
DEEPMAPSCALER: A WORKFLOW OF DEEP NEURAL NETWORKS FOR THE GENERATION OF GENERALISED MAPS.

This version of the article has been accepted for publication after peer review, but is not the Version of Record. The Version of Record is available online at: <https://doi.org/10.1080/15230406.2023.2267419>

Azelle Courtial

LASTIG

University Gustave Eiffel, ENSG, IGN,
F-94160 Saint-Mande, France
azelle.courtial@ign.fr

Guillaume Touya

LASTIG

University Gustave Eiffel, ENSG, IGN,
F-94160 Saint-Mande, France
guillaume.touya@ign.fr

Xiang Zhang

School of Geospatial Engineering and Sciences,
Sun Yat-Sen University,
Guangzhou 510275, China

ABSTRACT

The automation of map generalisation has been an important research subject for decades but is not fully solved yet. Deep learning techniques are designed for various image generation tasks, so one may think that it would be possible to apply these techniques to cartography and train a holistic model for end-to-end map generalisation. On the contrary, we assume that map generalisation is a task too complex to be learnt with a unique model. Thus, in this article, we propose to resort to past research on map generalisation and to separate map generalisation into simpler sub-tasks, each of which can be more easily resolved by a deep neural network. Our main contribution is a workflow of deep models, called DeepMapScaler, which achieves a step-by-step topographic map generalisation from detailed topographic data. First, we implement this workflow to generalise topographic maps containing roads, buildings, and rivers at a medium scale (1:50k) from a detailed dataset. The results of each step are quantitatively and visually evaluated. Then the generalised images are compared with the generalisation performed using a holistic model for an end-to-end map generalisation and a traditional semi-automatic map generalisation process. The experiment shows that the workflow approach is more promising than the holistic model, as each sub-task is specialised and fine-tuned accordingly. However, the results still do not reach the quality level of the semi-automatic traditional map generalisation process, as some sub-tasks are more complex to handle with neural networks.

Keywords Map Generalisation · Generative Adversarial Network · Deep Learning · Workflow

1 Introduction

Automatically generating a topographic map at a target scale from detailed geographic data has been the moonshot objective of map generalisation research for years. A process able of such a generation would benefit many applications of geo-visualization and cartography. However, it requires the use of complex generalisation processes: first, data is enriched to exhibit important structures and properties of the map, then different algorithms are applied to different portions of the map using complex workflows, and finally, data is symbolised and organised in a map [Mackaness et al., 2014].

With the emergence of image generation deep learning techniques applied to cartography, one might seek to design models to generate a symbolised and generalised topographic map image from an image of a detailed geographic

database [Kang et al., 2019]. However, the current state-of-the-art on deep learning-based map generalisation does not support such an optimistic goal [Courtial et al., 2021b], for the following reasons:

- Deep neural networks are able to use implicit information in the input to achieve generalisation. However, it seems complex to create a set of inputs containing all information needed for map generalisation. For example, in Figure 1 we observe that map tile images are not sufficient to represent all the information needed for road selection: the function of the road is necessary but not accessible in the raster; moreover, the limited tile size does not provide a good assessment of the role of the road section in the whole network.
- The design and evaluation of deep learning models for map generalisation is complex [Courtial et al., 2022a] and the separation into simple sub-tasks can help to achieve this goal. First, the evaluation of object generalisation and relation preservation is more challenging in images than with vector data, because they require the identification of different map objects from images. Furthermore, by improving the experiments of [Courtial et al., 2021b], we observed that it is complex to improve one aspect of the model performance without damaging the others. For example, adding a loss term to penalise structure alteration did reduce the occurrence of disconnected roads, but also led to a degradation in building generalisation quality. On the contrary, adding a loss to improve the shape of buildings often leads to disconnected roads.
- Deep learning processes are a kind of black box that infer knowledge from training examples. Thus, a model that tries to generate a complete generalised map is complex to control. In other words, different levels of complexity in map generalisation sub-tasks may require different numbers of training iterations and involve examples of different characteristics. For example, learning to transform a city centre into a built-up area and learning to enlarge buildings would require different examples and training.



Figure 1: The necessary information to learn the selection of road is not visible in tiled images. The section circled in green is selected because it is a dead-end but the tiling hides this context. The section circled in blue is selected because it is a local path but this information is not symbolised in the tile.

We, therefore, claim that an architecture that carries out map generalisation step-by-step is currently more appropriate than training a holistic generalisation model end-to-end. To construct such an architecture, we resort to past research on map generalisation [Brassel and Weibel, 1988] and propose separate sub-models for the main steps of map generalisation, each of which can be more easily resolved by a deep neural network. Our main contribution is a workflow of the deep models, called DeepMapScaler, which describes a step-by-step process for generalised map generation, where deep learning is used at most steps. It is organised in two main steps: the first step aims to adapt the level of detail of objects to the target scale. Then, the second stage compiles the layers together in a symbolised map.

The goal of this article is to illustrate the interest of such a workflow approach over the holistic approach for end-to-end map generation. For this comparison, we designed and tested both approaches for a topographic map composed of roads, buildings, and hydrography, generated at 1:50k in a region that includes urban and rural areas. The implemented workflow mixes image-based and graph-based deep learning models in order to generate a generalised map that is legible and preserves main geographic information and patterns. We perform an evaluation for each step of the workflow; then our approach is compared with the holistic generation.

The article is organised as follows: the next section reviews the existing research papers on the usage of deep learning for map generalisation. Then, Section 3 presents the general structure of the proposed workflow. Section 4 presents more details on the implementation and model used in our experiment. The experiment results are then presented in Section 5, with a focus on the evaluation of the deep learning predictions at each step of the workflow. Finally, in Section 6 we discuss the advantages, limitations, and possible improvements of the proposed workflow.

2 Related work

In this section, we present a brief overview of the diversity of approaches in map generalisation, which leads to the identification of three use cases where deep learning can benefit map generalisation. We then review the existing literature on each of these use cases.

2.1 Map generalisation

Map generalisation is a complex process that relies on multiple factors including geometry, semantics, context, etc. The information in the geographic databases is often not sufficient to perform map generalisation, and the first step is to make explicit some important patterns of the map [Mackaness and Edwards, 2002]. Then, map generalisation usually requires several steps. For instance, [Brassel and Weibel, 1988] proposes five steps: structure recognition, process recognition, process modelling, process execution and display. In another approach, [Shea and McMaster, 1989] proposes an approach of map generalisation based on "when", "why" and "how" to generalise. With such an approach the steps are the following: detection of local conflict, choice of the relevant algorithm for its resolution, local resolution and then integration of the corrected part to the map. This approach has been used for example for mountain road generalisation [Mustiere, 1998], or building displacement [Aslan et al., 2012].

The development of research on map generalisation went into three directions: 1) the definition of map constraint to formalise the requirements of map generalisation [Beard, 1991], 2) the design of operators, i.e. algorithms for resolving an atomic of the problem [Stanislawski et al., 2014] and 3) their orchestration in a map generalisation process that can deal with whole maps [Duchêne et al., 2018]. Currently, plenty of operators involving various techniques have been proposed to resolve different purposes in map generalisation. For instance, the problem of road selection was first resolved using rule-based processes, then it included notions from graphs [Thomson and Richardson, 1995, Porta et al., 2006b], centrality analysis [Weiss and Weibel, 2014], structures made from continuity and semantics [Thomson, 2006, Benz and Weibel, 2014]. Despite all these proposals, it is still complex to choose and implement the relevant algorithms, and even for some other problems, there is still no correct solution. That is why more and more researchers believe that deep learning could be an important tool in unlocking the research for map generalisation automation [Touya et al., 2019b]. The steps where deep learning can contribute to map generalisation are the following: 1) as a data enrichment process; 2) as a map generalisation operator; 3) for generalised map generation.

2.2 Deep learning as a data enrichment process

To preserve the key geographic properties of the map during the abstraction process, map generalisation usually requires a data enrichment step, where implicit properties, structures, or patterns are made explicit [Mackaness and Edwards, 2002]. For instance, this step can make roundabouts or dual carriageways explicit in a road network, to enable their proper typification [Touya, 2010]. Though techniques not based on machine learning were proposed in the past, the structures and patterns that we want to identify are well adapted to a description with examples, so machine learning has been used as a key technique for data enrichment [Sester, 2000, Steiniger et al., 2008, Garcia-Balboa and Ariza-López, 2008, Zhou and Li, 2017]. However, the range of applications with sufficient results is still limited, mostly due to the fact that we need to extract significant features to describe the pattern. In many fields, the recent deep learning approaches overcome this problem by directly extracting implicit information from the raw data and have enabled new applications with increased accuracy.

Among the many ways to model cartographic data enrichment as a deep learning task, several approaches are tested in the literature. The first approach is to convert cartographic data into small images, and have convolutional neural networks (CNN) classify the pixels of the image to identify the patterns. This approach was used to find the roads that belong to a highway interchange [Touya and Lokhat, 2020], to classify the different types of complex intersections in a road network [Li et al., 2020], or to classify the morphology of urban units, i.e. groups of buildings [Chen et al., 2021].

