

First-Order Rule Induction for the Recognition of Morphological Patterns in Topographic Maps

D. Malerba, F. Esposito, A. Lanza, F.A. Lisi

Dipartimento di Informatica, University of Bari
via Orabona 4, 70126 Bari, Italy
{malerba | esposito | lanza | lisi}@di.uniba.it

Abstract. Information given in topographic map legends or in GIS models is often insufficient to recognize interesting geographical patterns. Some prototypes of GIS have already been extended with a knowledge-base and some reasoning capabilities to support sophisticated map interpretation processes. Nevertheless, the acquisition of the necessary knowledge is still an open problem to which machine learning techniques can provide a solution. This paper presents an application of first-order rule induction to pattern recognition in topographic maps. Research issues related to the extraction of first-order logic descriptions from vectorized topographic maps are introduced. The recognition of morphological patterns in topographic maps of the Apulia region is presented as a case study.

1 Introduction

Handling digitized maps raises several research issues for the field of pattern recognition. For instance, raster-to-vector conversion of maps has received increasing attention in the community of graphics recognition [6]. In fact, obtaining vector data from a paper map is a very expensive and slow process, which often requires manual intervention. While supporting the map acquisition process is important, it is equally useful and even more challenging to automate the *interpretation* of a map in order to locate some geographic objects and their relations [12]. Indeed information given by map legends or given as basis of data models in Geographical Information Systems (GIS) is often insufficient to recognize not only *geographical objects* relevant for a certain application, but also *patterns* of geographical objects which geographers, geologists and town planners are interested in. Map interpretation tasks such as the detection of morphologies characterizing the landscape, the selection of important environmental elements, both natural and artificial, and the recognition of forms of the territorial organization require abstraction processes and deep domain knowledge that only human experts have.

Several studies show the difficulty of map interpretation tasks. For instance, a study on the drawing instructions of Bavarian cadastral maps (scale 1:5000) pointed out that symbols for road, pavement, roadside, garden and so on were defined neither in the legend nor in the GIS-model of the map [16]. In a previous work in cooperation with researchers from the Town Planning Department of the Polytechnic of Bari, an

environmental planning expert system was developed for administrators responsible for urban planning [2], [1]. The system was able to provide them with appropriate suggestions but presumed that they had good skills in reading topographic maps to detect some important ground morphology elements, such as system of cliffs, ravines, and so on. These are some examples of morphological patterns that are very important in many civil and military applications but never explicitly represented in topographic maps or in a GIS-model.

Empowering GIS with advanced pattern recognition capabilities would support effectively map readers in map interpretation tasks. Some prototypes of GIS have already been extended with a knowledge-base and some reasoning capabilities in order to support sophisticated map interpretation processes [20]. Nevertheless, these systems have a limited range of applicability for a variety of reasons mainly related to the knowledge acquisition bottleneck.

A solution to these difficulties can come from machine learning. In this paper we present an application of first-order rule induction to pattern recognition in topographic maps. Research issues related to the extraction of first-order logic descriptions from vectorized topographic maps are introduced. The task of topographic map interpretation as a whole is supported by INGENS (Inductive Geographic Information System), a prototypical GIS extended with a training facility and an inductive learning capability [16]. In INGENS, each time a user wants to retrieve geographic complex objects or patterns not explicitly modeled in the Map Repository, he/she can prospectively train the system to the recognition task within a special user view. Training is based on a set of examples and counterexamples of geographic concepts of interest to the user (e.g., ravine or steep slopes). Such (counter-) examples are provided by the user who detects them on stored maps by applying browsing, querying and displaying functions of the GIS interface. The symbolic representation of the training examples is automatically extracted from maps by the module Map Descriptor. The module Learning Server implements one or more inductive learning algorithms that can generate models of geographic objects from the chosen representations of training examples. In this paper, we will focus our presentation on the first-order rule induction algorithm ATRE [15].

