (Sester, 2000)
- choice by an expert of the best result among several possibilities (Plazanet et al., 1998; Harrie et al., 2015)

But compared to other scientific domains, such as automatic text summarization (Touya, 2015), machine learning techniques were a little underemployed in map generalization in the past fifteen years. In the recent years, the success of deep learning brought back the attention of researchers in map generalization (Ma, 2017; Sester et al., 2018) and this new interest is the very essence of this paper.

3. Why is Map Generalization a Good Application for Deep Learning?

We identified four main reasons why map generalization is good application for deep learning techniques:

- a) Map generalization is a graphical problem, and most of the constraints that guide generalization can be captured by an image of the area to generalize (minimum size, granularity, shape preservation, minimum distances between features, etc.). And deep learning is very well adapted to image processing.
- b) It is very complex to model human decisions when generalizing a map. As a consequence, there is a “knowledge acquisition bottleneck” that was identified in 1995 by Weibel et al., and

this bottleneck is still valid. Deep learning is adapted to such problems where knowledge is implicit in the data.

- c) Deep neural networks are composed of a huge number of parameters and they need massive training datasets to achieve an optimal performance. But we have massive datasets at our disposal, particularly with all the maps generalized manually, semi-automatically or automatically by national mapping agencies (Duchêne et al., 2014).
- d) There are still many problems to solve in map generalization, and new approaches are welcome to solve these remaining issues.

But there are also arguments that support the idea that map generalization is not adapted to deep learning techniques, at least the current techniques:

- a) Image vs Vector: convolutional neural networks (CNN) are effective with image or image-like data while map generalization techniques use geographic databases as input. Deriving generalized maps from images only, causes some information loss (e.g. overlapping features) that might prevent deep learning from being really effective.
- b) Deep learning model are very complex to parameterize to achieve optimal results, which would not solve one of the main reasons of the current stagnation.
- c) The causal mechanisms are hidden: a deep neural network that simplifies a building such as the one presented by Sester et al. (2018) cannot explain how the pixels of the resulting image were classified as buildings, we have to accept the result as it is, which can be a problem with semi-automatic processes with a human editing to finalize the map.

4. What Can We Learn?

(Weibel et al., 1995) stated that machine learning should be limited to learn very specific tasks of knowledge acquisition, in conjunction with other processes based on generalization algorithms (such as agent-based processes that did not exist at the time of the paper). This statement has been true for many years, as shown by the literature reviewed in the Section 2. In this section, we permit ourselves an extension of this scope, in order to discuss a broader use of machine learning in map generalization.

The first step of a generalization process consists in analysing the dataset to extract the implicit structures, patterns and spatial relations, to make them explicit to guide the orchestration of the algorithms. This step is

often called data enrichment. In this data enrichment step, there are many tasks that can be carried out with deep learning. First, the classification tasks such as the characterization of urban blocks (Steiniger et al. 2008) or finding the landmark buildings (Touya & Dumont, 2017) could greatly benefit from deep learning. Similarly, pattern recognition tasks that are common in data enrichment are totally adapted to the use of deep learning. For instance, the detection of highway interchanges is complex task using a classical analysis of the vector data (Touya, 2010), but a preliminary experiment showed the potential of deep recognition models to identify them in images of road networks (Figure 1). The results of this experiment are presented in Section 5.3.

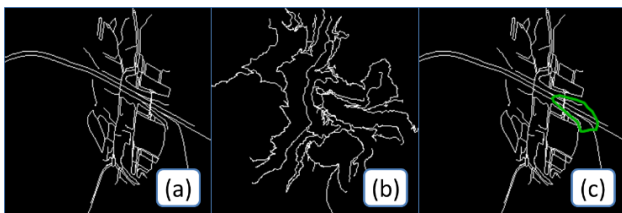


Figure 1. Highway interchange detection in images (a) image with interchange, (b) image without interchange, (c) interchange segmentation.