Then rather than using a CNN, researchers tend to favour graph deep learning architectures, as cartographic information is often better represented with a graph. Applications then differ in the way the graph is built from the cartographic data. The most simple approach is to derive a graph from the vertices and segments that compose a polygon or a line in the map. Thus, it is possible to classify the building shapes in relation to selected templates (e.g. rectangular

or L-shaped) [Yan et al., 2020a, Liu et al., 2021a, Hu et al., 2022], or to group the close buildings that have similar shapes [Yan et al., 2020b]. Graphs are also usual to model the cartographic networks [Mackness and Beard, 1993]. Thus, graph convolutions can be used to classify the geomorphological patterns of rivers [Yang et al., 2022a]. Finally, graphs can be used to represent different types of spatial relations between cartographic objects as a graph, where the cartographic objects are the nodes and where each pair of spatially related objects are linked by an edge [Iddianozie and Mcardle, 2021]. For instance, a proximity graph between buildings can be derived from a Delaunay triangulation to learn a grouping of buildings into specific patterns [Zhao et al., 2020]. With a similar proximity graph between buildings, [Ma et al., 2022] derive fractal indices of the urban forms, and [Yang et al., 2022b] classify building blocks (e.g. residential or industrial).

2.3 Deep learning as a map generalisation operator

Secondly, the deep learning ability to abstract information has led researchers to study its interest in performing map generalisation operations. In these attempts, a model is trained to reproduce one or a few generalisation operators. A generalisation operator is an atomic transformation of the cartographic objects, such as simplification, smoothing, or displacement. For instance, the graphic generalisation of buildings, coastlines and roads using an image-based approach has been investigated [Feng et al., 2019, Du et al., 2021, Courtial et al., 2022b]. These studies employed a segmentation model to predict pixels of generalised objects in the output image. These experiments give quite good results and highlight the potential of deep learning to produce enlargement, simplification, typification and smoothing operations. The reduction of building density also seems to be successful, while the selection and displacement of road sections could not be learnt with such an approach. Moreover, other approaches based on graphs or point series can be used to predict the simplification of line [Yu and Chen, 2022, Du et al., 2022] or a building [Zhou et al., 2022]. In these approaches, an artificial neural network is trained to predict if points are retained in the generalised shape of the object, and then the point set is post-processed in a generalised entity. From our observation, these proposals are promising as their formulation allows integration in traditional map generalisation mechanisms and avoids disconnection problems. But the quality of generated shapes is still a challenge in some parts that remain either too smooth, too simplified or contain too sharp angles.

2.4 Deep learning model for generalised map generation

Finally, the emergence of generative adversarial networks (GAN) has allowed the multiplication of map generation use cases. Indeed, the GANs are models designed for image-to-image transformation. The transformation from an aerial photograph into a map is a common use case of these models [Isola et al., 2017]. This use case demonstrated the ability of a GAN to generate map-looking images. However, it is based on GoogleMaps large-scale tiles, and the conversion includes very few changes in the level abstraction of data, and thus no generalisation. [Kang et al., 2019] have trained similar models for other map style transfer cases, in particular, to apply the style of GoogleMaps to OSM data presented in a raw stylised image. This work has shown that GANs not only are able to generate map-looking images but also to symbolise and abstract geographic information with scale changes. In a similar way, some experiments with the generation of generalised topographic maps from un-generalised images of these maps were carried out [Courtial et al., 2022c]. That demonstrates the ability of the model to learn some generalisation operators from image examples. This work also highlights that designing a training dataset that includes all necessary input information to learn all the generalisation operations, is a complex challenge. We claim that it is currently not possible to design an adapted training set for the whole generalisation and separating tasks would make this task simpler.

3 The DeepMapScaler Workflow

3.1 General organisation

The review of the literature shows that deep learning models are able to recognise implicit patterns, mimic some operators and generate map-looking images. However, they are not able to perform all these transformations together, and the creation of an adapted training set, or the choice of a good architecture would be very complex tasks. To enable the design of an adapted training set and model, we separate the problem into several simpler tasks. Thus, it is possible to apply relevant architectures to each smaller problem. Our workflow, called DeepMapScaler, describes a process for generalised map generation, making use of several deep-learning models. The aim is to use as input a vector database containing the desired layers (e.g. buildings, roads, rivers or others: labels, forests, isocontours, etc.) and produce as output an image of a generalised topographic map. It is organised in two main steps (illustrated in Figure 2): 1) adaptation of the level of detail for the desired scale and 2) map generation, which compiles in a generalised map.

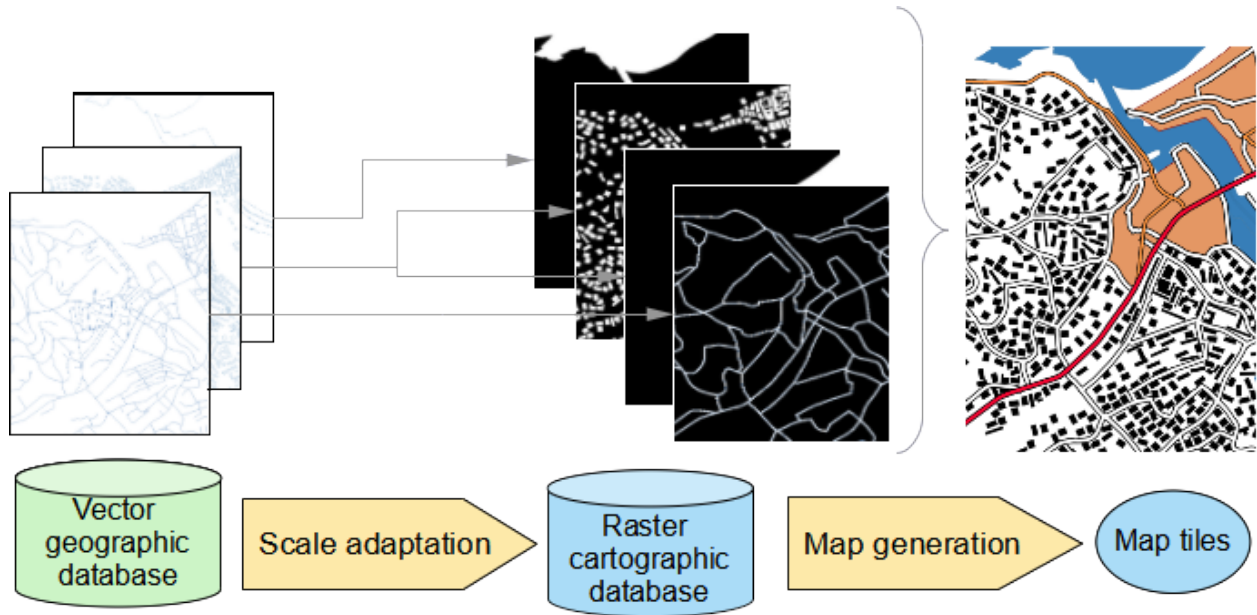


Figure 2: DeepMapScaler main steps: the scale adaptation transforms the vector geographic database into a raster cartographic database, and then the map generation transforms it into a map.

3.2 Scale adaptation

The first part of the workflow is called scale adaptation; it deals with the level of detail of the geographical data theme by theme. The goal of scale adaptation is to produce a raster cartographic database at the target scale, composed of one binary mask for each theme in the target map. In an analogy with traditional map generalisation, this step can be compared with the creation of a cartographic database with a certain level of detail (see Figure 3).

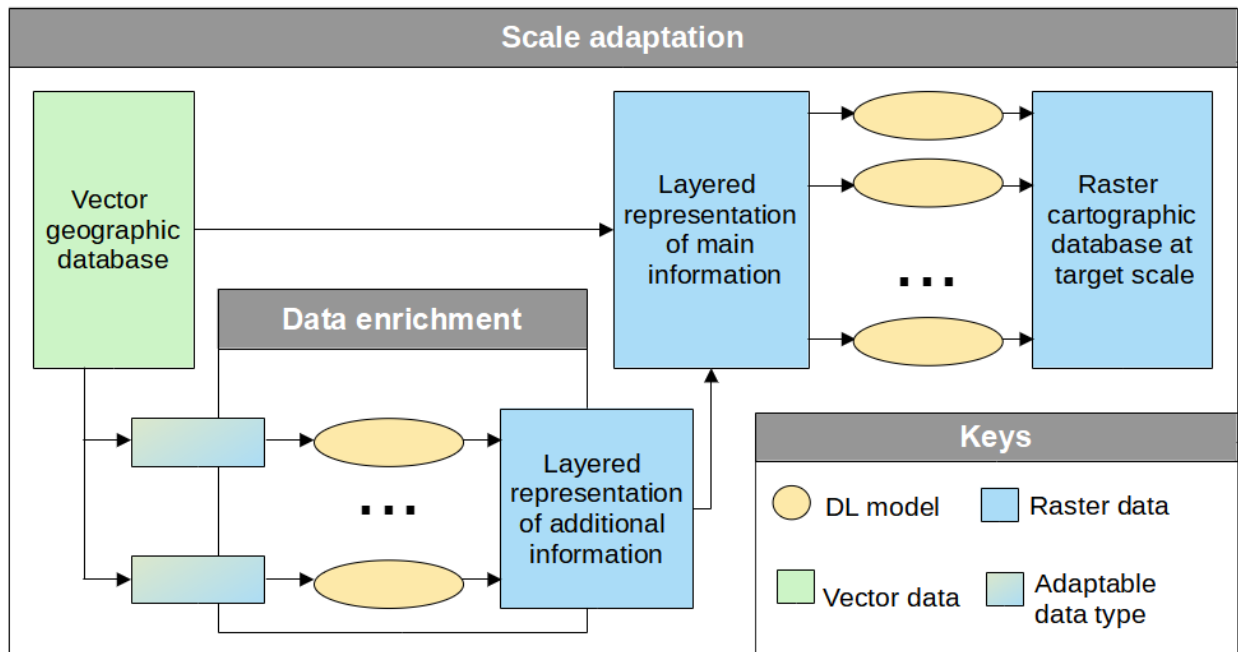


Figure 3: DeepMapScaler, detail of scale adaptation step: a subset of input data is enriched using the relevant model and representation, then it is combined with the main information to produce a raster cartographic database.