The data model for the Map Repository of INGENS is described in the next section. In Section 3, the feature extraction algorithms implemented in the Map Descriptor are sketched. Section 4 is devoted to the first-order rule induction algorithm ATRE made available in the Learning Server. A case study, namely the recognition of relevant morphological patterns on topographic maps of the Apulia region, is presented and discussed in Section 5. Conclusions and future work are reported in Section 6.

2 A Data Model for Topographic Maps

Many GIS store topographic maps. In the Map Repository of INGENS each map is stored according to a *hybrid tessellation – topological model*. The tessellation model follows the usual topographic practice of superimposing a regular grid on a map in order to simplify the localization process. Indeed each map in the repository is

divided into square cells of the same size. For each cell the raster image in GIF format is stored together with its coordinates and component objects. In the topological model of each cell it is possible to distinguish two different structural hierarchies: *physical* and *logical*.

The physical hierarchy describes the geographical objects by means of the most appropriate physical entity, that is: point, line or region. In different maps of the same geographical area, the same object may have different physical representations. For instance, a road can be represented as a line on a small-scale map, or as a region on a large-scale map. Points are described by their spatial coordinates, while (broken) lines are characterized by the list of line vertices, and regions are represented by their boundary line. Some topological relationships between points, lines and regions are modeled in the conceptual design, namely points inside a region or on its border, and regions disjoining/meeting/overlapping/containing/equaling/covering other regions. The meaning of the topological relationships between regions is a variant of that reported in the 9-intersection model by Egenhofer and Herring [7], in order to take into account problems due to approximation errors.

The logical hierarchy expresses the semantics of geographical objects, independent of their physical representation. Since the conceptual data model has been designed to store topographic maps, the logical entities concern geographic layers such as hydrography, orography, land administration, vegetation, administrative (or political) boundary, ground transportation network, construction and built-up area. Each of them is, in turn, a generalization meaning that, for instance, an administrative boundary must be classified in one of the following classes: city, province, county or state.

3 Feature Extraction from Vectorized Topographic Maps

In INGENS the content of a map cell is described by means of a set of *features*. Here the term feature is intended as a characteristic (property or relationship) of a geographical entity. This meaning is similar to that commonly used in Pattern Recognition (PR) and differs from that attributed by people working in the field of GIS, where the term feature denotes the unit of data by which a geographical entity is represented in computer systems and, according to the OGC terminology, is modelled through a series of properties [17], [21].

In PR, feature is a synonym for discriminatory property of objects which have to be recognised and classified. Obviously, the number of features needed to successfully perform a given recognition task depends on the discriminatory qualities of the chosen features. However, the problem of *feature selection* (i.e. what discriminatory features to select), is usually complicated by the fact that the most important features are not necessarily easily measurable. *Feature extraction* is an essential phase which follows the segmentation in the classical recognition methodology [11]. In PR, features are classified into three categories according to their nature: *physical*, *structural*, and *mathematical* [22]. The first two categories are used primarily in the area of image processing, while the third one includes statistical

means, correlation coefficients and so on. In map interpretation tasks a different category of features is required, namely *spatial* features.

Tables 1 and 2 show a taxonomy of spatial features that can be extracted from vectorized maps. The first distinction to be made concerns the *type* of feature: it can be an *attribute*, that is a property possessed by the spatial object, or a *relation* that holds among the object itself and other objects. Spatial relationships among geographic objects are actually conditions on object positions.

According to the *nature* of the feature, it is possible to distinguish among:

- *Locational* features, when they concern the position of the objects. The position of a geographic object will be represented by numeric values expressing coordinates for example in latitude/longitude or in polar coordinates or others.
- *Geometric* features, when they depend on some computation of metric/distance. Area, perimeter, length are some examples. Their domain is typically numeric.
- *Topological* features (actually only a relation can be topological), when they are preserved under topological transformations, such as translation, rotation, and scaling. Topological features are generally represented by nominal values.
- *Directional* features, when they concern orientation (e.g., north, north-east, and so on). Generally, a directional feature is represented by means of nominal values.