Besides data enrichment, we can learn how each map feature is generalized at a given scale. Incidentally, the first attempts to use deep learning in map generalization focused on this learning problem for buildings (Ma, 2017; Sester et al., 2018). These attempts are quite successful, which is not really surprising as individual building generalization is mainly driven by graphic legibility constraint. Similar approaches could be used on other map features that are individually generalized at some point of a complete generalization process, for instance on mountain roads that drew attention from many researchers in the past due to the complexity of the problem.

But the existing generalization models already perform very well on individual features, so the gain is minimal if deep learning can only mimic successful algorithms. This is why we need to consider its application to more complex problem, such as the generalization of blocks or groups of buildings. One difficulty to overcome would be the diversity of block sizes that would make it difficult for all blocks to fit in one image. It is currently not known how sizes of map objects or situations as encoded by pixels will influence the learning and prediction capabilities, i.e. what is the optimal image size and resolution to optimise this learning problem?

If we think even bigger, we could try to learn how to generalise a complete map, with all features. But with the diversity of situations that are generalized differently, it

would require an extremely large amount of training examples.

In addition to the data enrichment and the generalization steps, automatic evaluation, which gives feedbacks during the control of generalization and rates the overall quality of generalized maps, could also benefit from deep learning techniques. One of the most challenging tasks in generalization evaluation is the global assessment of a map, which currently still relied on human efforts (Stoter et al., 2014). For instance, Figure 2 shows two generalized outputs of the same area rated by an expert that can be used to learn what a globally well generalized map is.

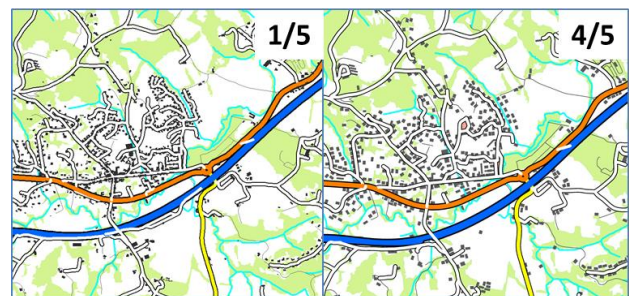


Figure 2. Rating maps as a global evaluation of generalization: the left map is rated 1/5 and the right map, better generalized is rated 4/5.

Deep learning proved very successful on problems of labelling images with descriptive labels, so similar models could be used to label the remaining conflicts or legibility problems in a generalized map (Figure 3).

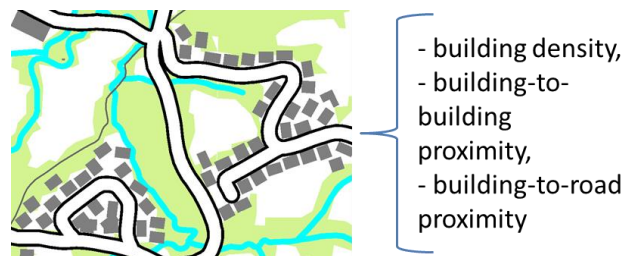


Figure 3. Annotation of the generalization constraints still unsatisfied in the map.

In this section, we discussed what we could learn with deep learning techniques, but we can also discuss what we cannot learn. Past attempts to use machine learning in map generalization mainly focused on the procedural knowledge required to monitor the orchestration and parametrization of algorithms. At first glance, classical deep learning techniques are not really adapted. In some cases, a macro view is necessary to generalize each feature. For instance, road network selection (Touya, 2010; Zhou & Li, 2017) often uses the complete network to compute the importance of one road section. In this case, a single image does not give enough context to make the best decision.

5. How Can We Learn?

While the previous section listed the tasks related to map generalization that could be automated or improved by the use of deep learning techniques, this section discusses the methods that could be used to adapt deep learning techniques to the case of map generalization.

According to Weibel et al. (1995), there are several possible sources of cartographic knowledge: human experts, maps, text documents, and process tracing. In this section, we explain how we can learn only from maps, while some parts of the discussion in Section 6 are related to the addition of the other sources of knowledge in a deep learning model.

