INFERRING THE SCALE AND CONTENT OF A MAP USING DEEP LEARNING

G. Touya[a, *], F. Brisebard[b], F. Quinton[b], A. Courtial[a]

[a] LASTIG, Univ Gustave Eiffel, ENSG, IGN, F-94160 Saint-Mandé, France - (guillaume.touya, azelle.courtial)@ign.fr
[b] Univ Gustave Eiffel, ENSG, IGN, F-77420 Champs-sur-Marne, France – (florentin.brisebard, felix.quinton)@ensg.eu

Commission IV, WG IV/1

KEY WORDS: Cartography, Scale, Deep Learning, Classification, Visually Impaired, Tactile Maps

ABSTRACT:

Visually impaired people cannot use classical maps but can learn to use tactile relief maps. These tactile maps are crucial at school to learn geography and history as well as the other students. They are produced manually by professional transcriptors in a very long and costly process. A platform able to generate tactile maps from maps scanned from geography textbooks could be extremely useful to these transcriptors, to fasten their production. As a first step towards such a platform, this paper proposes a method to infer the scale and the content of the map from its image. We used convolutional neural networks trained with a few hundred maps from French geography textbooks, and the results show promising results to infer labels about the content of the map (e.g. "there are roads, cities and administrative boundaries"), and to infer the extent of the map (e.g. a map of France or of Europe).

1. INTRODUCTION

Visually impaired people cannot use classical maps but can learn to use tactile relief maps, in order to understand the geography of their place of residence, hometown, region, country, or of the world. These tactile maps are useful to teach to visually impaired people how to make their daily travel autonomously. But they are also crucial at school to learn geography and history as well as the other students. Today, these maps are manually produced by professional transcriptors, and the process can be very long. Geography teachers use many maps from their textbooks in class, and when they have visually impaired students, they need tactile equivalents, and even when they have access to a transcriptor, which is not always the case, it would take too long to create all these equivalents. This is why it would be useful to apply advanced cartography techniques to accelerate tactile mapmaking (Lobben, 2015), or even to automate it (Touya et al., 2019, Wabiński and Mościcka, 2019).

In this context, our long term goal is an automated online platform where geography teachers are able to send an image (e.g. scanned from a textbook) of a map, and then receive the tactile map, as a file ready to be 3D printed, which is as close as possible as the original map. Another use of the platform would be to provide a first draft of the map to a professional transcriptor, which would just have to finalize and polish the map, saving a lot of time in the tactile map production process. In order to provide a tactile map from a map image, several steps are necessary: (1) infer the map characteristics from the image (the scale, the content, the style, the geographical extent); (2) collect the geographic data necessary to draw this map; (3) simplify and generalise this geographic data to adapt it to tactile map reading (Touya et al., 2019); (4) transform the map into a digital model able to be printed with one of the existing techniques (Lobben, 2015, Brock et al., 2015) As a first step towards this online platform, we propose in this paper a method to infer the scale and the content of the map image, based on deep learning techniques.

The paper is structured as follows. The second section presents past research related to the adaptation of maps for visually impaired people. Section 3 describes the method proposed to infer

the scale and the extent of the maps. Section 4 describes the method proposed to infer the type of features contained in the map. Section 5 discusses the methods and their results. Finally, Section 6 draws some conclusions and presents future research towards automated tactile cartography for visually impaired people.

2. ADAPTING MAPS FOR VISUALLY IMPAIRED STUDENTS

As touch perception is far less precise than vision, the specifications for good tactile maps command maps that are extemely simplified. Figure 1 shows for instance a map of Australia for people with normal vision, and a version for visually impaired people, where the content is reduced to minimum and the geometry of the boundaries is also simplified and even schematised (Mackaness and Reimer, 2014). There is a long history of studies to understand the limits of tactile graphics (Edman, 1992), even in the context of maps (Rowell and Ungar, 2003). Similarities between cartographic design and tactile map design have been highlighted with a translation of Bertin's visual variables into tactile variable for maps (Vasconcellos, 1996).
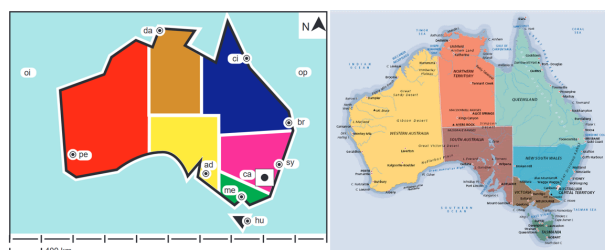

Figure 1: Two administrative maps of Australia: for visually impaired people on the left (©INSHEA-SDADV-2019-2020); for people with normal vision on the right (source: http://maps.unomaha.edu).