Clearly, a geo-referenced object also has *aspatial* features, such as the name, the layer label, and the temperature. Many other features can be extracted from maps, some of which are hybrid in the sense that merge properties of two or more categories. For instance, the features that express the conditions of parallelism and perpendicularity of two lines are both topological and geometrical. They are topological since they are invariant with respect to translation, rotation and stretching, while they are geometrical since their semantics is based on the size of their angle of incidence. Another example of hybrid spatial feature is represented by the relation of “faraway-west”, whose semantics mixes both directional and geometric concepts. Finally, some features might mix spatial relations with *aspatial* properties, such as the feature that describes coplanar roads by combining the condition of parallelism with information on the type of spatial objects.

The problem of extracting features from maps has been mainly investigated in the fields of document processing and graphics recognition, nevertheless most of the

Table 1. A classification of attributive features.

<i>ATTRIBUTES</i>			
SPATIAL			ASPATIAL
LOCATIONAL	GEOMETRIC	DIRECTIONAL	Name Layer Type Others (temperature, no. inhabitants, ...)
Co-ordinate (x,y) of a point (centroid, extremal points, bounding rectangles, ...)	<ul style="list-style-type: none"> ▪ Area ▪ Perimeter ▪ Length of axes Other shape properties 	Orientation of major axis	

Table 2. A classification of relational features.

<i>RELATIONS</i>			
SPATIAL			ASPATIAL
GEOMETRIC	TOPOLOGICAL	DIRECTIONAL	<ul style="list-style-type: none"> ▪ Instance-of ▪ Hierarchical relation (sub-type, super-type) ▪ Aggregation/Composition
<ul style="list-style-type: none"> ▪ Distance ▪ Angle of incidence 	<ul style="list-style-type: none"> ▪ Region-to-Region ▪ Region-to-Line ▪ Region-to-Point ▪ Line-to-Line ▪ Line-to-Point ▪ Point-to-Point 	<ul style="list-style-type: none"> ▪ Neighbouring relations 	

work reported in the literature concerns raster maps, where the issues are how to isolate text from graphics as in the work by Pierrot *et al.* [18], or how to extract particular geographical objects, such as contour lines as in [6], or points and lines as in the work by Yamada *et al.* [23] or land-use classes using thematic maps as in [3].

The lack of works on vectorized representations can be attributed to the main usage of topographic maps made in the field of GIS: only for rendering purposes. The rare applications to vectorized maps reported in the literature refer to cadastral maps, as in [5].

A first application of feature extraction algorithms to vectorized topographic maps can be found in the work by Esposito *et al.* [8]. This work is a natural evolution of the collaboration already established between a research group on Machine Learning of the University of Bari with the Town Planning Department of the Polytechnic of Bari in order to develop an expert system for environmental planning [2], [1]. For environmental planning tasks, fifteen features were specified with the help of domain experts (see Table 3). Being quite general, they can also be used to describe maps on different scales. In INGENS they are extracted by the module *Map Descriptor*, which generates first-order logic descriptions of the maps stored in the Map Repository.

Actually, feature extraction procedures working on vectorized maps are far from being a simple “adaptation” of existing graphics recognition algorithms. In fact, the different data representation (raster vs. vector) makes the available algorithms totally unsuitable to vectorized maps, as it is the case of all filters based on the mathematical morphology [23]. Each feature to be extracted needs a specific procedure to be developed basing upon the geometrical, topological and topographical principles, which are involved in the semantics of that feature.

For instance, the relation *distance* between two “parallel” lines is computed by means of the following algorithm. Let O_1 and O_2 be two geographical linear objects represented by n and m coordinate pairs, respectively. Without loss of generality, let us assume that $n \geq m$. The algorithm first computes d_{\min_h} as the minimum distance