5.1 How to Create Training Datasets?

As cartographers, our contribution is more on the derivation of optimal training datasets than on the design of novel deep neural networks. A training dataset is a set of training example, each training example being a pair of input vector (often an image in deep learning) and an output vector. The output vector is a class label in classification problems, a mask in image segmentation problem, and can also be an image (this is the case when learning how to generalize). This subsection describes four different ways to derive such training datasets for learning how to output generalized results.




Figure 4. Three training examples extracted from our building generalization use case (initial buildings on top and generalized ones with agents below).

The first way to create a large training dataset is to use *existing algorithms/processes*. In this case, initial data is used as input, and then generalized by a given algorithm or more complex process; the generalized data is used as output of the training dataset. This is the method used for learning building generalization by Sester et al. (2018) and also the one used in our own experiment on building generalization (Figure 4) where buildings are generalized by an agent-based process (Ruas & Duchêne, 2007). The main advantage of this method is the possibility to easily generate very large datasets, as only initial topographic data and a generalization software tool are required. The main drawback of this method is that the machine learning model can only learn to mimic a process that

already provides automatic generalization and may sometimes give unsatisfactory results. A plugin was developed in the CartAGen open generalization platform (Touya et al., 2019) to generate such training examples from a dataset and a given generalization process.

The second way to create training datasets is to use *existing maps at multiple scales*. The main advantage of this method is that it might capture the knowledge of human cartographers that generalized the maps manually or that edited the output of a semi-automatic process, so it goes further than just mimicking an existing automatic process. The main drawback of this method is the availability of such multi-scale generalized maps. Even if multi-scale generalized maps are now available, for instance on NMA geoportals, the vector dataset is required to generate optimal input/output (e.g. change symbols, or remove features, see the following subsections).



Figure 5. Two representations of highway interchanges (road lines and points) in a MRDB, that can be used to generate a training dataset to classify images that contain highway interchanges.

The third way create training examples is to use *multiple representation databases* (MRDB). MRDBs can provide different abstractions of the same real world geographic features, and these multiple abstractions can be used to generate the proper input/output. For instance, in our experiment on highway interchange detection, we have to generate input images for the two classes (interchange/no interchange), and we used the multiple representation of highway interchange to automatically derive images where there is an instance and images where there is none (Figure 5). The advantage of this method is the possibility to generate large training datasets automatically. The main drawback is the same as the previous one, there is a need for a MRDB, and this is not common. Another drawback is that less detailed abstractions were often subjected to some kind of generalization: for instance, in the highway interchange example, only the main instances were retained in the point representation, so there is no example showing what was considered as less important instances.

The fourth and final way to create training datasets is to use *human annotations*. The main drawback of this

method is the cost in time and money. 20,000 images used to train the U-Net model from (Sester et al., 2018), 1.2 million manually annotated images to train the famous Inception model from Google (Szegedy et al., 2015). In the case of learning generalization, this is not a simple annotation problem, as a training example requires some manual generalization, so the cost of human annotation is even higher. Using crowdsourced annotation, rather than asking expert cartographers, is a possibility but it might introduce some biases in the training dataset with problems of visualization literacy (Borgo et al., 2017). The only advantage of this method is that there is no prerequisite data or software to derive the training dataset.

5.2 Just an Image as Input?

Deep convolutional neural networks are very adapted to input vectors that have a grid format, just like images. This is why images are used as the main input in most deep learning applications.

To generate input and output images from vector initial and generalized data, there are several choices to make. The first one is the size of the image in pixels. Figure 6 shows three images derived from the same building with different resolutions. When the number of pixels decreases, the building boundaries are not sharply displayed which might be a problem as generalization is often a detail reduction issue. But larger images make the learning process much difficult to compute for a computer, as the size of the tensors manipulated by the model depends on the number of pixels in the image. As a comparison, the MNIST dataset that contains images of handwritten digits uses a 28x28 pixels format (LeCun et al., 1998).