The optimal specifications for a tactile map vary according to the type of visual impairement (Brock et al., 2013), but also depending on the printing techniques (e.g. relief embossing, tactile screen with audio descriptions, or 3D printing). Tactile maps can

*Corresponding author

also be interactive, as well as current web maps, and it was shown that such interactivity increases the chances of understanding the geography behind the map for visual impaired people (Brock et al., 2015).

Now, considering our four-step process to automatically generate tactile maps from images of maps for geography and history lessons, some of these steps have been approached in recent research projects. For step (2), data collection, the main task is to select the appropriate layers of vector data from the existing geographical databases. But some information necessary for visually impaired people might be lacking in geographic databases collected for topographic mapping (Touya et al., 2019). For instance, pavements and zebra crossing are rarely collected in such databases but are necessary when maps of daily travels are created. In this case, it is possible to generate this missing information from high resolution aerial images with deep learning techniques (Fillières-Riveau et al., 2019).

At step (3), the vector geographic data collected at the previous step needs to be simplified to meet the drastic specifications of tactile maps. Many tools already exist for this task, many of them based on OpenStreetMap data, a quite complete list of such applications can be found in (Wabiński and Mościcka, 2019). But as illustrated by the Mapy.cz application (Červenka et al., 2016), the maps remain often too complex to be fully usable for visually impaired people, or even as base maps for transcriptors. The use and adaptation of map generalisation techniques, firstly designed for scale reduction in topographic maps seems really promising to achieve tactile maps that are simple enough (Štampach and Mulíčková, 2016, Touya et al., 2019, Wabiński and Mościcka, 2019).

But, before being able to transform geographic data into tactile maps for visually impaired people, the first step is to infer the characteristics of the map, from which the specifications of the map can later be derived: scale, projection, map content, geographic extent, style and symbols, etc. If we consider that the input map is an image, as it is the most practical way to interact with teachers, inferring map characteristics can be seen as an image labeling and classification problem. Deep learning techniques, and particularly convolutional neural networks (CNN) (LeCun et al., 2015) are now obvious solutions for such image classification problems. However, there are very few applications of image classification techniques to map images. The use of CNNs was proven useful to classify geographic images (e.g. an aerial image, a topographic map, a landcover map, or a 3D scene) (Zhou et al., 2018). In their model to transfer the style of Google Maps to OpenStreetMap data with generative adversarial networks, Kang et al. used an intermediate CNN called isMap to discriminate an image that is a realistic map from other types of images that could be generated by the model (Kang et al., 2019). These two research projects show the potential of convolutional neural networks to infer label and classify maps, and this is why we used such techniques to solve our problem of map characteristics inference. The map scale inference problem is described in the following section, and the inference of map content is presented in Section 4..

3. MAP SCALE INFERENCE

This section focuses on the first issue of map scale inference from map images. We first explain how the issue was modelled as a deep learning classification problem. Then, data preparation is described, and results are presented.

3.1 Problem Formalisation

One of the most important characteristics of maps, and maybe the most important regarding generalisation and simplification, is the scale of the map. The scale of the map first helps to define the size of the symbols in the map (Ruas, 2004). For instance, if a line needs to be longer than 12.7 mm to be understood as a line in the tactile map (Štampach and Mulíčková, 2016), defining the scale of the map helps to define the minimum length of lines that should be kept from our geographic database. But map scale is also related to the use and the content of the map (Ruas, 2004, Mackaness, 2007). For instance, topographic maps at the 1:25,000 scale are often used for hiking and contain information about terrain and all the individual buildings are represented.

We propose three different ways to formalise our problem of inference of the map scale:

- learn a regression function that gives the numerical value of the scale as output of the CNN;

- classify the maps into scale categories;

- classify the maps into specific geographic extents from which the scale can be derived: if the extent of the map is the whole France and the map is to be printed on a 20 * 20 cm device, the scale of the map can be deduced.

In the first case, the regression model, the numerical value to be learned can be the scale ratio (e.g. 0.00004 for 1:25,000 scale), or the denominator of the scale (e.g. 25,000 for 1:25,000 scale). After several tests, we opted for a third solution, a normalised value between 0 and 1. We used the minimum and maximum scales from the WMTS standard [1] to normalise scale values, where 1 is the 1:500,000,000 scale.

The second formalisation proposal is to classify maps into a small number of scale categories. Scale ratios were great tools when maps were only printed on paper, at a fixed size, but now that maps are used more and more digitally, these scale ratios are less meaningful and should be replaced by scale categories when possible (Goodchild and Proctor, 1997). As an extension of a first proposal dedicated to the level of detail of OpenStreetMap map features, (Touya and Reimer, 2015), we propose to separate maps into 7 scale categories: street, city, county, region, country, continent, and world. Approximate scale ranges corresponding to these categories are provided in Table 1. Figure 2 shows three maps at different scale categories.