First, the necessary information for scale reduction is identified and prepared in a layered representation. This representation was introduced by [Courtial et al., 2022c], it is a simple way to encode cartographic information with deep models. It is composed of two tensors: one that stacks tiles of the "main information" for cartography (the geometry of geographic entities), and one that stacks in tiles of the same extent additional information required to construct the map (e.g. semantics, contexts, etc.). There are two ways to construct the additional information tensor: if it is already explicit in the input vector data (e.g. semantic information such as a function), the information is tiled to be integrated into the relevant tensor of the layered representation. Otherwise, some data enrichment process is performed using the relevant encoding (raster, vector, graph etc.) and model, to be then integrated into the layered representation.

Secondly, a deep model is trained (for each theme) to predict the generalisation of this theme. For instance, the generalised shape of the buildings can be learnt from 1) primarily, the geometry of the buildings and of other themes, and 2) additionally, the function of the building (derived from attributes) and the limits of the city centre (learnt from a graph or image segmentation model). This deep model can be image-based or graph-based, depending on the theme.

Scale adaptation is performed with one or several models for each theme and may include data enrichment, object selection, and graphic generalisation. Thus, this step mainly focuses on individual theme generalisation and does not resolve inter-theme conflicts, such as overlaps between symbols of roads and buildings.

3.3 Map generation

The second step, called map generation, assembles images of generalised objects from the previous step, into a symbolised map. It is responsible for the preservation of inter-theme spatial relations, the management of style-related conflicts and for the global legibility of the map. This step is performed by an unpaired image generation deep neural network such as [Zhu et al., 2017]. These models are trained to generate images that represent the content of the input with the style of the target. In our case: producing an image of the generalised map from raw independently generalised data.

4 Experiment

In this section, we propose to test the general workflow with a specific implementation for a simple use case.

4.1 Use case

In the proposed use case, we chose a study area that includes both urban and rural areas, located near the town of Saint-Jean-de-Luz, which covers 30 * 15 kilometres in the southwest of France.

The input is a detailed topographic database produced by the French National Mapping Agency (IGN) composed of roads, buildings, and hydrography (river lines and water areas). The level of detail of the database approximately corresponds to a 1:25k scale. The target is a map generalised and symbolised at 1:50k. It is the result of a traditional semi-automatic map generalisation process using the open-source map generalisation software CartAGen, [Touya et al., 2019a]. In the target map, the hydrography is only slightly changed: some minor rivers and water surfaces are deleted using rules-based selection on attribute values. The geometry of the kept rivers and water areas is not changed. The road network is selected (about 30% of road sections are selected) and some intersections are schematised according to the strategy proposed by [Touya, 2010]. The shape of a selected road is globally maintained: very few sections are smoothed, simplified or displaced to avoid overlapping or increase legibility at this scale. The representation of buildings at this scale is more challenging. This is a medium-large scale and contrary to a smaller scale, the map still represents the individual buildings in most cases. Thus, a satisfying balance between constraints of minimum size and minimum spacing is complex to find. In particular, the number of buildings has to be reduced while preserving patterns and relations. The used generalisation process is combining the AGENT model [Barrault et al., 2001] and a typification algorithm [Burghardt and Cecconi, 2007]. Finally, city centres are small and dense blocks that cannot be generalised with enlarged individual buildings and they are converted into built-up areas [Touya and Dumont, 2017]. Built-up areas must be areas in the city centre, where the density of buildings is important compared to the rest of the city. It might be a coherent area formed of one or more blocks.

4.2 General design

Thus, our implementation includes the following sub-tasks, organised as presented in Figure 4.

- creation of a block covering mask (data enrichment)
- creation of a road value map (data enrichment)

- creation of generalised road tiles (scale adaptation)
- creation of generalised building tiles (scale adaptation)
- map generation

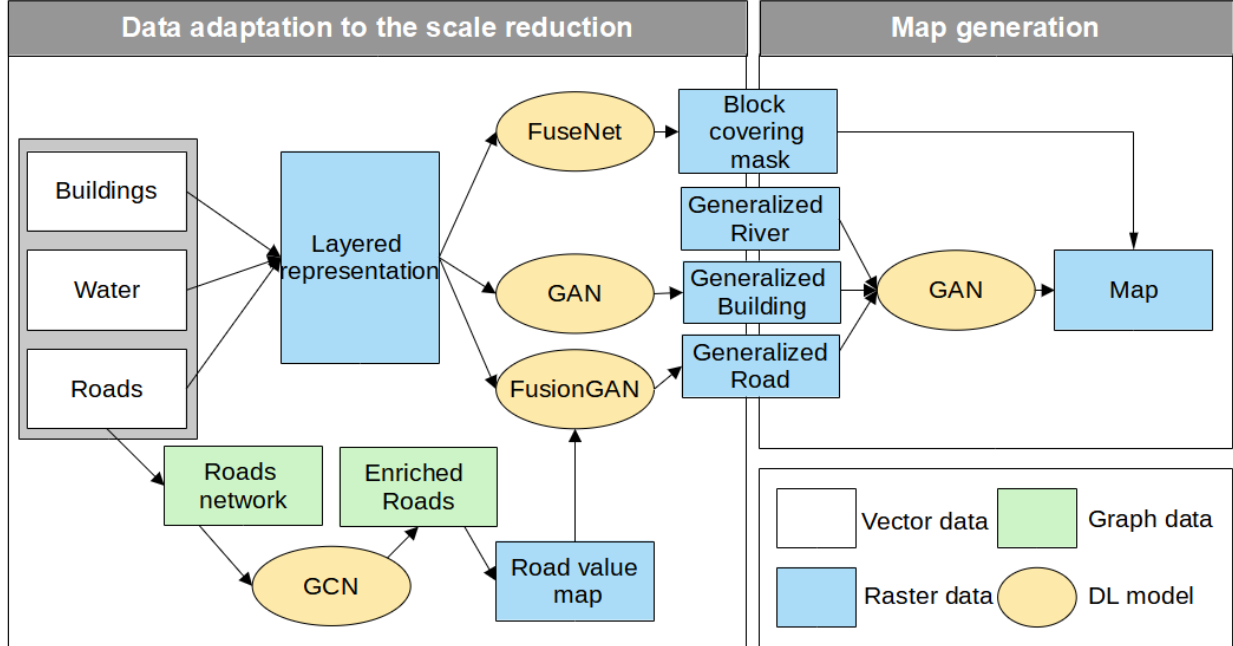


Figure 4: The implementation of DeepMapScaler used for our experiments, buildings, roads, and rivers are generalized and symbolized in a topographic map at medium scale.

To summarise, the scale adaptation step for roads at 1:50k scales is performed in two sub-steps as it is a complex task. Traditionally, it may first require data enrichment to detect particular structures (e.g. interchange, roundabout, etc.), then the road network is selected to reduce its density, and finally, the road shape can be simplified or smoothed [Touya, 2010]. Here, the steps are the following: 1) a graph-based learning approach to encode and analyse the structure of the road network and learn the probability for each section to be selected; 2) an image model to perform the graphic generalisation of the roads. This part of the workflow is illustrated in Figure 5. Then, building adaptation is also a two-step process (block covering and individual building generalisation) but contrary to road adaptation, the two steps of building adaptation are not connected and are performed independently, which means that the first step generalises all buildings including those in covered blocks. Similarly to traditional map generalisation, the block covering information is only used at the end of the workflow to cover the relevant part of the output map.

Finally, in this implementation, we did not implement any model for water generalisation using deep learning as it is really simple for the chosen scale gap.

Each model is designed and set up independently, with an adapted training set. We used a PyTorch implementation of the different deep learning models (FusionGAN, pix2pix, CycleGAN). The models used are detailed in the next paragraphs.

4.3 Training set creation

To create a training set of tiled images, the study area is regularly separated into square tiles using a sliding window [Courtial et al., 2020]. The code for raster tile creation is published as a plugin of CartAGen open-source map generalisation software [Touya et al., 2019a]. Each tile is composed of 512×512 pixels and represents a surface of 500×500 meters. Each tile has a 40% overlap with the neighbouring tiles. We obtain around 3600 tiles. This process was applied to generate the training sets for most of the tasks performed in raster format, but not for block covering, as it is quite rare in our use case, and we did not have enough tiles for training. Thus, we annotated data in other regions of France, to obtain around 1200 tiles containing block-covering examples. Road data enrichment was performed using a graph directly derived from connectivity in the vector data.

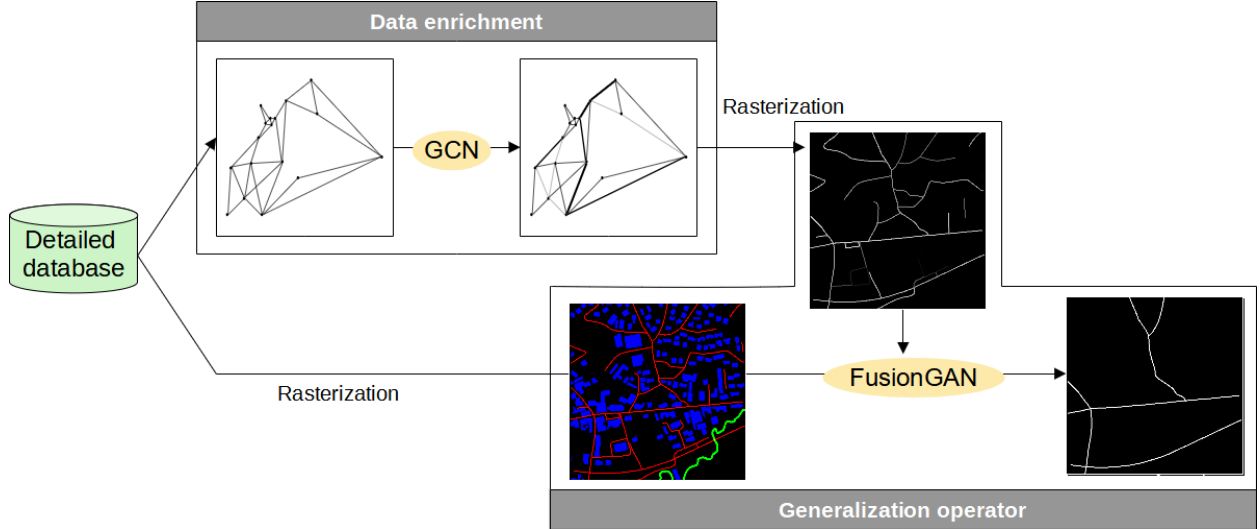


Figure 5: Road scale adaptation is a two-step process: first, a graph model is trained for data enrichment, then a FusionGAN gathers initial and enriched information to generalise the road shape.