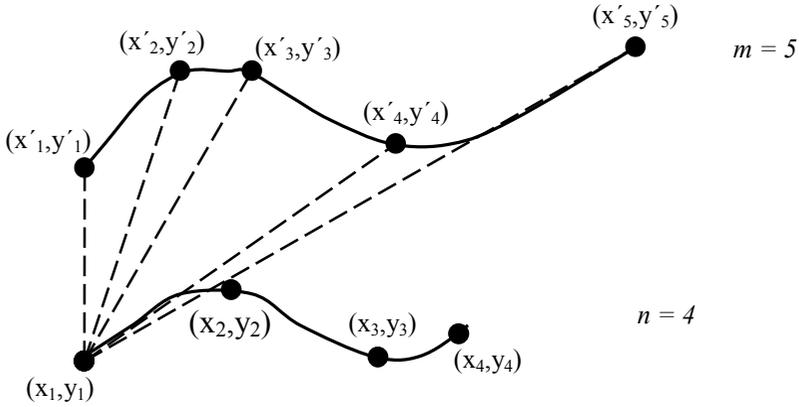


Fig. 1. Computation of the distance between two “parallel” lines.

between the h -th point of O_1 and any point of O_2 (see Figure 1). Then, the distance between O_1 and O_2 is computed as follows:

$$distance = \frac{\sum_{h=1}^n d \min_h}{n} \tag{1}$$

The complexity of this simple feature extraction algorithm is $O(mn)$ though less computationally expensive solutions can be found by applying multidimensional access methods [10].

The descriptions obtained for each cell are quite complex, since some cells contain dozens of geographic objects of various types. For instance, the cell shown in Figure 2 contains one hundred and eighteen distinct objects, and its complete description is a clause with more than one thousand literals in the body.

4 The Induction of First-Order Rules with ATRE

Sophisticated end users may train INGENS to recognize geographical patterns that are not explicitly modeled in the Map Repository. To support this category of users, the module *Learning Server* places some inductive learning systems at their disposal. We will focus our attention on the first-order rule induction algorithm ATRE [14].

The distinguishing feature of ATRE is that it can induce recursive logical theories from a set of training examples. Here the term *logical theory* (or simply *theory*) denotes a set of first-order definite clauses. An example of logical theory is the following:

- $downtown(X) ?high_business_activity(X), onthesea(X).$
- $residential(X) ?close_to(X,Y), downtown(Y), low_business_activity(X).$
- $residential(X) ?close_to(X,Y), residential(Y), low_business_activity(X).$

Table 3. Features extracted for the generation of map descriptions.

Feature	Meaning	Type	Domain	
			Type	Values
CONTAIN(X,Y)	Cell X contains object Y	Topologic relation	boolean	{true, false}
TYPE_OF(Y)	Object Y type	Aspatial attribute	nominal	33 nominal values
SUBTYPE_OF(Y)	Specialization of object Y type	Aspatial attribute	nominal	101 nominal values that are specializations of type_of domain
COLOR(Y)	Object Y color	Aspatial attribute	nominal	{blue, brown, black}
AREA(Y)	Object Y area	Geometrical attribute	linear	[0..MAX_AREA]
DENSITY(Y)	Object Y density	Geometrical attribute	ordinal	Symbolic names chosen by expert user
EXTENSION(Y)	Object Y extension	Geometrical attribute	linear	[0..MAX_EXTENTION]
GEOGRAPHIC_DIRECTION(Y)	Geographic direction of Y	Directional attribute	nominal	{north, east, north_west, north_east}
LINE_SHAPE(Y)	Shape of the linear object Y	Geometrical attribute	nominal	{straight, curvilinear, cuspidal}
ALTITUDE(Y)	Altitude of Y	Geometrical attribute	linear	[0.. MAX_ALTITUDE]
LINE_TO_LINE(Y,Z)	Spatial relation between two lines Y and Z	Hybrid relation	nominal	{almost parallel, almost perpendicular}
DISTANCE(Y,Z)	Distance between objects Y and Z	Geometrical relation	linear	[0..MAX_DISTANCE]
REGION_TO_REGION(Y,Z)	Spatial relation between two regions Y and Z	Topological relation	nominal	{disjoint, meet, overlap, covers, covered_by, contains, equal, inside}
LINE_TO_REGION(Y,Z)	Spatial relation between a line Y and a region Z	Hybrid relation	nominal	{along_edge, intersect}
POINT_TO_REGION(Y,Z)	Spatial relation between a point Y and a region Z	Topological relation	nominal	{inside, outside, on_boundary, on_vertex}