Category	Scale denominator min	Scale denominator max
World	50,000,000	500,000,000
Continent	8,000,000	50,000,000
Country	750,000	8,000,000
Region	150,000	750,000
County	50,000	150,000
City	15,000	50,000
Street	1	15,000

Table 1: The scale categories proposed, as an extension of the categories from (Touya and Reimer, 2015).

The third solution to undirectly learn map scale derives from the maps collected in our dataset (we have many maps of Europe, but not so many of other continents, many maps of France, but no so many of other countries), but also from the global set of characteristics we plan to infer on maps. Indeed, in addition to

---

[1] https://www.ogc.org/standards/wmts

Figure 2: Three maps of at different categories of scale: (a) country (b) continent (c) world.

scale, we need to know the geographical extent of the map, and we can derive scale from the extent and the size of the output. In this case, the categories we want to learn correspond to the geographic extents that were redundant enough in our dataset: World, Europe, France, Paris.

We tested several CNN architectures to test our different formalisms to infer scale, and we finally opted for rather simple CNN that is a mix between LeNet models(Lecun et al., 1998) used for handwritten character recognition, and AlexNet (Krizhevsky et al., 2012). The architecture, presented in Figure 3 is deeper than LeNet-5, but does not use the parallelism mechanisms of AlexNet because our dataset is much smaller than ImageNet (Russakovsky et al., 2015), for which AlexNet was designed.
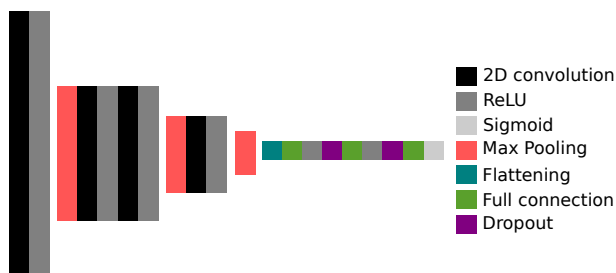

Figure 3: The layers of convolutional neural network used to learn map scales from images.

3.2 Data Preparation

Our main source of maps is a French collaborative project where teachers create open source textbooks, including geography and history textbooks for all classes (https://www.lelivrescolaire.fr/). We extracted maps from the textbooks with screen capture tools, keeping only the map and not the text that describes the map. When a legend was included along the map, it was included in the image, as it was anticipated that the legend could be used by the CNN to capture the map characteristics better. As the number of maps collected this way was quite small, we added maps that are not extracted from geography or history textbooks: we extracted several types of maps from Google Maps, OpenStreetMap, and the geoportal of the French mapping agency[2]. In most cases, the scale of the map was not available, so we used the available scale bars to measure the scale of the map; then, we stored for each map its scale (i.e. the denominator of the scale ratio), and the geographic extent of the map. The scale category was derived from the stored scale denominator using the equivalences of Table 1.

[2]https://www.geoportail.gouv.fr

Our CNN requires images with fixed dimensions (here 256*256 pixels). But the collected maps have very diverse dimensions. There are two possibilities to make all these map images square (Figure 4):

- resizing the images;
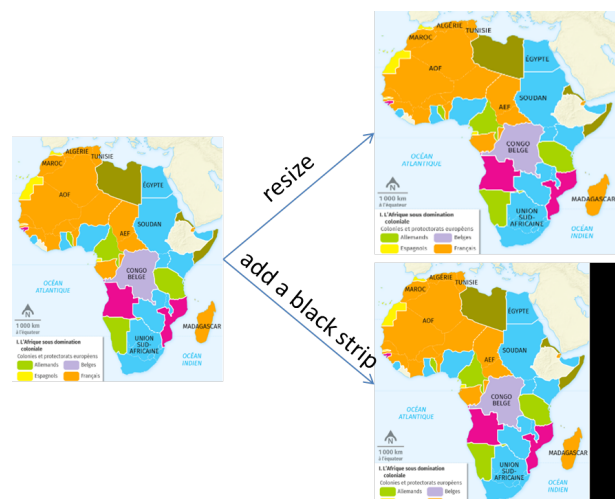- adding a strip on an edge of the image.

Figure 4: Two ways of squaring rectangle images: resizing or adding a strip (black here) on an edge of the image.

Resizing the image can cause important distortions when the width of the rectangle is very different from the length. As these distortions can bias how the model learns how to infer scale, we opted for adding a strip in the image. We tested several options for this strip, as we wanted to be sure that the model "understands" that it is not part of the map. The black strip was the one that provided the best results.

3.3 Results and Evaluation

The model was implemented with Keras Python library, and run in the Google Colaboratory platform. From the 450 maps in our dataset, we kept 93% for training, and only 7% (32 images) for the evaluation of the model. These 32 evaluation images were carefully chosen to be as diverse as possible in terms of scale and geographic extent.

