A deep learning approach to Cadastral Computing

Sofia Ares Oliveira, Isabella di Lenardo, Bastien Tourenc, Frederic Kaplan
{sofia.oliveiraares, isabella.dilenardo, bastien.tourenc, frederic.kaplan}@epfl.ch

Digital Humanities Laboratory, Ecole Polytechnique Fédérale de Lausanne

1 Introduction

Among all the diverse typologies of administrative systems, the fiscal-cadastral sources retracing the ownership of lands are undoubtedly the richest and, in a sense, the most coherent records. The need for efficient tax collection processes pushed different administrations and political regimes in the 18th century to develop scale-invariant coherent property tracing system. These systems were progressively generalised at the turn of the 19th century with the abolition of privileges for certain classes of citizens and the establishment of republican administrations for managing cities and countries [1].

The cadastre, by introducing a constant collection of taxes calculated on a fixed percentage, defines a new conception of the city and of the urban space. It marks the transition from the Ancient Regimes to the Modern State. The introduction of a constant, proportional and impartial tax determined a new conception of the city, adapted to statistical computations. In some sense, for the first time, the cadastre transforms the city into a *computational object*. We use the neologism "cadastral computing" to refer to the operations performed on such primary sources.

This article presents a generic approach for processing automatically the information contained in the Napoleonic cadastres. The cadastres established during the first years of the 19th century cover a large part of Europe. For many cities they give one of the first geometrical surveys, linking precise parcels with identification numbers. These identification numbers point to register lines with the names of the parcel's owners (Fig. 1). As the Napoleonic cadastres include millions of parcels, it therefore offers a detailed snapshot of large part of Europe's population at the beginning of the 19th century [2]. To develop a generic approach adapted to the processing of these administrative
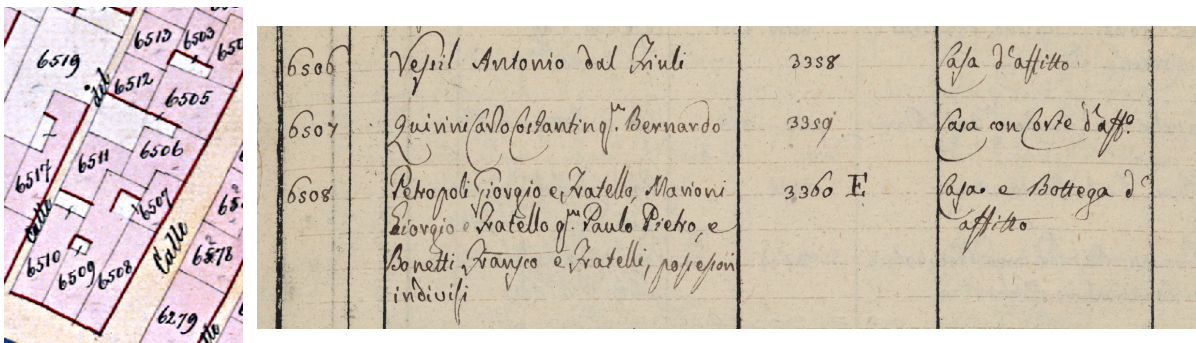

Figure 1: Identification numbers in parcels and in the register

documents, one needs to solve two difficult challenges : (1) developing algorithms capable of robustly segmenting maps into parcels and administrative tables into cells (2) developing solutions for transcribing handwritten text containing people or places mentions and identification numbers. Until recently, these two problems were considered much beyond the state-of-the-art. The results of this article are based on the important progress recently made on both issues, using deep learning architectures. In 2017, an initial study on cadastre extraction showed promising results in parcel extraction and identifier recognition. However it was entirely designed as an ad hoc pipeline fine-tuned for a particular cadastre [3]. A year later, a generic deep-learning segmentation engine, relying on a convolutional neural network [4] demonstrated that it was possible to design a generic architecture to segment many typologies of documents. In this article we use the genericity of this approach to develop the first fully automatic pipeline to transform the Napoleonic Cadastres into an information system.