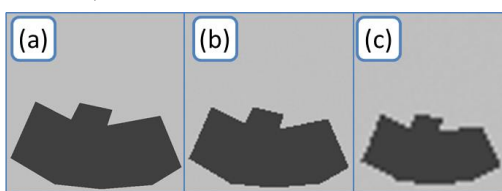


Figure 6. A same building displayed in images with different sizes: (a) 512x512, (b) 128x128, (c) 64x64.

So, what is the optimal size/resolution of images for deep learning map generalization? We studied different possible sizes in Table 1, to find the best compromise for our building generalization use case at a given scale. It should be noted that power of two dimensions are preferred to optimize the use of pooling layers in the convolutional neural network. We used a classical threshold of building minimum size (0.16 mm² in the map) from (Stoter et al., 2009). (Stoter et al., 2009) also shows that more than 95% of the buildings in a

generalized topographic map at the 1:50k are below three times the area threshold, so we used 3×0.16 as the maximum size of a building, to compute the size of a pixel in Table 1. We selected the 128x128 resolution, even at the 1:25k, as a good compromise between pixel size and computing performance.

	1:25k	1:50k	1: 100k
256x256	0.07	0.14	0.27
128x128	0.14	0.27	0.55
64x64	0.28	0.55	1.09

Table 1. Pixel dimensions (in m) according to map scale and training image dimensions for a building to fit in one image with a minimum area of 0.16 mm² in the map.

When an image resolution is chosen, the next issue is the partitioning method. A simple method is to use a regular grid (Figure 7a) as proposed in the literature (Xu et al., 2017; Sester et al., 2018). The drawback of this method is that features can be cut on borders and it will require some merging afterwards. Another method is to derive one image per feature with the feature on the center of the image (Figure 7b). The last method was used in both experiments we carried out. It would be interesting to compare both methods on a same dataset.

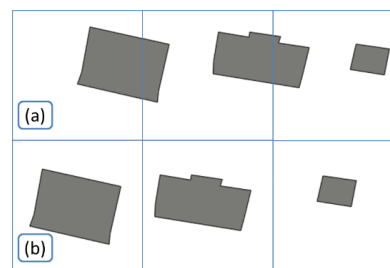


Figure 7. Two ways to generate training images: (a) partition with a regular grid, (b) an image centered on each feature.

The last issue to derive input (and output) images is the background of the image. We identified three different types of background: (1) an empty background (Figure 8a); (2) a map background with the features that influence generalization rendered in the image (Figure 8b); (3) the aerial image as background (Figure 8c). The first solution was used for both experiments, but there should be comparisons to assess which one performs best for different generalization problems.

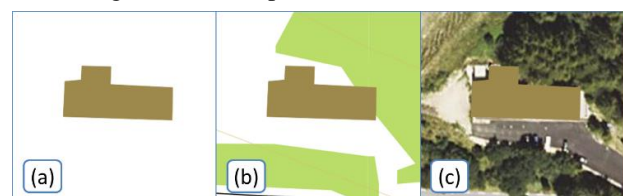


Figure 8. A same building displayed in training images with different backgrounds: (a) empty background, (b) map background, (c) aerial image background.

Even if we try to derive the best image possible, using only an image as the input is a clear limitation compared to the richness of vector data. First of all, many of the constraints that drive map generalization are related to some kind of occlusion/overlap/coalescence between map symbols. For instance, buildings are displaced when they overlap road symbol, which might not be visible on an image (Figure 9a). Using a specific style to highlight these occlusion/overlap problems could be a partial solution (Figure 9b).

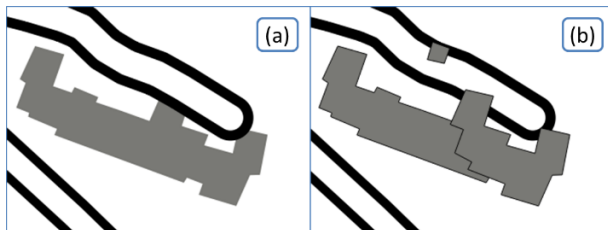


Figure 9. Conveying occlusions and overlaps in training images: (a) one or several buildings? Overlap with road symbol or not? (b) Feature styling to better convey occlusions and overlaps.