The results of the regression model are really poor and we decided not to push them further. In the best cases, the predicted

normalized scale is half the reference value, but most of the time, the predicted value is three times smaller than the reference scale. These results can be explained by a training dataset that contains a very small number of unique scale values, from which the regression was clearly difficult.

The results of scale classification are also disappointing. As we were lacking maps at large scales, we only trained the model to classify maps into the four small scale categories (world to region scales in Table 1). The best classification accuracy obtained on the test maps was around 60%. The results tend to show that with maps that can be so diverse, there is no graphical feature that can differentiate scales, beyond the scale bars and the geographical extent of the map. To go further these disappointing results, we tested the same architecture with simpler images of mountain roads symbolized and generalized at two scales: 1:25,000 and 1:250,000. This use case was chosen because mountain roads have a very simplified geometry at the 1:250,000 scale, and we believed that the graphical difference between scales would be clearer than with our geography maps (Figure 5). We trained the model with 1,500 images, but the results are once again disappointing with a similar 60% accuracy of classification on the evaluation images.
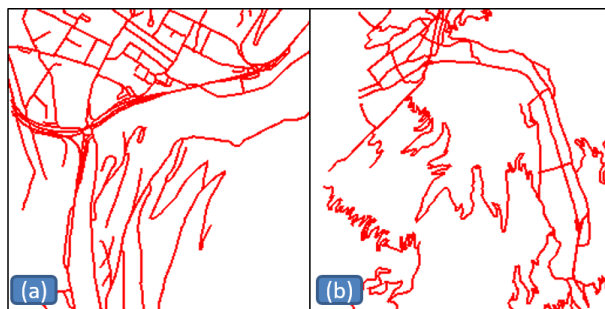

Figure 5: Two images of maps with only mountain roads at different scales: (a) 1:25,000 scale (b) 1:250,000 scale.

Table 2 shows the results obtained on the evaluation dataset when learning the geographical extent of the maps. Globally, the classification accuracy is around 71%, but this accuracy is heterogeneous between classes. The accuracy raises to 81% for the two classes with the most instances in the training set ("World" and "France"), but is much lower for the other classes that lack training examples. These results are promising and we can expect much higher accuracies if we augment the dataset with new training examples.

	World	Europe	France	Paris
World	10	0	2	0
Europe	1	3	2	1
France	0	2	7	0
Paris	0	0	2	2

Table 2: Results of geographic extent prediction on the 32 maps of our evaluation dataset.

4. MAP CONTENT INFERENCE

In this section, we describe how map content inference was formalised as an image classification problem. Similarly to the previous section, we explain how the images were prepared, and then present the results of the model.

4.1 Problem Formalisation

If we want to derive a tactile version of geography textbook maps, we need to know what is contained in the map, i.e. the map legend entries. So infering map content is similar to rebuilding the legend of the map. Here, we try to infer the content only, whatever the symbols used for this content in the legend. For instance, roads with wide red symbols in one map, and thin gray symbols in another one should similarly be infered as roads; we are not interested in their representation as it will be necessarily different in a tactile map. This remark also applies for generalisation: two maps of Germany with roads might contain different selections of important roads, but we do not aim at identifying such differences, because the generalisation will be different in a tactile map anyway.

The content of a map can be diverse, and there have been attempts to create exhaustive ontologies of map content (Iosifescu and Hurni, 2007, Abadie et al., 2010, Balley and Regnauld, 2011). In this paper, we only target a proof of concept with a small amount of training maps. Based on these collected maps, we defined a simple set of map legend entries: administrative boundaries, cities, roads, hydrography (both lakes and rivers), relief (contour lines, hypsometry, shaded relief...), vegetation (or other natural features), thematic flows (migrations, wars, commercial trades, etc.). Figure 6 shows three maps containing all these seven categories of map content.

Infering the presence of these seven content types in a map is not strictly an image classification problem, but can be seen as similar to image multi-label annotation, where several labels on the content of the image are infered (Gong et al., 2013). Multilabel annotation can easily be achieved with CNN architectures, by changing the activation function of the final fully connected layer from softmax to sigmoid. Then, rather than giving a probability to be classified into one of the (seven here) classes, the model output is a probability for each class, e.g. 84% boundaries, 13% cities, 92% hydrography, which means that the map probably contains boundaries and hydrography but not cities.

After testing several CNN architectures, we opted for the VGG-16 model (Simonyan and Zisserman, 2015), because of its past success on multi-labeling tasks with complex images, and also because it was easy to access to a pre-trained version of the model. The VGG-16 main drawback is that it is very slow to train, but using a pre-trained version prevents this drawback. The first layers of the network learn to detect low level features of the images such as contours, so they can be trained with other images than maps. Then, only the final fully connected layers are trained with our maps to infer the different labels of map content.