2 Method

The automatic processing of the cadastre maps aims at extracting the parcels as geometrical shapes and also at transcribing the parcel's identification numbers. It is composed of three main steps:

1. Training of the deep neural network on manually annotated data

2. Segmentation of the maps into meaningful regions and objects

3. Transcription of the identification numbers

Our segmentation network is a fully convolutional neural network inspired by the U-Net architecture [5]. It uses a ResNet-50 pretrained network [6] as encoder, which speeds up the training, reduces the amount of training data needed and helps generalization. The full details of the architecture and open-source implementation can be found in [4].
The network is trained to extract the parcel contours and text using annotated data from the Venetian cadastre (Fig. 2). The training data corresponds to roughly 1/3 of one map sheet among the 26 maps of the city of Venice (roughly 800 parcels).

The transcription network is a convolutional recurrent neural network (CRNN) [7] which pro-

Figure 2: Sample of training data for cadastre maps segmentation. Parcels contour are in red, text is in green.

duces a chain of characters when given an image segment containing text. The network is trained on samples of numbers from the Venetian archives and on numbers synthetically generated with MNIST digits [8] (Fig. 3).

(a) synthetically generated numbers
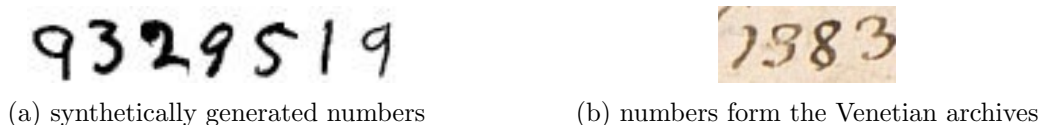

(b) numbers form the Venetian archives

Figure 3: Example of training data for the transcription system.

The segmentation model obtained after the training is able to predict the parcel contours and text region at pixel level (Fig. 4). Watershed by flooding algorithm [9] is applied on parcel contours predictions, which allows the extraction of parcel objects as polygonal shapes. Text regions are cropped, horizontally aligned and converted into grayscale image segments for further processing by the transcription system. The image segments containing text are fed to the transcription network,

(a) text extraction

(b) contour extraction
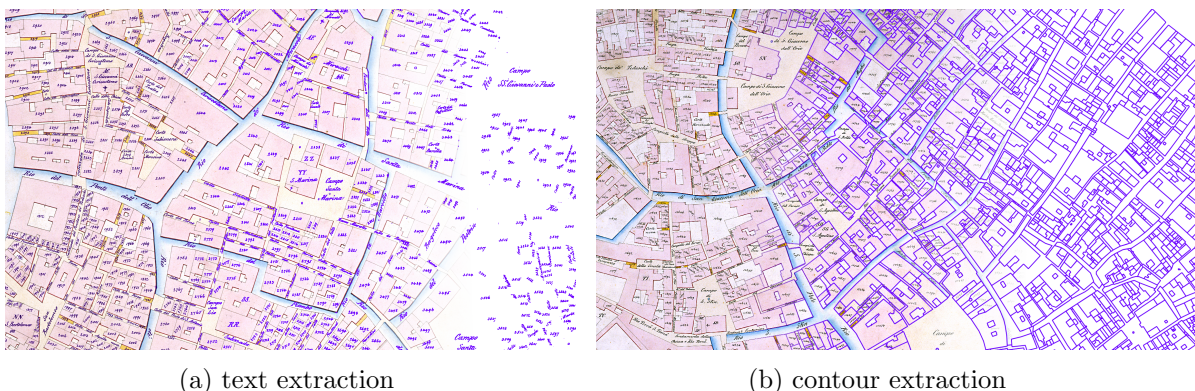
Figure 4: Output of the segmentation network (overlay in purple)

which outputs a prediction of a number, i.e a chain of digits. Each transcription is then linked to its corresponding parcel.

The contours of the parcels (whether they contain an identifier number or not) are saved as polygonal shapes and are exported into JSON format. In our case, since the images have previously been georeferenced, the coordinates are exported as geographical coordinates and can therefore directly be imported in any geographic information system.

3 Results

Two evaluations are performed in order to assess the performances of the system: the geographical accuracy of the extracted parcels and the transcription of the identification numbers.

3.1 Geometrical evaluation

The number of geometrical shapes extracted and manually annotated are listed in Tab. 1. After a filtering step, which keeps only shapes which area range from $2m^2$ to $15000m^2$, the total number of extracted parcels is 28711, among which 18138 contain a transcription.