4.4 Creation of generalised building tiles

We expect the model to perform several operations on buildings: enlargement, displacement, typification, and deletion. We propose to use pix2pix [Isola et al., 2017] for this step, as it gave good results on building generation in the past, with a reasonable computation speed [Feng et al., 2019, Kang et al., 2020, Courtial et al., 2021b]. This architecture is made of two CNNs that are trained in an adversary manner: a generator that is trained to predict an image that looks like the target domain (i.e. a topographic 1:50k map), and a discriminator that is trained to classify a pair of images (input, output) between real (from the training set) and fake (with an output that is generated by the generator). The generator then learns to generate images that are increasingly realistic to fool the discriminator.

Previous experiments with this model sometimes result in blurred or unrealistic shapes [Feng et al., 2019]. We hypothesise that the separation in simpler sub-tasks may reduce the quantity of change that the model must learn and may yield results that are less blurred. Furthermore, we also propose to add a geometry consistency loss [Fu et al., 2019] to the default pix2pix loss. This loss compares the rotation of the predicted image and the prediction of the rotated input image (eq. 1). It aims to speed up learning and should result in a less blurred shape.

$$G(t(i)) \approx t(G(i)) \quad (1)$$

where $G(i)$ is the prediction of the model on image i and $t(i)$ is the transformation of image i the model should be invariant to (rotation in this case).

The input of the model is the layered representation of buildings, roads, and water, and the target is a mask of generalised buildings. This representation stacks the mask of each main theme of the target map at the input scale. We choose a layered representation of all themes in addition to the building theme. This is because the density and the nature of the cartographic features around a building or a building group are important to understand good generalisation decisions. In particular, it is necessary to know where there is free space around the building to be generalised.

4.5 Creation of a block covering mask

We also expect the output to include a built-up area that covers buildings in dense city centres to avoid overlapping and cluttered views. This task concerns a small number of tiles in our training set and cannot be learnt together with the individual building generalisation [Courtial et al., 2021b]. We believe that the geometries of buildings, roads and water features are necessary for this task but not sufficient. Indeed, this block-covering task is usually guided also by the function of the block and the density of the neighbouring blocks [Touya and Dumont, 2017]. Thus, we propose to provide as additional information a classification of the block [Touya, 2021] obtained by a multi-criteria decision, computed on all the blocks of a town or city.

Then, a FuseNet [Hazırbaşı et al., 2016] is trained to predict a mask of covered blocks. This is a CNN architecture with two encoding paths that fuse two levels of information. The information from the additional information path is fused, by element-wise summation, with the main information at a different encoding step. Finally, the fused information is decoded into an output of the desired shape.

4.6 Encoding road importance

A graph-convolution network is trained to classify road sections using the BasqueRoads dataset [Touya and Courtial, 2021] that considers the same scale change. The target is a label "kept" or "deleted". The probability that the label is "kept" is used as an indication of the role of each section in the global network and, particularly, regarding its selection on a generalised map (e.g. if the section belongs to a long stroke, to an important itinerary, etc.). This part is inspired by the experiments of [Courtial et al., 2021a].

The model used is a GCN. GCN are deep neural networks designed to deal with graphs as input. In these models, the graph convolution is defined using the node attribute matrix X (that describes for each node the list of its features) and the adjacency matrix A (that describes the relationship between nodes). For each layer, $l + 1$, the convolution result is given by eq. 2 with $H(l)$ being the result of the previous layer and a weight matrix $W(l)$. σ is the activation function. The goal of the model is to optimise the weights to minimise differences between the target class and the prediction for all the nodes in the training set. Thus, the attributes of neighbouring nodes are summed at each layer, and, the model learns an embedding for each node. The deeper the model is, the more neighbours are considered for each prediction on a node.

$$\{ H(0) = X \quad H(l + 1) = \sigma(A.H(l).W(l)) \quad (2)$$

Our model takes as input a graph representing the dual of the road network, where each node is a road section and connected road sections are linked in the graph [Porta et al., 2006a]. The nodes have the following features:

- Section length.
- Hierarchy classes from one to six derived from attributes.
- Length of the stroke each section belongs to.
- Betweenness centrality, that is, the number of shortest paths between two pairs of nodes that use each section [Jiang and Claramunt, 2004].
- Area of both adjacent faces of the network. Sections in the city are surrounded by only small faces; in the rural areas by mostly large faces; at the border of the city by faces of very different sizes; dead ends are surrounded by one unique face.
- Sinuosity, i.e. length divided by base length. This feature informs on the role of the road, moreover, more sinuous roads may contribute to the global structure of the road but are less useful than a straight one in constructing an itinerary.

4.7 Creation of generalised road tiles

A model is trained to predict the generalised road mask from the geometry masks of roads, buildings, and rivers, and the additional information produced in Section 4.6. The FusionGAN architecture is proposed to combine two levels of geographic information from the layered representation. It has the general architecture and objective of a classical paired-conditioned generative adversarial network based on pix2pix [Isola et al., 2017], where the generator was replaced by a FuseNet [Hazırbaşı et al., 2016].

Finally, in order to add some explicit knowledge to the process [Janowicz et al., 2020, Kuhn, 2012], we tested two alternative versions of this step: 1) the addition of a connectivity loss [Courtial et al., 2022b] and 2) the addition of a topology loss [Chen et al., 2020]. Both aim at the reduction of structure alteration, i.e. disconnections of the network in the images. The connectivity loss for a generator G , and an input x computes the difference between the number of road pixel sets in the target and the prediction. The topology loss computes the gradient of the predicted image and compares it with the gradient of the input image. It was used in a style transfer problem to remove disconnections in the road network of the output map.

4.8 Map generation

This is the final step that generates the target 1:50k scale topographic map. We propose to train a CycleGAN [Zhu et al., 2017] to generate a symbolised and generalised map from the cartographic data derived in the previous

steps. CylceGAN is an unpaired GAN that learns the A to B transformation and its inverse altogether. This model is unpaired, because the model is given a set of input training images and a set of output images, but it never knows which output image corresponds to a given input image. The model learns how the output domain (i.e. the 1:50k scale map) looks like rather than the transformation to produce the output image from the input image. Unpaired models are usually more adapted than paired models in cases where paired data does not exist or where the input and target have very different aspects (that is the case for our task). Concretely, the architecture is made of two generators (F which learns transformation from A to B and G learns transformation from B to A) and two discriminators (for A and B domains). For each image in the training set (from both domains A or B), the two generators are applied consecutively (with the relevant order) and the appropriate discriminator evaluates if the prediction of the second model applied (reconstruction) is real or fake.

The input is constituted of tiles of different themes, with a level of detail adapted to the scale of the target map, stacked into the different dimensions of the input tensor, similar to the layered representation used in the previous steps. In practice, to demonstrate the interest and efficiency of this step, we need to test it with relevant input of the correct quality. Thus, we included in the input the raw results of building scale adaptation. However, for roads, predictions from the previous step were not good enough to be used in a mapping process, we rather included tiles made from a classical selection of road sections, to test the effectiveness of the map generation without bias inherited from previous steps.

5 Results

In this section, we present the results of each stage of the experiment.

5.1 Road generalisation

For a quantitative evaluation of the road importance enrichment, we compute the confusion matrix (Table 1) of selection. The GCN gives quite satisfying results on the test set. The recall is around 78% and the precision is 92%. However, the model is under-selecting the road network and kept too many undesired sections, but, it deletes very few important sections.

Prediction \ Annotation	Deleted	Kept	Total
Deleted	580	150	730
Kept	476	1770	2246
Total	1056	1920	2976

Table 1: Confusion matrix for the selection of road sections using a GCN.

We think this quantitative evaluation is not sufficient to guarantee the quality of the prediction. Indeed, structure alterations are not visible in a confusion matrix. Thus we use a visual examination of the selected network for determining if spatial relations and the structure of the network are preserved. Figure 6 presents an extract of the road network, its target generalisation at the 1:50k scale, and the predicted probability to be kept at the target scale. We observe that the model can clearly identify the main and local connections: these roads respectively appear with a significant probability of being kept (in red) or erased (in yellow). However, most roads have a medium probability. These roads are neither important nor useless, and the prediction gives unclear indications about which one to select, notably because it is also complex for human cartographers to make this choice. This close probability value can also lead to disconnections and irrelevant choices, this is why this GCN cannot be a direct selection operator but is used as an enrichment to guide a selection process that will take connection issues into account. Moreover, the predictions in urban areas contain more errors and seem to suffer from a lack of examples.

For the generalisation of road tiles at the 1:50k scale step. We observe in Figure 7 that the continuity of the network in the prediction images is the major issue. The model is able to understand the selection and selects appropriate edges but it does not preserve them entirely, and intersections are disconnected. The proposed improvements for connectivity are not really conclusive. The topology loss gives a similar quality of results for most situations and often creates more disconnections for complex intersections where simplification (and thus change in the image structure) is required. The connectivity loss slightly reduces the number of disconnections, but the numerous pixel artefacts present in the images generated with this loss reduce the legibility of the image.