4.2 Data Preparation

We used the same maps as training examples as the ones used for map scale inference. The maps were labelled manually using a custom GUI developed in Python to select and store the labels of each map. In the cases displayed in Figure 6, the labeling process is straightforward, but there are maps where the content is not easily modelled by seven binary labels. Figure 7 shows two examples where the labeling process is complex. In Figure 7a, the hydrography is drawn in China, but not in the other countries. As a tactile map with hydrography all over the map would not be so different from the original map, we decided to put the label "hydrography" for this map. In Figure 7b, there is only one city, Vienna, in the map. We decided not to put the label "cities" for this map, because it would bias the learning process, but such choice forbids a faithful reproduction of the map.

The problem of the required square size of the images that occurs with scale inference also occurs with the CNN used for content inference. But as VGG-16 is a deeper model than the one used for scale inference, we need much more training examples to avoid over-fitting. To solve both issues, rather than adding a black strip
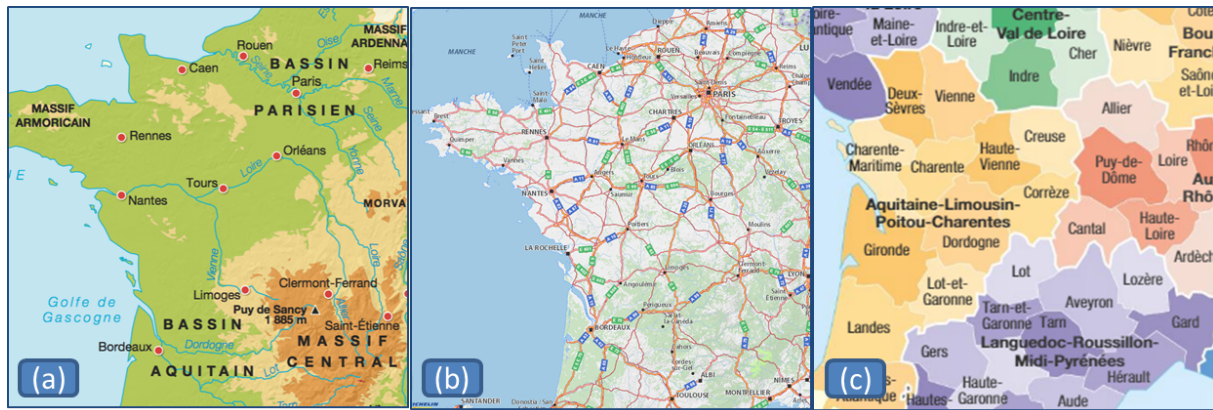
Figure 6: Three maps of with different types of content: (a) cities, relief, and hydrography (b) cities, roads, and vegetation (c) administrative boundaries.

Figure 7: Some maps are complex to label: (a) hydrography is drawn in China but not in the other countries; (b) There is only one city in the map (Vienna).

on an edge of the image, we cut each map in four squares that overlap to preserve a large portion of the map in each part (Figure 8). The length of the side of the square is 80% the length of the shortest side of the initial rectangle.
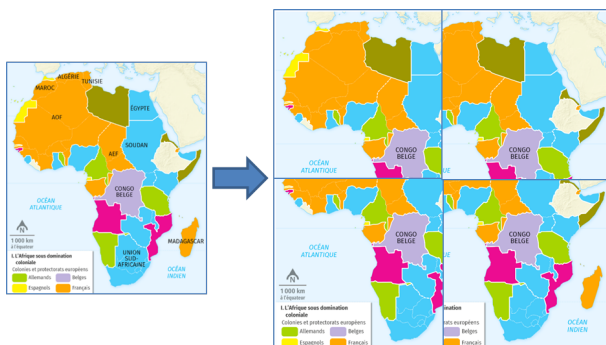

Figure 8: The dataset is augmented by cutting each map in four square images with overlaps.

When the map images are augmented by cutting them in four, there might be some images with a large portion of the legend, and only a small portion of the image depicting the map (Figure 9). To avoid this problem the maps used for content inference are cropped to remove the legend as much as possible (in the map of Figure 9, the bottom of the image is cropped). After the augmentation process, our dataset contains 1,600 map images, which remains quite small for a model as deep as VGG-16.

4.3 Results and Evaluation

The VGG-16 model was loaded with the Keras Python library, and the last fully connected was changed to fit our multi-label

Figure 9: This map of Africa contains a large legend, but when cut in four for data augmentation, the bottom left part (red dashes) is mainly the legend and the sea, and we do not see the map so much.

output. Only this new final layer was trained with our maps, allowing us to benefit from the effective training of the other layers with a much higher number of training images. Table 3 shows the best results we obtained for content inference. We did not test the label "vegetation" because there was not enough training examples. From the 1,600 images of our dataset, we used 88% of the dataset for training and 12% (192 images) for evaluation.