There is also a problem if we try to convey spatial relations and meso patterns in the input image. Map generalization needs some explicit modelling of the spatial relations that play a key role in map generalization (Touya et al., 2014). With just an image as input, these spatial relations remain implicit. Deep learning methods allow the addition of metadata dimension in the input tensor, and this might be the solution to give explicit relations and patterns as input of the learning model.

5.3 What Type of Network to Use?

For image classification problems, convolutional neural networks (CNN) are very effective. For our interchange classification problem we used a simple CNN architecture and the results are very positive. 2000 examples of interchange and 800 examples without interchange led to a 95% classification, with 98% of good classification for our remaining 600 interchange classification and 77% of good classification for our remaining 300 examples without interchange. Additional experiments with more examples without interchange are on-going.

When the output is another image, which is the case when we learn how to generalize, or when segmentation is the target, the output of the neural network is not a class for the whole image but a class for each pixel of the image (class “background” or class “generalized building” in our use case). U-net architectures are usually used for such input/output pair (Ronneberger et al., 2015). For instance, a U-Net architecture was used to simplify sketchy drawings (Simo-Serra et al., 2016), and this is

also the case for the first building generalization experiment by Sester et al. (2018). We also tested a U-Net architecture for our building generalization use case, but the results were not conclusive (Figure 10). These disappointing results can be explained by our processing power limitations that prevented the use of training datasets larger than 2,000 buildings.

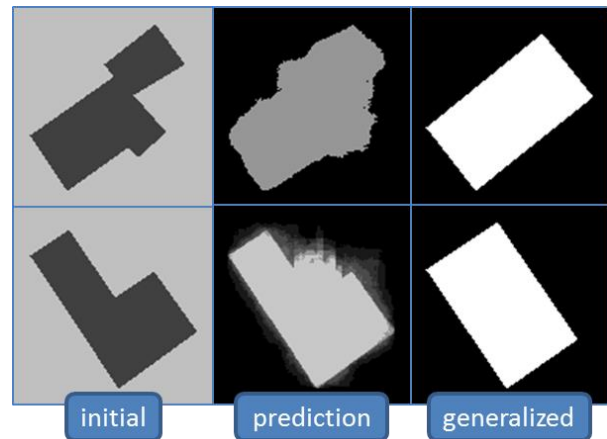


Figure 10. Two examples of building generalization prediction with a U-Net trained with quite few examples (2,000) due to lack of processing power. The paler the pixel is, the higher the probability of being a building pixel is.

Unlike the results of Figure 10, Sester et al. (2018) obtained very promising results with a similar U-Net architecture trained with much more examples, which shows that U-Nets are an interesting kind of model for map generalization. We also successfully used a U-Net to segment the roads belonging to a highway interchange (Figure 1).

But, to achieve even better results on more complex problems (e.g. urban block generalization), we believe that Generative Adversarial Networks (GAN) are the best network architecture. A GAN combines two networks, one that generates many outputs, and one that discriminates the best possible outputs from the ones generated by the first network (Goodfellow et al., 2014). In our case, a GAN could combine a U-Net as the generative network and a CNN to discriminate valid from valid generalized shapes. An interesting use of a GAN is the pix2pix model (Isola et al., 2017) that was successfully used to generate a GoogleMaps-like map image from an aerial photograph of the same area.

Recurrent Neural Networks (RNN) could also be very useful. RNNs are neural networks that contain recurrent connections between neurons, and are adapted to problems where an input i can be better classified using the classification of input $i-1$, $i-2$, etc. RNNs are effective on time series, or for language processing. As the analogy between time and scale already proved its success in the past with the ScaleMaster, RNNs could be useful when

multi-scale, or continuous, generalization is performed, with the generalization at one scale influencing the generalizations at smaller scales.