Figure 6: Visual representation of road enrichment results.

5.2 Building generalisation

Figure 8 presents the results of the generalisation of individual buildings. We observe satisfying results, with realistic and generalised buildings. However, the comparative evaluation of such buildings between the basic GAN and the addition of the geometry consistency loss is visually difficult: some buildings are better with the loss, others are worse, and the improvement of geometry consistency loss in terms of shape regularity is not visually significant.

Thus, we propose a quantitative evaluation of building shape regularity that includes a vector conversion. To vectorize the buildings from the tiles, we use the GDAL function 'polygonize' which creates vector polygons for all connected regions of pixels in the raster that share a common pixel value. Then we apply a slight Douglas-Peucker filter (threshold = 2m) to correct the notched shape due to the pixel limits. This simplification does not generalise the building and globally preserves the generalisation problems of the building (when there is a problem), but it smooths the pixel border. Finally, we eliminated the truncated buildings at the border of the image, which are not relevant for shape evaluation. Then, we propose and implemented using Python, several measures to compare the quality of the prediction:

- **Granularity:** the length of the smallest building edge for each building (the value should be as large as possible).
- **Number of edges:** the more edges a building has, the less credible its shape is. In fact, even a very complex real building often has very few edges once it is generalised at the 1:50k scale (the smallest the value is, the best the prediction is).
- **Rectangularity coefficient:** the area ratio between the building geometry and its minimal bounding box. This measure is only relevant in quasi-rectangular-shaped buildings, and in the case of more complex buildings, this measure is not used (the biggest the value is, the best the prediction is).
- **Deviation to the most credible angle**[Touya and Lokhat, 2016]: it sums up how far each angle is from a flat or right angle, for each pair of the consecutive edge of the building (the smallest the value is, the best the prediction is).

We compared the distribution of these measures for both buildings generalised with and without the geometry consistency with the target building generalisation in Table 2. The measures give quite similar results: none of the models has a significantly different distribution. The quantitative evaluation could not prove the effectiveness of the geometry consistency loss for building generalisation. The comparison with reference confirms that the building shapes have the expected granularity, and complexity (number of edges), but they suffer from irregular shapes (rectangular coefficient is worse than the target).

Input		Target	Prediction
Main information	Roads and buildings		

Figure 7: Prediction of the proposed model for the road adaptation at 1:50k.

5.3 Block covering generation

Finally, Figure 9 presents the results for the covering of the city centre blocks. We observe that the results are globally bad: the model covers too many portions of space, and it does not learn to find the city centre but just creates a kind of built-up area region. We think that the tile representation does not cover a large enough area to allow a block analysis. Larger tiles or the use of attention-based mechanisms could help to solve this block-covering task [Liu et al., 2021b].

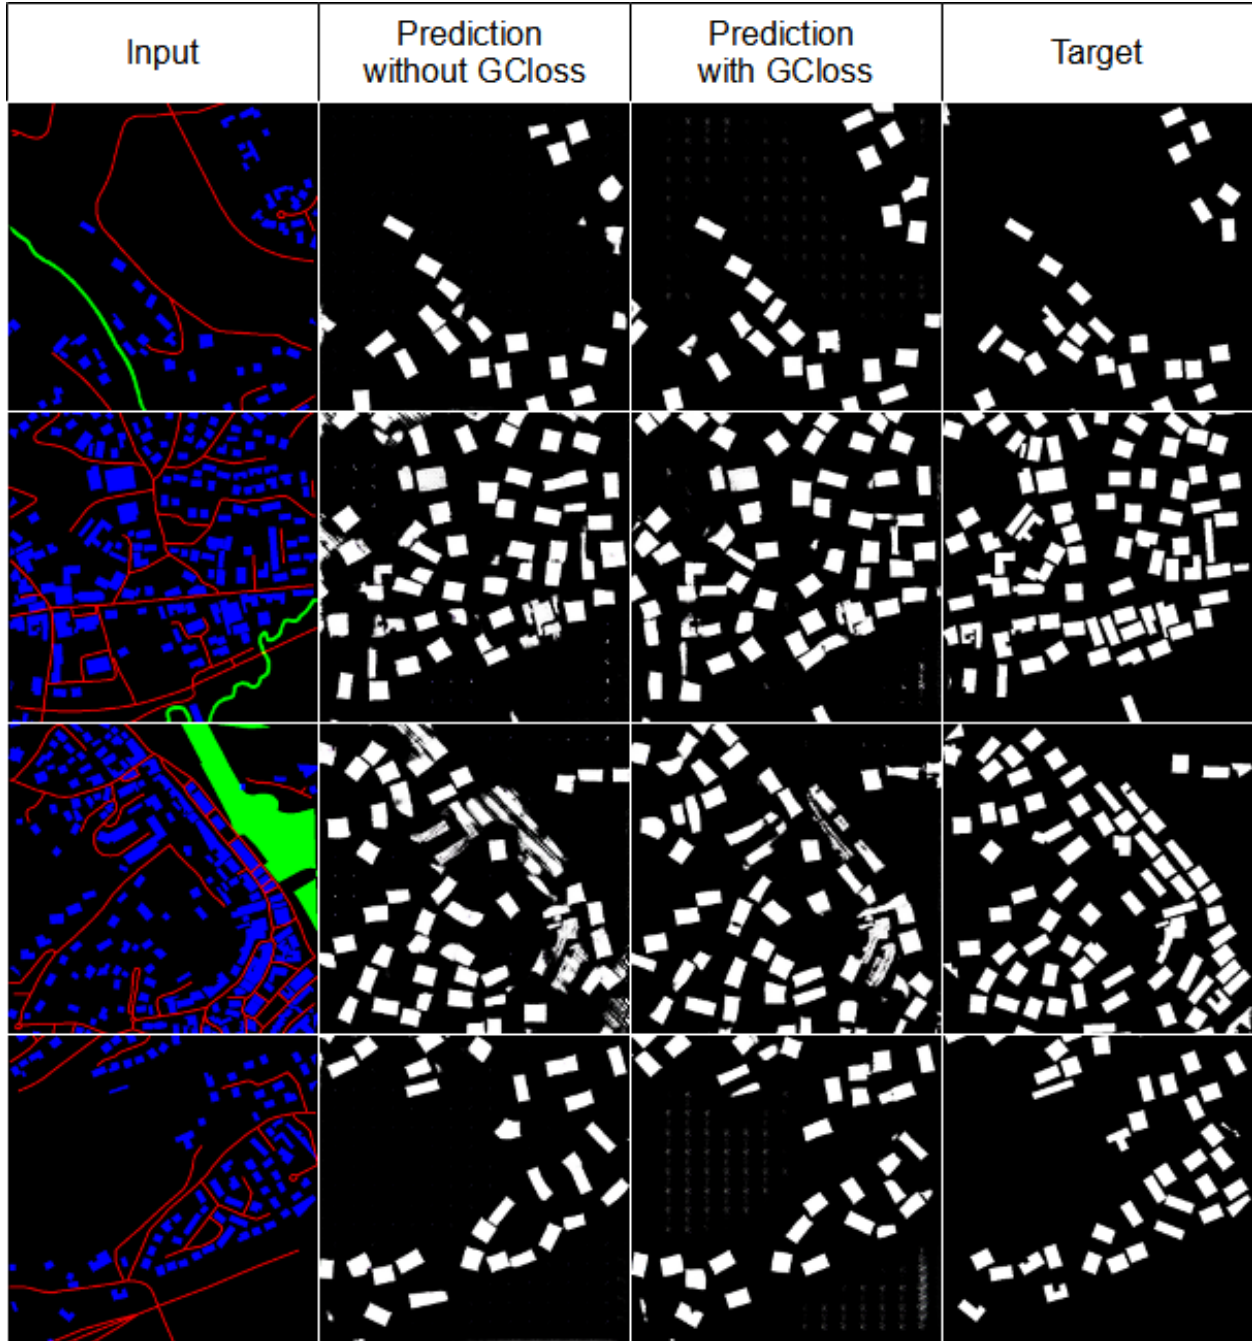


Figure 8: Prediction of proposed models for the building adaptation at 1:50k.

5.4 Map generation

The last step of the workflow is the generation of legible symbolised map tiles without overlapping data at the target scale. Figure 10 presents an extract of the results of this generation. The outputs have a satisfying style and are credible as topographic maps. We observe that the shape of buildings is sometimes modified mostly in two cases: 1) when the separation between buildings or with roads is not respected, the buildings are refined or enlarged to create separation or an amalgamation; 2) when a building or an artefact from the previous model is too small, it is enlarged.

Some interesting elements in the prediction are illustrated in Figure 11. The amalgamation of close buildings is illustrated in purple. In green (and pink), we observe the displacement and distortion of buildings to avoid overlaps with

Table 2: Statistics on building regularity measures for prediction of the model trained without geometry consistency loss (NoGC), with geometry consistency loss (GC) and for the target semi-automatic generalisation.

Constraint	Model	min	max	mean	med
Granularity (metre)	NoGC	2	38	8	4.5
	GC	2	32	7.7	4.5
	Target	2	27	7.6	4.5
Number of edges	NoGC	4	70	7.5	6
	GC	4	70	7.6	6
	Target	4	59	7.6	7
Rectangular coefficient	NoGC	0.37	1	0.79	0.81
	GC	0.33	1	0.79	0.81
	Target	0.41	1	0.81	0.86
Deviation coefficient (degree)	NoGC	0	45	15.7	14.0
	GC	0	45	15.6	14.0
	Target	0	45	17.1	15.2




Figure 9: Prediction block covering generation.