	TP	FP	TN	FN	Precision	Recall
Boundaries	23	4	25	2	0.85	0.92
Cities	29	4	18	3	1	0.94
Hydrography	16	0	37	1	0.88	0.91
Relief	62	12	107	11	0.84	0.85
Flows	27	11	143	11	0.71	0.71
Roads	9	1	40	4	0.90	0.69

Table 3: Results of multi-labeling on our evaluation dataset. TP is for True Positive, FP for False Positive, TN for True Negative, FN for False Negative.

Table 3 shows the best results we obtained to label the maps of our evaluation dataset. The three labels that appear the most in the training maps are the boundaries, the cities, and hydrography, and these three labels are not surprisingly the ones with the best precision and recall values. Regarding flows, the diversity of flow representations coupled with the small number of training examples explains the low values of precision and recall. Regarding

roads, it is interesting to note the difference between the high precision and the rather low recall. It means that there are few false positive results, i.e. the model does not infer roads in maps where it should not, and that there are many false negative results, i.e. the model does not label all the maps that do contain roads.

5. DISCUSSION

In this paper, the problem of map characteristics inference has been simplified because our first goal was to achieve a proof of concept rather than solving the complete problem. In this section, the remaining issues to allow a further derivation of tactile maps from these infered characteristics are discussed.

Administrative or country boundaries are often represented in the maps of geography or history textbooks, but if it is easy to collect the current boundaries, they are not always the ones that are used in the map, in particular with maps from history textbooks. Figure 10 illustrates this problem with three maps of Europe at different periods, with different boundaries between countries/realms/empires. To automatically reproduce these maps, we need to use the right boundaries and this is far from simple because the date needs to be infered and the ancient boundaries need to be available at vector format. We do not believe that infering the date with deep learning techniques would work (not enough training examples for each date), so in this case, an achievable target would be a semi-automatic process where a base map would be provided to a transcriptor, who would be responsible for drawing the correct boundaries to finalise the map.

We also simplified map characteristics by ignoring map projections. But map projection can play a key role in the way the information is conveyed, particularly for small scale maps that are frequent in geography textbooks. For instance, an equal-area projection of the world does not represent it the same way as the classical conformal Mercator projection. Figure 11 shows another example where the projection should be infered because it is important in the way the cartographic message is conveyed: the polar projection allows a really insightful view of the flows, which would not be possible with a classical projection. Contrary to dates, we do believe that projections can be infered by trained CNNs, at least family of projections for small scale maps (e.g. from country to world scale categories).

The style of the map is also one of its important characteristics as style is largely responsible for the way the map information is conveyed to the map reader, and creative or original styles can make maps better (Christophe, 2012, Christophe et al., 2016). Infering the style of the map in addition to the other infered characteristics can be useful for two reasons: (1) many visually impaired people do see some colors, so colored and visually styled maps can be useful in addition to the tactile graphics; (2) it might be possible to translate some styles into tactile counterparts, to increase the expressivity of tactile maps. In their classification of map images, (Zhou et al., 2018) defined some map categories that differ because of their style, so it seems possible to infer style categories (e.g. black and white map), or style choices (e.g. flows in red), which can later be used to derive the tactile map.

The text in the map was also completely ignored in these first experiments. Even if Braille is not understood by all visually impaired people, it is still used to add textual information such as names that cannot be graphically conveyed (Miller et al., 2010). Extracting the text from images of maps has long been a research topic (Pierrot Deseilligny et al., 1995, Pouderoux et al., 2007, Yao-Yi Chiang and Knoblock, 2010, Gobbi et al., 2019), and seems feasible with current deep learning techniques given the

progresses made on other optical character recognition problems with such techniques (Lecun et al., 1998, Zhou et al., 2017).

Finally, the examples of training images showed in this paper often include the legend of the map, but this legend was not specifically used while it contains a lot of the information that we want to infer. It should be easy to automatically extract the legend for the images and train different networks with the map and legend separately. The same remark applies to scale bars, that seem be the only reliable graphical hint to infer scale. Another reason to specifically address the legend of the map is to be able to derive a tactile legend as well as a tactile map. To analyse the legend with computer vision techniques and to derive a tactile one, it will be necessary to formally model the legend of a map, as proposed in (Christophe, 2012). And the legend is not the only information that can be useful to improve our inferences on map characteristics, as maps in geography textbooks are often accompanied by some text that describes the map. Past research on relating the text of journal articles with accompanying maps could be useful to adapt here (Brun et al., 2015).