6. Discussion

Despite its potential, deep learning techniques are not within the reach of any researcher in map generalization. The first reason is economical as deep learning requires high performance computing facilities, and most of research team in map generalization do not have an easy access to such facilities (at least we do not). Labs will have to invest either on computing time in cloud computing platform such as AmazonWebService, or expensive computers with high-performance graphical processing units. The second reason why it would not be an easy transition for map generalization researchers is that deep learning is mainly driven by researchers in computer vision and pattern recognition. As the raw material of most researchers in map generalization is a vector geographic dataset, we will have to learn how to also work with images. We will also have to learn how to map the generalization tasks correctly to the right deep learning models and architectures.

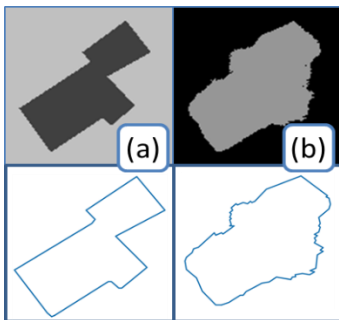


Figure 11. (a) vectorization of a building from the training dataset. (b) vectorization of a fuzzy prediction for the same building.

One of the important consequences of working with images is that the output of a deep learning model for map generalization is not digital cartographic model (Grünreich, 1985), but an image. Even if the images generated as inputs are geolocated tiff images, a vectorization step is required to re-inject the generalized information into the map. Figure 11 shows two examples of a basic vectorization algorithm based on contour detection. In the case (a), the image output of one training example is vectorized, and the algorithm is really able to cope with the low resolution to generate straight lines. In case (b), it is clear that vectorizing a fuzzy output gives extremely bad results. So vectorization should not be a big problem as long as the image result is very good. However, this assertion needs to be verified with other map features, or with multiple features (e.g. an urban block) in the image.

As the generalization of a complete map with a deep learning model is out of reach for now, such techniques will have to be integrated with classical generalization models. For instance, it is possible to integrate specific deep learning models as available operations in an orchestrated generalization model such as the agent-based ones. Compared to other operations of such systems, there should be a focus on the diffusion of the operation on nearby features that were outside the image.

Transfer learning (Pan & Yang, 2010) is also an important issue to discuss. If a model learns how to generalize with examples from a specific place (type of landscape, region, country...), it might give disappointing results on places that are different. The problem is similar with output scale: a model trained for 1:25k might not be able to properly work at the 1:35k scale. Also, a network trained for a specific generalization purpose can be difficult to be tuned for a different purpose (e.g. on-demand and adaptive generalization). Transfer learning techniques enable to re-train models at a minimum cost to adapt them to new locations, scales or purposes, and research will be necessary on the transfer of models for map generalization.

Finally, deep learning techniques would be easier to use in map generalization if they were able to handle vector data rather than images. There is no inevitability in the use of images as input. For instance, the SuperPoint Graph model (Landrieu & Simonovsky, 2018) processes Lidar 3D point clouds as graphs rather than images. Using a graph as input instead of an image could be a way to overcome the limits of the image highlighted in Section 5.2 (overlaps, explicit patterns and spatial relations...).

7. Conclusion and Future Work

To conclude, this paper discussed a possible paradigm shift in map generalization research with the use of deep learning techniques to solve some of our remaining issues. The main contribution of the paper is the detailed description of what could be learned with deep learning and how we can use deep learning for map generalization. Two preliminary experiments (detection of highway interchanges and building generalization) were carried out to illustrate this discussion. Our conclusion is to recognise a huge potential of improvements for map generalization, but there are still important limitations and deep learning is clearly not a magic remedy in our domain.

Our future plan is to continue investigating the topic step by step: (1) we want to improve the results for simple problems such as individual feature generalization or data enrichment; (2) we need to tackle the vectorization/fuzzy

output problem, which is mandatory to integrate deep learning within existing generalization systems (using a GAN architecture?); (3) we plan to try more complex generalization problems such as block/city generalization with multiple features; (4) we want to investigate novel techniques to include information from vector data (i.e. symbol overlaps and spatial relations) into the model.