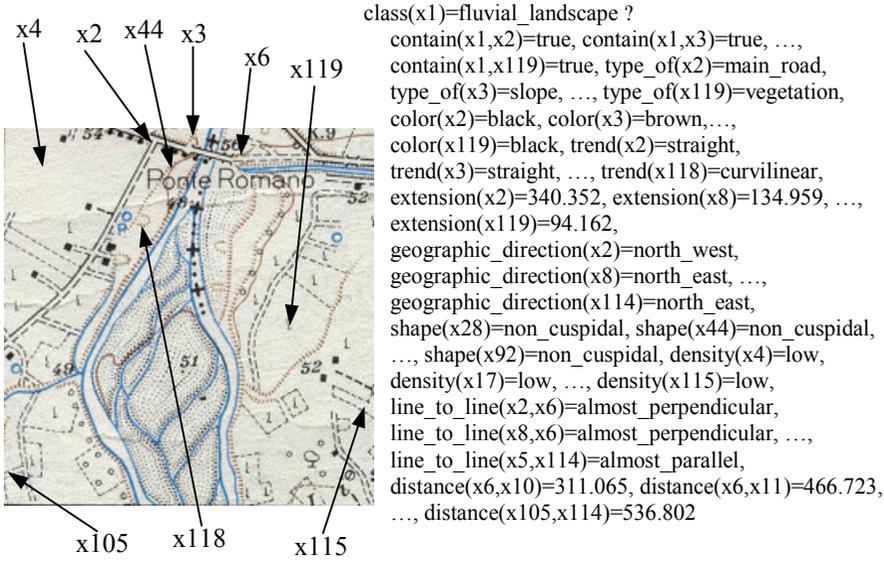


Fig. 2. A partial logical description of a cell. The constant $x1$ represents the whole cell, while all other constants denote the one hundred and eighteen enclosed objects. Distances and extensions are expressed in meters.

It expresses sufficient conditions for the two concepts of “main business center of a city” and “residential zone,” which are represented by the unary predicates *downtown* and *residential*, respectively.

The learning problem solved by ATRE can be formulated as follows:

Given

- a set of concepts $C_1, C_2, ?, C_r$ to be learned,
- a set of observations O described in a language L_O ,
- a background knowledge BK described in a language L_{BK} ,
- a language of hypotheses L_H ,
- a generalization model Γ over the space of hypotheses,
- a user’s preference criterion PC ,

Find

a (possibly recursive) logical theory T for the concepts $C_1, C_2, ?, C_r$, such that T is complete and consistent with respect to O and satisfies the preference criterion PC .

The *completeness* property holds when the theory T explains all observations in O of the r concepts C_i , while the *consistency* property holds when the theory T explains no counter-example in O of any concept C_i . The satisfaction of these properties guarantees the correctness of the induced theory with respect to O .

As regards the representation languages L_O, L_{BK}, L_H , the basic component is the *literal*, which takes two distinct forms:

$$f(t_1, ?, t_n) = \text{Value} \quad (\text{simple literal}) \quad f(t_1, ?, t_n) ? [a..b] \quad (\text{set literal}),$$

where f and g are function symbols called *descriptors*, t_i 's and s_i 's are terms, and $[a..b]$ is a closed interval. Descriptors can be either *nominal* or *linear*, according to the ordering relation defined on its domain values. Some examples of literals are:

$color(X)=blue$, $distance(X,Y)=63.9$, $width(X) \in [82.2 .. 83.1]$, and $close_to(X,Y)=true$. The last example points out the lack of predicate symbols in the representation languages adopted by ATRE. Thus, the first-order literals $p(X,Y)$ and $?p(X,Y)$ will be represented as $f_p(X,Y)=true$ and $f_p(X,Y)=false$, respectively, where f_p is the function symbol associated to the predicate p . Henceforth, for the sake of simplicity, we will adopt the usual notation $p(X,Y)$ and $?p(X,Y)$. Furthermore, the interval $[a..b]$ in a set literal $f(X_1, \dots, X_n) \in [a..b]$ is computed according to the same information theoretic criterion used in INDUBI/CSL [13].