6. CONCLUSION AND FUTURE WORK

In order to ease the access to tactile maps for geography teachers of visually impaired children, we want to build an online platform where teachers submit the image of a map from geography (or history) textbooks, and further receive a tactile version of this map created automatically or semi-automatically with a professional transcriptor. As a first step towards the design of such a platform, this paper presents experiments to infer the scale and the content of such maps with deep convolutional neural networks. Scale is infered as a category or scale range rather than as an exact numerical ratio, and content is infered as multi-labeling problem. The preliminary results are really promising considering the small amount of maps collected for now to train our models.

To go further, several possible improvements have been discussed in the previous section. Beside the obvious necessity for a larger and more diverse training dataset, the most important topics to tackle seem to be the projections and the legend (or even the text accompanying the map).

This paper presents the starting point of a large project[3] and there are many issues left to automatically derive a tactile map from the image of the map for people with normal vision. Besides the possible improvements of step (1) described above, step (2) on data collection, (3) on the simplification and generalisation of the collected data, and (4) on the derivation of a 3D model from the 2D simplified map, all need to be addressed in future years. Experiments with professional transcriptors and visually impaired people wil be carried out throughout the project to verify the usability of all our current and future propositions.

ACKNOWLEDGEMENTS

REFERENCES

Abadie, N., Mechouche, A. and Mustière, S., 2010. OWL-based formalisation of geographic databases specifications. In: Pro-

---

[3]https://activmap.limos.fr/

Figure 10: Three maps of Europe at different periods of time for history classes: the country borders are not the same.
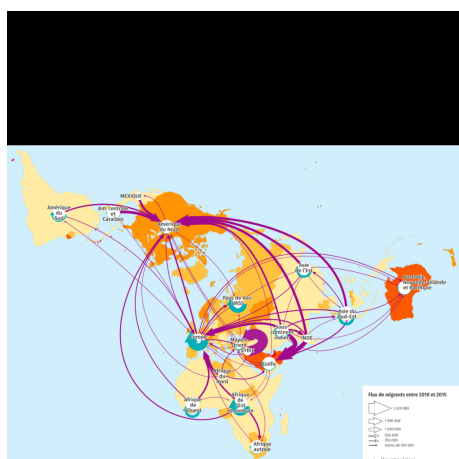

Figure 11: World map with a polar projection.

ceedings of 17th International Conference on Knowledge Engineering and Knowledge Management (EKAW'10), Lisbon (Portugal).

Balley, S. and Regnauld, N., 2011. Models and Standards for On-Demand Mapping. In: Proceedings of 25th International Cartographic Conference (ICC'11), Paris, France.

Brock, A., Oriola, B., Truillet, P., Jouffrais, C. and Picard, D., 2013. Map design for visually impaired people: past, present, and future research. Médiation et Information, Handicap et Communication 36, pp. 117–129.

Brock, A., Truillet, P., Oriola, B., Picard, D. and Jouffrais, C., 2015. Interactivity Improves Usability of Geographic Maps for Visually Impaired People. HumanComputer Interaction 30, pp. 156–194.

Brun, G., Domingues, C. and Van Damme, M.-D., 2015. TEXTOMAP : determining geographical window for texts. In: Proceedings of the 9th Workshop on Geographic Information Retrieval, GIR'2015, ACM, Paris, France.

Christophe, S., 2012. Cartographic Styles between traditional and original (towards a cartographic style model). In: Proceedings of AutoCarto Conference, Columbus (Ohio), USA.

Christophe, S., Dumenieu, B., Turbet, J., Hoarau, C., Mellado, N., Ory, J., Loi, H., Masse, A., Arbelot, B., Vergne, R., Brédif, M., Hurtut, T., Thollot, J. and Vanderhaeghe, D., 2016. Map Style Formalization: Rendering Techniques Extension for Cartography. In: D. Mould and P. Bénard (eds), Proceedings of Non-Photorealistic Animation and Rendering, Lisbon, Portugal.

Edman, P., 1992. Tactile Graphics. American Foundation for the Blind.

Fillières-Riveau, G., Barra, V., Favreau, J.-M. and Touya, G., 2019. Enrichissement d'orthophotographie par des données OpenStreetMap pour l'apprentissage machine. In: SAGEO, Clermont-Ferrand, France.

Gobbi, S., Ciolli, M., La Porta, N., Rocchini, D., Tattoni, C. and Zatelli, P., 2019. New Tools for the Classification and Filtering of Historical Maps. ISPRS International Journal of Geo-Information 8(10), pp. 455. Number: 10 Publisher: Multidisciplinary Digital Publishing Institute.

Gong, Y., Jia, Y., Leung, T., Toshev, A. and Ioffe, S., 2013. Deep convolutional ranking for multilabel image annotation.

Goodchild, M. and Proctor, J., 1997. Scale in a digital geographic world. Geographical and Environmental Modelling 1(1), pp. 5–23.

Iosifescu, I. and Hurni, L., 2007. Towards Cartographic Ontologies or " How Computers Learn Cartography". In: Proceedings of 23rd International Cartographic Conference, ICA, Moscow, Russia.