Otherwise, if we want to reproduce the successes of deep learning in other fields, we need to create shared training datasets and benchmarks to foster research propositions and comparisons. The datasets we used in these first experiments and in future ones will be made available as components of the CartAGen open generalization platform¹ and we hope that they will be the building blocks of future benchmarks.

References

- Borgo, R., Lee, B., Bach, B., Fabrikant, S., Jianu, R., Kerren, A., Kobourov, S., McGee, F., Micallef, L., von Landesberger, T., Ballweg, K., Diehl, S., Simonetto, P., and Zhou, M. (2017). Crowdsourcing for information visualization: Promises and pitfalls. In Archambault, D., Purchase, H., and Hoßfeld, T., editors, *Evaluation in the Crowd. Crowdsourcing and Human-Centered Experiments*, pages 96-138. Springer.
- Burghardt, D. and Neun, M. (2006). Automated sequencing of generalisation services based on collaborative filtering. In Raubal, M., Miller, H. J., Frank, A. U., and Goodchild, M. F., editors, *Geographic Information Science - 4th International Conference GIScience*, 41-46. IFGI prints, Münster, Germany.
- Duchêne, C., Baella, B., Brewer, C. A., Burghardt, D., Battenfield, B. P., Gaffuri, J., Käuferle, D., Lecordix, F., Maugeais, E., Nijhuis, R., Pla, M., Post, M., Regnaud, N., Stanislawski, L. V., Stoter, J., Tóth, K., Urbanke, S., van Altena, V., and Wiedemann, A. (2014). Generalisation in practice within national mapping agencies. In Burghardt, D., Duchêne, C., and Mackaness, W. (Eds.), *Abstracting Geographic Information in a Data Rich World*, pages 329-391. Springer International Publishing.
- Duchêne, C., Touya, G., Taillandier, P., Gaffuri, J., Ruas, A., and Renard, J. (2018). Multi-Agents systems for cartographic generalization: Feedback from past and on-going research. Technical report, IGN, France.
- Cheng, C., Liu, Q., Li, X., and Wang, Y. (2013). Building simplification using backpropagation neural networks: a combination of cartographers' expertise and raster-based local perception. *GIScience & Remote Sensing* 50(5): 527-542.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 27*, 2672-2680. Curran Associates, Inc.
- Grünreich, D. (1985). Ein vorschlag zum aufbau einer grossmassstäbigen topographischkartographischen-Datenbank unter besonderer berücksichtigung der grundrissdateides ALK-systems. *Nachrichten aus dem Karten- und Vermessungswesen, Series I*, 95:55+.
- Harrie, L. E. (2003). Weight-Setting and quality assessment in simultaneous graphic generalization. *The Cartographic Journal* 40(3):221-233.
- Harrie, L., Stigmar, H., and Djordjevic, M. (2015). Analytical estimation of map readability. *ISPRS International Journal of Geo-Information*, 4(2):418-446.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. (2017). Image-to-Image translation with conditional adversarial networks. *CVPR 2017*.
- Karsznia, I. and Weibel, R. (2018). Improving settlement selection for small-scale maps using data enrichment and machine learning. *Cartography and Geographic Information Science*, 45(2):111-127.
- Kilpelainen, T. (2000). Knowledge acquisition for generalization rules. *Cartography and Geographic Information Science*, 27(1):41-50.
- Landrieu, L. and Simonovsky, M. (2018). Large-scale point cloud semantic segmentation with superpoint graphs. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, Salt Lake City, Utah, USA.
- LeCun, Y., Bottou, L., Bengio, Y. and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278-2324
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436-444.
- Ma, L. (2017). Features extraction of buildings and generalization using deep learning. In *Proceedings of 28th International Cartographic Conference*, Washington, DC, USA.
- Mustière, S., Zucker, J.-D., and Saitta, L. (2000). Abstraction-Based machine learning approach to cartographic generalisation. In *Proceedings of 9th International Symposium on Spatial Data Handling*, volume 1a, pages 50-63, Beijing, China.