Figure 10: Some input and prediction images for the map generation step.

road symbols (and with neighbouring buildings). Finally, some drawbacks remain, highlighted in Figure 12. Some undesired overlaps remain between roads and hydrography (circled in red) and with buildings (circled in blue). These errors seem to be due to the inability of the model to sometimes decide between amalgamation and displacement for resolving the conflict, and the model performs both which results in shapes that are not realistic. These results are also the output of the whole workflow and are later discussed and compared with state-of-the-art methods in Section 6.2.

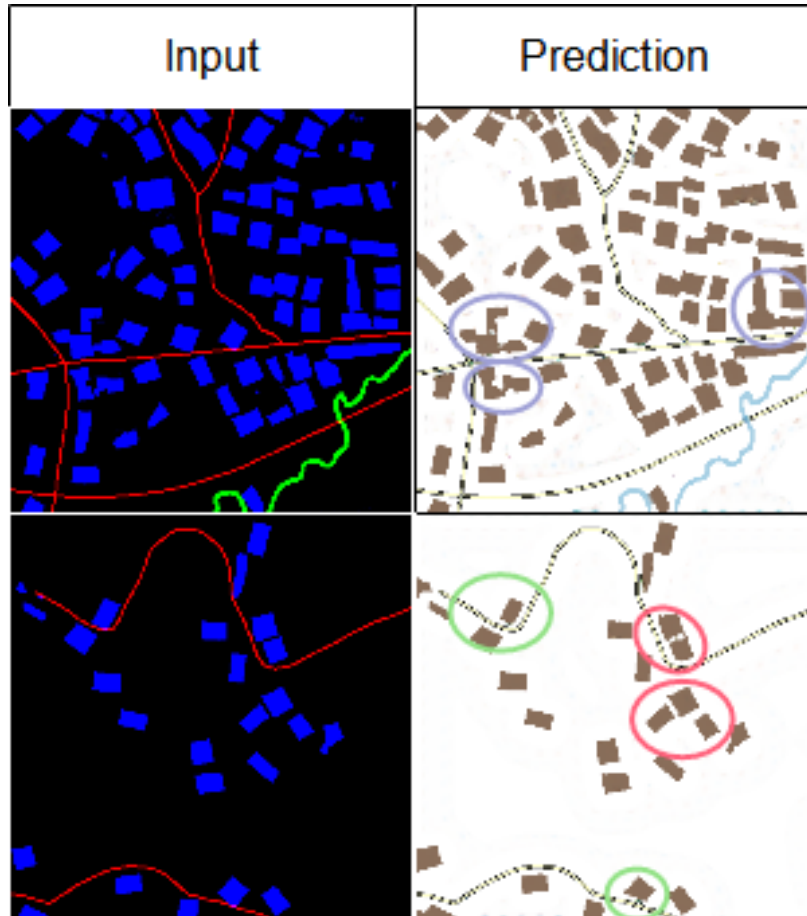


Figure 11: Some predictions of our model with circles showing some specific strong points: in purple, two close buildings are generalised with an amalgamation strategy; in green, some buildings are displaced to a minimum distance with the road; in pink, two close buildings are displaced to ensure a minimum distance with other buildings.

6 Discussion

6.1 Advantages and limitations of the workflow

Our experiment tests the effectiveness of a workflow that combines several deep neural networks for the generation of a complete generalised topographic map. With a workflow organisation, the objective of each deep neural network is more clearly defined which has shown some advantages over the usage of a holistic end-to-end model. These are discussed, as well as limitations in the following subsections.

6.1.1 The workflow allows for an independent formulation of each sub-task

Map generalisation is a complex task, thus finding a relevant formulation and data representation to learn this task is also complex. The input must contain (at least implicitly) all necessary information for generalisation.

For example, the use of additional information enables the learning of road selection, but it globally makes the input more complex and is useless for the generalisation of other themes [Courtial et al., 2022c]. Consequently, when applied to a whole map it degraded the global quality of prediction. On the contrary in such a workflow, we can determine the appropriate formulation, and design an adapted training set for each task, and the additional information is only used when it is beneficial. Thus, each model is trained with an adapted input in terms of content, representation, level of detail, size of context, etc.

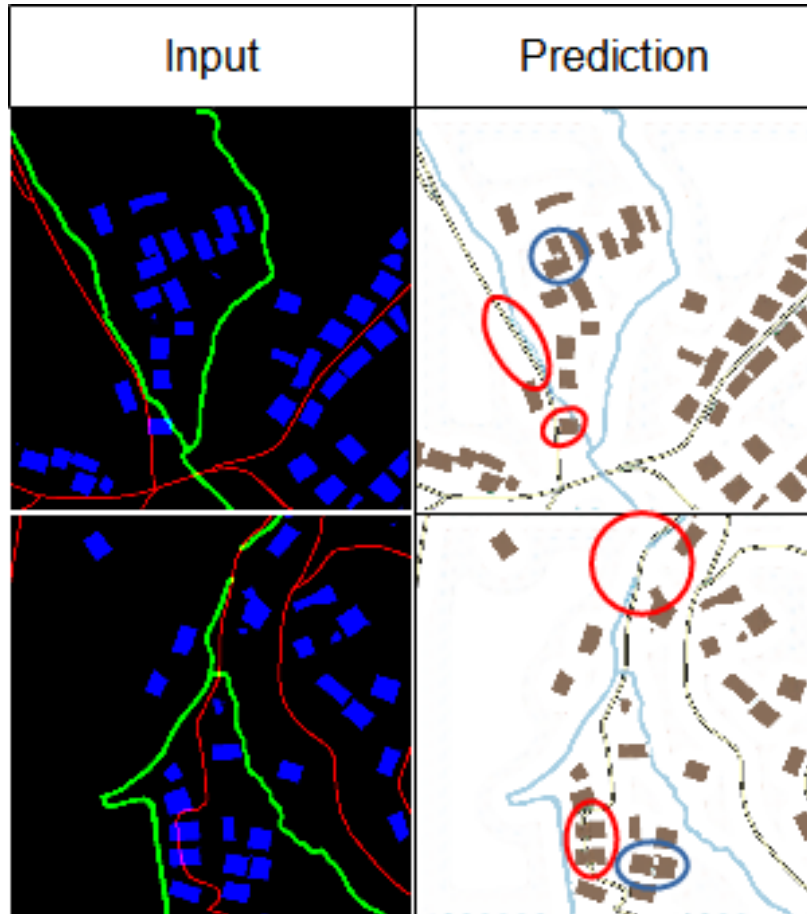


Figure 12: Some predictions of our model with circles showing some flaws: in red, the displacement to ensure the minimal distance between the road and the river is not performed; in blue, both displacement and amalgamation operations are performed instead of one, which leads to an unrealistic building shape.

Another example is the choice of the best deep learning architecture for a task. Architectures working on image tiles such as pix2pix or CycleGAN seem to be adapted for a graphic generalisation, but not for semantic enrichment or network generalisation.

Lastly, the block-covering task illustrates this advantage. The scarcity of tiles showing block covering in the training set made it impossible to properly learn this task with the general set designed for buildings. By separating tasks, we could provide additional examples from other areas without creating a new complete training dataset for the whole map generation in these areas. Then, with the workflow organisation, we could also suggest changing the tile size for this task without changing the global tile size for other tasks. The tile size depends on the level of detail and context which may vary for different generalisation sub-tasks.

6.1.2 The workflow allows for simpler evaluation and post-processing

Secondly, the workflow organisation allows for a simpler evaluation for two reasons. First, if the task and its objective are clearly defined, it is more simple to estimate if they are satisfied. Indeed, there are much more constraints to assess the quality of a complete topographic map, than for the assessment of a single operator. And the larger the number of constraints is, the more complex a global estimation of quality is [Touya, 2012]. Secondly, the estimation of intra and inter-theme conflicts is separated, which makes their automatic evaluation simpler. For example, the evaluation of building shape is impossible if the building and road overlap, because the road might hide some parts of the building, but with the workflow organisation we could evaluate first the shape of buildings and then if they overlap with other themes. In the same manner, the identification of remaining errors and corrections with post-processes is simpler when tasks are separated (e.g. elimination of pixel artefacts, correction of disconnections, etc.).

6.1.3 The workflow is more pragmatic

Finally, the workflow organisation allows us to train and improve each model independently which is more pragmatic [Scheider and Richter, 2023]. Indeed, the different sub-tasks of map generalisation have different characteristics and complexity. Thus, they may require different models and training, in particular the most adapted model architectures, losses, optimization methods and the number of epochs can vary for each sub-task. For instance, during our experiments, the map generation takes 300 epochs, while building generalisation only needs between 150 and 200 epochs to reach an optimal state. In a similar way, a quick experiment on loss functions has shown that a geometry consistency loss could improve the shape of buildings but decrease the quality of road generalisation.

6.1.4 Limitations

Despite these advantages, the proposed workflow has some limitations.

- It requires more time and storage capacity than a unique model;
- The propagation of errors is more probable. For instance, in our experiment, the road selection is not performed correctly during the scale adaption step and disconnections occur. If we would have used this prediction in the map generation step, not only the prediction would have included the same disconnections, but also the model would have learned to make map-looking images from this erroneous data and have interpreted disconnected roads as dead ends. This may lead to bad placement of other themes relative to the erroneous roads.

6.2 Comparative evaluation

In Figure 13, we compare the results obtained with the DeepMapScaler workflow, with the results of the holistic unique model trained to generate end-to-end generalised topographic maps proposed by [Courtial et al., 2021b], and with a traditional semi-automatic map generalisation process (target).