Kang, Y., Gao, S. and Roth, R. E., 2019. Transferring Multiscale Map Styles Using Generative Adversarial Networks. International Journal of Cartography 5(2-3), pp. 115–141.

Krizhevsky, A., Sutskever, I. and Hinton, G. E., 2012. Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems, pp. 1097–1105.

LeCun, Y., Bengio, Y. and Hinton, G., 2015. Deep learning. Nature 521(7553), pp. 436–444.

Lecun, Y., Bottou, L., Bengio, Y. and Haffner, P., 1998. Gradient-based learning applied to document recognition. Proceedings of the IEEE 86(11), pp. 2278–2324.

Lobben, A., 2015. Tactile Maps and Mapping. Journal of Blindness Innovation and Research.

Mackaness, W. A., 2007. Understanding Geographic Space. In: W. A. Mackaness, A. Ruas and T. Sarjakoski (eds), The Generalisation of Geographic Information: Models and Applications, Elsevier, Amsterdam, Netherlands, pp. 1–10.

Mackaness, W. and Reimer, A., 2014. Generalisation in the Context of Schematised Maps. In: D. Burghardt, C. Duchêne and W. Mackaness (eds), Abstracting Geographic Information in a Data Rich World, Lecture Notes in Geoinformation and Cartography, Springer International Publishing, pp. 299–328.

Miller, I., Pather, A., Milbury, J., Hasty, L., O'Day, A. and Spence, D., 2010. Guidelines and Standards for Tactile Graphics. Technical report, The Braille Authority of North America.

Pierrot Deseilligny, M., Le Men, H. and Stamon, G., 1995. Character string recognition on maps, a rotation-invariant recognition method. Pattern Recognition Letters 16(12), pp. 1297–1310.

Pouderoux, J., Gonzato, J.-C., Pereira, A. and Guitton, P., 2007. Toponym Recognition in Scanned Color Topographic Maps. In: Ninth International Conference on Document Analysis and Recognition (ICDAR 2007), Vol. 1, pp. 531–535. ISSN: 2379-2140.

Rowell, J. and Ungar, S., 2003. The world of touch: an international survey of tactile maps. Part 2: design. British Journal of Visual Impairment 21(3), pp. 105–110.

Ruas, A., 2004. Le changement de niveau de détail dans la représentation de l'information géographique. Université de Marne-la-Vallée.

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A. C. and Fei-Fei, L., 2015. ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (IJCV) 115(3), pp. 211–252.

Simonyan, K. and Zisserman, A., 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv:1409.1556 [cs]. arXiv: 1409.1556.

Touya, G. and Reimer, A., 2015. Inferring the Scale of OpenStreetMap Features. In: J. Jokar Arsanjani, A. Zipf, P. Mooney and M. Helbich (eds), OpenStreetMap in GIScience, Lecture Notes in Geoinformation and Cartography, Springer International Publishing, pp. 81–99.

Touya, G., Christophe, S., Favreau, J.-M. and Ben Rhaiem, A., 2019. Automatic derivation of on-demand tactile maps for visually impaired people: first experiments and research agenda. International Journal of Cartography 5(1), pp. 65–91.

Vasconcellos, R., 1996. Tactile Mapping Design and the Visually Impaired User. In: C. H. Wood and C. P. Keller (eds), Cartographic Design: Theoretical and Practical Perspectives, Wiley, New York, USA, pp. 91–102.

Wabiński, J. and Mościcka, A., 2019. Automatic (Tactile) Map Generation—A Systematic Literature Review. ISPRS International Journal of Geo-Information 8(7), pp. 293.

Yao-Yi Chiang and Knoblock, C. A., 2010. An approach for recognizing text labels in raster maps. In: 2010 20th International Conference on Pattern Recognition, pp. 3199–3202.

Zhou, X., Li, W., Arundel, S. T. and Liu, J., 2018. Deep Convolutional Neural Networks for Map-Type Classification. arXiv:1805.10402 [cs, stat]. arXiv: 1805.10402.

Zhou, X., Yao, C., Wen, H., Wang, Y., Zhou, S., He, W. and Liang, J., 2017. EAST: An Efficient and Accurate Scene Text Detector. In: CVPR 2017, IEEE, Honolulu, HI, USA. arXiv: 1704.03155.

Červenka, P., Břinda, K., Hanousková, M., Hofman, P. and Seifert, R., 2016. Blind Friendly Maps. In: K. Miesenberger, C. Bühler and P. Penaz (eds), Computers Helping People with Special Needs, Lecture Notes in Computer Science, Vol. 9759, Springer International Publishing, pp. 131–138.

Štampach, R. and Mulíčková, E., 2016. Automated generation of tactile maps. Journal of Maps 12(sup1), pp. 532–540.