The quality of the parcel's extraction is evaluated by measuring the intersection over union (IoU) between the geometries produced automatically and almost 17000 manually annotated shapes. Precision and recall with three different IoU thresholds 0.5 (acceptable), 0.7 (good), 0.9 (excellent) are reported in Tab. 2. The recall value shows that a large majority of parcels are extracted (85% in the most strict case). The low precision value is mainly due to the incorrect extraction of streets,

Figure 5: Visualization of the results of the automatic extraction of parcels. The red rectangle indicates the parcels used as training data for the segmentation network.

Geometries	Number
Geometries extracted automatically	31 342
Geometries remaining after filtering	28 711
with ID number	18 138
Manually annotated geometries	16 946
with ID number	15 634

Table 1: Total number of geometries in the automatic extraction and manual annotation. The first three rows relate to automatically extracted parcels, the two last rows show the statistics for manually annotated parcels.

squares, canals, etc. that are currently not filtered out (example in Fig. 6).

3.2 Transcription evaluation

We assess the performance of the transcription of parcel's identifier numbers by computing the number of correct predictions and report the precision and recall values in Tab. 3. The current method assumes that the identifiers are located within the parcel, thus, identifiers partially or completely outside the geometrical shape are not correctly transcribed (Fig. 7), resulting in a lower recall.

In order to increase the precision value and since we can assume that spatially close parcels will have numerically close identifiers, we tried to discard false predictions by determining if a transcription was 'plausible' or not, using information from its spatial neighbourhood. Thus, a transcription is considered as an outlier if the (numerical) difference between the predicted number and the median of its 5 neighbouring transcriptions is greater than 10. This results in a significant increase in

Figure 6: Example of false extraction of streets and canals (in blue)

IoU	Correct parcels	Precision	Recall
t=0.5	15999	0.557	0.944
t=0.7	15292	0.533	0.902
t=0.9	14440	0.503	0.852

Table 2: Evaluation of the geometrical shape extraction with different Intersection over Union (IoU) thresholds

precision (up to 93%), but at the expense of a decrease in recall.

	Correct transcriptions	Precision	Recall
Transcriptions	11101	0.612	0.710
Transcriptions after outlier detection	8070	0.927	0.516

Table 3: Evaluation of the transcriptions of parcel's identifiers numbers

4 Perspectives

This initial study demonstrates on a particular case that 90% of the urban geometrical structure and more than 50% of a city population can be automatically remapped with high precision using only generic pipelines. Even if these numbers need to be confirmed on the basis of other case studies, the genericity of the methods used makes us optimistic about the possibility of conducting a large scale study in the coming years. Such datasets call for a confrontation with the large number of historical hypotheses that have been formulated on the urban society at the beginning of the 19th century based on much smaller sets of evidence. Thanks to the standardisation processes of the Napoleonic administration, we hope in the coming months to extend this systematic processing beyond Venice to a large part of Europe.
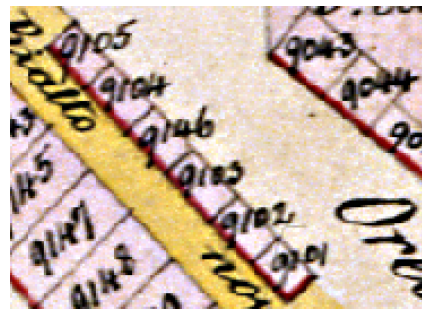
Figure 7: Example of identifiers numbers outside or partially outside the parcel

References

[1] R. J. P. Kain and E. Baigent, *Cadastral Map in the Service of the State: History of Property Mapping.* Univ. Chicago P., 1993.

[2] P. c. Clergeot, *Cent millions de parcelles en France : 1807, un cadastre pour l'Empire.* Éditions Publi-Topex, 2007.

[3] S. Ares Oliveira, I. di Lenardo, and F. Kaplan, "Machine vision algorithms on cadaster plans," in *Premiere Annual Conference of the International Alliance of Digital Humanities Organizations (DH 2017)*, 2017.

[4] S. Ares Oliveira, B. Seguin, and F. Kaplan, "dhsegment: A generic deep-learning approach for document segmentation," in *Proceedings of 16th International Conference on Frontiers in Handwriting Recognition (ICFHR-2018)*, 2018.

[5] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241, Springer, 2015.

[6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

[7] B. Shi, X. Bai, and C. Yao, "An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 11, pp. 2298–2304, 2017.

[8] Y. LeCun, C. Cortes, and C. J. Burges, "The mnist database of handwritten digits," 1998.

[9] S. Beucher, "Use of watersheds in contour detection," *Proceedings of the International Workshop on Image Processing*, 1979.