The workflow predictions are particularly satisfying: the road selection stays connected and the proximity constraint between roads and buildings is satisfied. The realism of the prediction is also improved compared to the holistic model, the prediction with the workflow contains less noise and blurred parts since the block covering and building shape are not learned together. The workflow model also has less oversimplifies the shape of the complex buildings. We also observe that the workflow can also learn to reproduce generalisation operators that are not learnt by the holistic model such as displacement and amalgamation. Thus, the prediction better preserves the structure of buildings (i.e. on the top image the square sparse building structure is preserved while the unique model generalised it in the building bar).

Then, compared with the semi-automatic generalisation, the quality of building generalisation is more homogeneous: with the semi-automatic method there is no unrealistic shape but some buildings remain too detailed or too small, because the algorithm used failed, while with the generalisation by deep learning, each building contains small imperfections but very few building have a very bad generalisation.

6.3 Improving the workflow

In this section, we discuss how the workflow can be improved in the future. We discuss the improvement of the different proposed steps, and then we discuss the improvement of the transition between steps.

6.3.1 Improving the models responsible for each step of the workflow

The road scale adaptation is the step with the worst results, in particular, the generation of generalised road maps includes too many disconnections to be considered satisfying. To improve this point, we proposed a loss function which should discourage disconnections. However, this strategy is not significantly successful. We believe that the formulation of this step must be changed, and the selection of the roads should not be performed with an image-based architecture. The image-based architecture should only be used for a graphic generalisation (if needed). We also think that the addition of a step for the simplification of complex road structures, such as highway interchanges or complex intersections, is required.

The block-covering model also needs to be improved. We believe that both the training set and the architecture proposed are not adapted for this task. We already suggested producing bigger tiles for this task. We could also suggest another approach with a graph where each node describes an urban block and an edge describes adjacency relations between blocks. The model is then trained to predict a label (covered or not) for each block according to the characteristics (stored as node attributes) and proximity of blocks.




Figure 13: Comparison of predictions made by a workflow and by a unique model, as well as the ground truth obtained with traditional generalisation techniques.

For buildings, the prediction is global and of good quality, as the slightly fuzzy shapes can be corrected by fine-tuning the model or by post-processing the output. There are also some remaining buildings with too much granularity. This problem may be due to the few errors remaining in the training set that make the model learn that it is sometimes possible to leave buildings with high granularity in the geometry.

6.3.2 Transition between steps

In this article, we trained each step independently, with the input of intermediary steps derived from the output of previous steps. In this section, we discuss this choice.

In the current workflow, transitions are limited to rasterization and graph creation. Our data is initially composed of vector cartographic data while the output and the input of each model is either a graph or an image. This choice allows us to visualise and analyse each step independently. However, it forces us to encode and decode multiple times the same information (e.g. the inputs of the map generation step are decoded from the last steps of the scale adaptation, and the model begins by encoding them). In particular, when the step aims to provide information about an image or an object (e.g. the data enrichment steps), we think that this redundancy is not useful. Our idea is to use an encoding-only model for this step, that provides an embedding representation of attributes or context information to the next model directly. Such an approach skips the decoding for human understanding, which may be biased and degrade information (see Figure 14). Moreover, such transitions would chain processes based on any kind of object representation and would give more freedom in the steps of the workflow.




Figure 14: Illustration of the alternative transition between steps based on embedding.

7 Conclusion

We proposed and tested a workflow for generalised map generation with deep learning techniques, called DeepMapScaler. The proposed workflow is organised in two steps, the data adaptation for the desired scale (that is made theme by theme) and the map generation. These two steps do not exactly correspond to classical steps in traditional map generalisation: the data adaptation includes data enrichment, object selection and some graphic generalisation that relates to intra-theme constraints, while the map generation deals with solving inter-theme conflicts and conflicts linked to symbolisation.

The comparison with literature results shows that the workflow unlocks the learning of generalisation operators that were important challenges in the literature on generalised map generation using deep learning: displacement, road selection and amalgamation. However, some steps (the generalisation of roads and the covering of city centre blocks), and the transition between steps still require to be improved. With the results obtained with DeepMapScaler, we believe that the creation of a workflow that combines deep learning (and others) models is more reasonable than the design of a unique map generalisation model for now.

Data availability statement

The data that support the findings of this study are available at the link <https://doi.org/10.5281/zenodo.7957430>.

References

- [Aslan et al., 2012] Aslan, S., Bildirici, I., Simav, O., and Cetinkaya, B. (2012). An incremental displacement approach applied to building objects in topographic mapping. In *13th ICA workshop on generalisation and multiple representation*, Istanbul.
- [Barrault et al., 2001] Barrault, M., Regnauld, N., Duchêne, C., Haire, K., Baeijs, C., Demazeau, Y., Hardy, P., Mackaness, W. A., Ruas, A., and Weibel, R. (2001). Integrating multi-agent, object-oriented, and algorithmic techniques for improved automated map generalisation. In *20th International Cartographic Conference*, volume 3, pages 2110–2116, Beijing, China. ICA.
- [Beard, 1991] Beard, M. K. (1991). Constraints on rule formation. In Buttenfield, B. P. and McMaster, R. B., editors, *Map Generalization: Making Rules for Knowledge Representation*, pages 121–135. Longman, London.
- [Benz and Weibel, 2014] Benz, S. A. and Weibel, R. (2014). Road network selection for medium scales using an extended stroke-mesh combination algorithm. *Cartography and Geographic Information Science*, 41(4):323–339. Publisher: Taylor & Francis.
- [Brassel and Weibel, 1988] Brassel, K. E. and Weibel, R. (1988). A review and conceptual framework of automated map generalization. *International Journal of Geographical Information Systems*, 2(3):229–244. Publisher: Taylor & Francis.
- [Burghardt and Cecconi, 2007] Burghardt, D. and Cecconi, A. (2007). Mesh simplification for building typification. *International Journal of Geographical Information Science*, 21(3):283–298.
- [Chen et al., 2021] Chen, W., Wu, A., and Biljecki, F. (2021). Classification of urban morphology with deep learning: Application on urban vitality. *Computers, Environment and Urban Systems*, 90:13.
- [Chen et al., 2020] Chen, X., Chen, S., Xu, T., Yin, B., Peng, J., Mei, X., and Li, H. (2020). SMAPGAN: Generative Adversarial Network-Based Semisupervised Styled Map Tile Generation Method. *IEEE Transactions on Geoscience and Remote Sensing*, pages 1–19.
- [Courtial et al., 2020] Courtial, A., El Ayedi, A., Touya, G., and Zhang, X. (2020). Exploring the Potential of Deep Learning Segmentation for Mountain Roads Generalisation. *ISPRS International Journal of Geo-Information*, 9(5):338. Number: 5 Publisher: Multidisciplinary Digital Publishing Institute.

- [Courtial et al., 2021a] Courtial, A., Touya, G., and Zhang, X. (2021a). Can Graph Convolution Networks Learn Spatial Relations? *Abstracts of the ICA*, 3:1–2.
- [Courtial et al., 2021b] Courtial, A., Touya, G., and Zhang, X. (2021b). Generative Adversarial Networks to Generalise Urban Areas in Topographic Maps. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume XLIII-B4-2021, pages 15–22. Copernicus GmbH. ISSN: 1682-1750.
- [Courtial et al., 2022a] Courtial, A., Touya, G., and Zhang, X. (2022a). Constraint-Based Evaluation of Map Images Generalized by Deep Learning. *Journal of Geovisualization and Spatial Analysis*, 6.
- [Courtial et al., 2022b] Courtial, A., Touya, G., and Zhang, X. (2022b). Deriving map images of generalised mountain roads with generative adversarial networks. *International Journal of Geographical Information Science*, 0(0):1–30. Publisher: Taylor & Francis.
- [Courtial et al., 2022c] Courtial, A., Touya, G., and Zhang, X. (2022c). Representing Vector Geographic Information As a Tensor for Deep Learning Based Map Generalisation. In Parseliunas, E., Mansourian, A., Partsinevelos, P., and Suziedelyte-Visockiene, J., editors, *AGILE 2022*, volume 3, page 32, Vilnius, Lithuania. Copernicus Publications.
- [Du et al., 2021] Du, J., Wu, F., Xing, R., Gong, X., and Yu, L. (2021). Segmentation and sampling method for complex polyline generalization based on a generative adversarial network. *Geocarto International*, pages 1–23. Publisher: Taylor & Francis.
- [Du et al., 2022] Du, J., Wu, F., Yin, J., Liu, C., and Gong, X. (2022). Polyline simplification based on the artificial neural network with constraints of generalization knowledge. *Cartography and Geographic Information Science*, 49(4):313–337. Publisher: Taylor & Francis.
- [Duchêne et al., 2018] 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.
- [Feng et al., 2019] Feng, Y., Thiemann, F., and Sester, M. (2019). Learning cartographic building generalization with deep convolutional neural networks. *International Journal of Geo-Information*, 8(6).
- [Fu et al., 2019] Fu, H., Gong, M., Wang, C., Batmanghelich, K., Zhang, K., and Tao, D. (2019). Geometry-Consistent Generative Adversarial Networks for One-Sided Unsupervised Domain Mapping. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2422–2431, Long Beach, CA, USA. IEEE.
- [Garcia-Balboa and Ariza-López, 2008] Garcia-Balboa, J. L. and Ariza-López, F. J. (2008). Generalization-oriented road line classification by means of an artificial neural network. *Geoinformatica*.
- [Hazırbaş et al., 2016] Hazırbaş, C., Ma, L., Domokos, C., and Cremers, D. (2016). FuseNet: Incorporating Depth into Semantic Segmentation via Fusion-Based CNN Architecture.
- [Hu et al., 2022] Hu, Y., Liu, C., Li, Z., Xu, J., Han, Z., and Guo, J. (2022). Few-Shot Building Footprint Shape Classification with Relation Network. *ISPRS International Journal of Geo-Information*, 11(5):311.
- [Iddianozie and Mcardle, 2021] Iddianozie, C. and Mcardle, G. (2021). Transferable Graph Neural Networks for Inferring Road Type Attributes in Street Networks. *IEEE Access*, 9:158331–158339. Conference Name: IEEE Access.
- [Isola et al., 2017] 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)*, pages 5967–5976. ISSN: 1063-6919.
- [Janowicz et al., 2020] Janowicz, K., Gao, S., McKenzie, G., Hu, Y., and Bhaduri, B. (2020). GeoAI: spatially explicit artificial intelligence techniques for geographic knowledge discovery and beyond. *International Journal of Geographical Information Science*, 34(4):625–636.
- [Jiang and Claramunt, 2004] Jiang, B. and Claramunt, C. (2004). A Structural Approach to the Model Generalization of an Urban Street Network. *Geoinformatica*, 8(2):157–171.
- [Kang et al., 2019] Kang, Y., Gao, S., and Roth, R. E. (2019). Transferring Multiscale Map Styles Using Generative Adversarial Networks. *International Journal of Cartography*, 5(2-3):115–141.
- [Kang et al., 2020] Kang, Y., Rao, J., Wang, W., Peng, B., Gao, S., and Zhang, F. (2020). Towards Cartographic Knowledge Encoding with Deep Learning. In *Autocarto*, page 6, Redlands, California.
- [Kuhn, 2012] Kuhn, W. (2012). Core concepts of spatial information for transdisciplinary research. *International Journal of Geographical Information Science*, 26(12):2267–2276. Publisher: Taylor & Francis.
- [Li et al., 2020] Li, C., Zhang, H., Wu, P., Yin, Y., and Liu, S. (2020). A complex junction recognition method based on GoogLeNet model. *Transactions in GIS*, 24.