¹ <https://github.com/IGNF/CartAGen>

- Pan, S. J., and Yang, Q. (2010). A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345-1359.
- Plazanet, C., Bigolin, N., and Ruas, A. (1998). Experiments with learning techniques for spatial model enrichment and line generalization. 2(4):315-333.
- Ronneberger, O., Fischer, P., and Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597.
- Ruas, A. and Duchêne, C. (2007). A prototype generalisation system based on the Multi-Agent system paradigm. In Mackaness, W. A., Ruas, A., and Sarjakoski, L. T., editors, *Generalisation of Geographic Information*, pages 269-284. Elsevier.
- Sester, M. (2000). Knowledge acquisition for the automatic interpretation of spatial data. *International Journal of Geographical Information Science*, 14(1):1-24.
- Sester, M., Feng, Y., and Thiemann, F. (2018). Building generalization using deep learning. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-4:565-572.
- 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*, 35(4).
- Steiniger, S., Lange, T., Burghardt, D., and Weibel, R. (2008). An approach for the classification of urban building structures based on discriminant analysis techniques. *Transactions in GIS*, 12(1):31-59.
- Stoter, J., Burghardt, D., Duchêne, C., Baella, B., Bakker, N., Blok, C., Pla, M., Regnauld, N., Touya, G., and Schmid, S. (2009). Methodology for evaluating automated map generalization in commercial software. *Computers, Environment and Urban Systems*, 33(5):311-324.
- Stoter, J., Zhang, X., Stigmar, H., and Harrie, L. (2014). Evaluation in generalisation. In Burghardt, D., Duchêne, C., and Mackaness, W., editors, *Abstracting Geographic Information in a Data Rich World*, Lecture Notes in Geoinformation and Cartography, pages 259-297. Springer International Publishing.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 1-9.
- Taillandier, P., Duchêne, C., and Drogoul, A. (2011). Automatic revision of rules used to guide the generalisation process in systems based on a trial and error strategy. *International Journal of Geographical Information Science*, 25(12):1971-1999.
- Touya, G. (2010). A road network selection process based on data enrichment and structure detection. *Transactions in GIS*, 14(5):595-614.
- Touya, G. (2015). Lessons learned from research on multimedia summarization. In *Proceedings of 18th ICA Workshop on Generalisation and Multiple Representation*, Rio de Janeiro, Brazil.
- Touya, G. (2017). *Vers l'automatisation de la production de cartes*. Habilitation thesis, Université Paris Est.
- Touya, G. and Dumont, M. (2017). Progressive block graying and landmarks enhancing as intermediate representations between buildings and urban areas. In *Proceedings of 20th ICA Workshop on Generalisation and Multiple Representation*, Washington DC, USA.
- Touya, G., Bucher, B., Falquet, G., Jaara, K., and Steiniger, S. (2014). Modelling geographic relationships in automated environments. In Burghardt, D., Duchêne, C., and Mackaness, W., editors, *Abstracting Geographic Information in a Data Rich World*, pages 53-82. Springer, Berlin.
- Touya, G., Lokhat, I. and Duchêne, C. (2019) CartAGen: an Open Source Research Platform for Map Generalization. Submitted to ICC'19.
- Weibel, R., Keller, S., and Reichenbacher, T. (1995). Overcoming the knowledge acquisition bottleneck in map generalization: The role of interactive systems and computational intelligence. In Frank, A. U. and Kuhn, W., editors, *Spatial Information Theory, A Theoretical Basis for GIS*, pages 139-156. Springer, Berlin.
- 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):1929-1951.
- Zhou, Q. and Li, Z. (2016). Empirical determination of geometric parameters for selective omission in a road network. *International Journal of Geographical Information Science*, 30(2):263-299.
- 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):254-264.