- [Liu et al., 2021a] Liu, C., Hu, Y., Li, Z., Xu, J., Han, Z., and Guo, J. (2021a). TriangleConv: A Deep Point Convolutional Network for Recognizing Building Shapes in Map Space. *ISPRS International Journal of Geo-Information*, 10(10):687. Number: 10 Publisher: Multidisciplinary Digital Publishing Institute.
- [Liu et al., 2021b] Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. (2021b). Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. arXiv:2103.14030 [cs].
- [Ma et al., 2022] Ma, L., Seipel, S., Brandt, S. A., and Ma, D. (2022). A New Graph-Based Fractality Index to Characterize Complexity of Urban Form. *ISPRS International Journal of Geo-Information*, 11(5):287.
- [Mackaness and Beard, 1993] Mackaness, W. and Beard, K. (1993). Use of Graph Theory to Support Map Generalization. *Cartography and Geographic Information Science*, 20:210–221.
- [Mackaness et al., 2014] Mackaness, W., Burghardt, D., and Duchêne, C. (2014). Map Generalisation: Fundamental to the Modelling and Understanding of Geographic Space. In Burghardt, D., Duchêne, C., and Mackaness, W., editors, *Abstracting Geographic Information in a Data Rich World: Methodologies and Applications of Map Generalisation*, Lecture Notes in Geoinformation and Cartography, pages 1–15. Springer International Publishing, Cham.
- [Mackaness and Edwards, 2002] Mackaness, W. and Edwards, G. (2002). The Importance of Modelling Pattern and Structure in Automated Map Generalisation. Join ISPRS/ICA workshop on multi-scale representations of spatial data. In *Proceedings of the Joint ISPRS/ICA Workshop on Multi-Scale Representations of Spatial Data*, pages 7–8.
- [Mustiere, 1998] Mustiere, S. (1998). GALBE : Adaptative generalization. the need for an adaptative process for automated generalisation an exemple on roads. *Procesding of 1st GIS PlaNet conference*.
- [Porta et al., 2006a] Porta, S., Crucitti, P., and Latora, V. (2006a). The network analysis of urban streets: A dual approach. *Physica A: Statistical Mechanics and its Applications*, 369(2):853–866.
- [Porta et al., 2006b] Porta, S., Crucitti, P., and Latora, V. (2006b). The Network Analysis of Urban Streets: A Primal Approach. *Environment and Planning B: Planning and Design*, 33(5):705–725.
- [Scheider and Richter, 2023] Scheider, S. and Richter, K.-F. (2023). Pragmatic GeoAI: Geographic Information as Externalized Practice. *KI - Künstliche Intelligenz*.
- [Sester, 2000] Sester, M. (2000). Knowledge acquisition for the automatic interpretation of spatial data. *International Journal of Geographical Information Science*, 14(1):1–24.
- [Shea and McMaster, 1989] Shea, K. S. and McMaster, R. B. (1989). When And How To Generalize.
- [Stanislawski et al., 2014] Stanislawski, L. V., Bittenfield, B. P., Bereuter, P., Savino, S., and Brewer, C. A. (2014). Generalisation Operators. In Burghardt, D., Duchêne, C., and Mackaness, W., editors, *Abstracting Geographic Information in a Data Rich World*, Lecture Notes in Geoinformation and Cartography, pages 157–195. Springer International Publishing.
- [Steiniger et al., 2008] 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.
- [Thomson, 2006] Thomson, R. C. (2006). The 'stroke' concept in geographic network; generalization and analysis. In Riedl, A., Kainz, W., and Elmes, G. A., editors, *Progress in Spatial Data Handling 12th International Symposium on Spatial Data Handling*, pages 681–697.
- [Thomson and Richardson, 1995] Thomson, R. C. and Richardson, D. (1995). A Graph Theory Approach to Road Network Generalisation. In *17th International Cartographic Conference*, pages 1871–1880. ICA.
- [Touya, 2010] Touya, G. (2010). A Road Network Selection Process Based on Data Enrichment and Structure Detection. *Transactions in GIS*, 14(5):595–614.
- [Touya, 2012] Touya, G. (2012). Social Welfare to Assess the Global Legibility of a Generalized Map. In Xiao, N., Kwan, M.-P., Goodchild, M. F., and Shekhar, S., editors, *Geographic Information Science*, Lecture Notes in Computer Science, pages 198–211, Berlin, Heidelberg. Springer.
- [Touya, 2021] Touya, G. (2021). Multi-Criteria Geographic Analysis for Automated Cartographic Generalization. *Cartographic Journal*. Publisher: Maney Publishing.
- [Touya and Courtial, 2021] Touya, G. and Courtial, A. (2021). BasqueRoads: A Benchmark for Road Network Selection. volume 4, 2021. Copernicus Publications.
- [Touya and Dumont, 2017] 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 and Lokhat, 2016] Touya, G. and Lokhat, I. (2016). Enhancing building footprints with squaring operations. *Journal of Spatial Information Science*, (13).

- [Touya and Lokhat, 2020] Touya, G. and Lokhat, I. (2020). Deep Learning for Enrichment of Vector Spatial Databases: Application to Highway Interchange. *ACM Transactions on Spatial Algorithms and Systems*, 6(3):21. Publisher: ACM.
- [Touya et al., 2019a] Touya, G., Lokhat, I., and Duchêne, C. (2019a). CartAGen: an Open Source Research Platform for Map Generalization. *Proceedings of the ICA*, 2:1–9.
- [Touya et al., 2019b] Touya, G., Zhang, X., and Lokhat, I. (2019b). Is deep learning the new agent for map generalization? *International Journal of Cartography*, 5(2-3):142–157.
- [Weiss and Weibel, 2014] Weiss, R. and Weibel, R. (2014). Road network selection for small-scale maps using an improved centrality-based algorithm. *Journal of Spatial Information Science*, 9(9):71–99.
- [Yan et al., 2020a] Yan, X., Ai, T., Yang, M., and Tong, X. (2020a). Graph convolutional autoencoder model for the shape coding and cognition of buildings in maps. *International Journal of Geographical Information Science*, 0(0):1–23. Publisher: Taylor & Francis.
- [Yan et al., 2020b] Yan, X., Ai, T., Yang, M., Tong, X., and Liu, Q. (2020b). A graph deep learning approach for urban building grouping. *Geocarto International*, 0(0):1–24. Publisher: Taylor & Francis.
- [Yang et al., 2022a] Yang, M., Jiang, C., Yan, X., Ai, T., Cao, M., and Chen, W. (2022a). Detecting interchanges in road networks using a graph convolutional network approach. *International Journal of Geographical Information Science*, 0(0):1–21. Publisher: Taylor & Francis.
- [Yang et al., 2022b] Yang, M., Kong, B., Dang, R., and Yan, X. (2022b). Classifying urban functional regions by integrating buildings and points-of-interest using a stacking ensemble method. *International Journal of Applied Earth Observation and Geoinformation*, 108:102753.
- [Yu and Chen, 2022] Yu, W. and Chen, Y. (2022). Data-driven polyline simplification using a stacked autoencoder-based deep neural network. *Transactions in GIS*, 26(5):2302–2325.
- [Zhao et al., 2020] Zhao, R., Ai, T., Yu, W., He, Y., and Shen, Y. (2020). Recognition of building group patterns using graph convolutional network. *Cartography and Geographic Information Science*, 47(5):400–417.
- [Zhou and Li, 2017] 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.
- [Zhou et al., 2022] Zhou, Z., Fu, C., and Weibel, R. (2022). Building simplification of vector maps using graph convolutional neural networks. *Abstracts of the ICA*, 5:1–2. Publisher: Copernicus GmbH.
- [Zhu et al., 2017] 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)*, pages 2242–2251, Venice. IEEE.