Observations in ATRE are represented as ground multiple-head clauses, called *objects*, which have a conjunction of simple literals in the head. Multiple-head clauses present two main advantages with respect to definite clauses: higher comprehensibility and efficiency. The former is basically due to the fact that multiple-head clauses provide us with a compact description of multiple properties to be predicted in a complex object such as those we may have in map interpretation. The second advantage derives from the possibility of having a unique representation of known properties shared by a subset of observations.

The *background knowledge* defines any relevant problem domain knowledge. It is expressed by means of *linked*, *range-restricted* definite clauses [4] with simple and set literals in the body and one simple literal in the head. The same constraints are applied to the language of hypotheses.

ATRE implements a novel approach to the induction of recursive theories [9]. To illustrate how the main procedure works, let us consider the following instance of the learning problem:

Observations	O $downtown(zone_1) \ ? \ residential(zone_1) \ ? \ residential(zone_2) \ ?$ 1 $?downtown(zone_2) \ ? \ ?downtown(zone_3) \ ? \ residential(zone_4) \ ?$ $?downtown(zone_4) \ ? \ ?downtown(zone_5) \ ? \ ?residential(zone_5) \ ?$ $?residential(zone_6) \ ? \ downtown(zone_7) \ ? \ ?residential(zone_7) \ ?$ $onthesea(zone_1), high_business_activity(zone_1),$ $close_to(zone_1, zone_2),$ $low_business_activity(zone_2), close_to(zone_2, zone_4),$ $adjacent(zone_1, zone_3), onthesea(zone_3),$ $low_business_activity(zone_3), low_business_activity(zone_4),$ $close_to(zone_4, zone_5), high_business_activity(zone_5),$ $adjacent(zone_5, zone_6), low_business_activity(zone_6),$ $close_to(zone_6, zone_8), low_business_activity(zone_8),$ $close_to(zone_1, zone_7), onthesea(zone_7),$
BK	$close_to(X,Y) \leftarrow adjacent(X,Y)$ $close_to(X,Y) \leftarrow close_to(Y,X)$
Concepts	C $downtown(X)=true$ 1 $residential_zone(X)=true$ C 2
PC	2 Minimize/maximize negative/positive examples explained by the theory

The first step towards the generation of inductive hypotheses is the saturation of all observations with respect to the given *BK* [19]. In this way, information that was

implicit in the observation, given the background knowledge, is made explicit. In the example above, the saturation of O_1 involves the addition of the nine literals logically entailed by BK, that is $close_to(zone_2, zone_1)$, $close_to(zone_1, zone_3)$, $close_to(zone_3, zone_1)$, $close_to(zone_7, zone_1)$, $close_to(zone_4, zone_2)$, $close_to(zone_5, zone_4)$, $close_to(zone_5, zone_6)$, $close_to(zone_6, zone_5)$, and $close_to(zone_8, zone_6)$.

Initially, all positive and negative examples are generated for every concept to be learned, the learned theory is empty, while the set of concepts to be learned contains all C_i . With reference to the above input data, the system generates two positive examples for C_1 ($downtown(zone_1)$ and $downtown(zone_7)$), two positive examples for C_2 ($residential(zone_2)$ and $residential(zone_4)$), and eight negative examples equally distributed between C_1 and C_2 ($?downtown(zone_2)$, $?downtown(zone_3)$, $?downtown(zone_4)$, $?downtown(zone_5)$, $?residential(zone_1)$, $?residential(zone_3)$, $?residential(zone_6)$, $?residential(zone_7)$).

Once the observations have been saturated and examples have been generated, the separate-conquer loop starts. The step of *parallel conquer* generates a set of consistent clauses, whose minimum number is defined by the user. Since clauses are consistent, they should explain no negative example. For instance, by requiring the generation of at least one consistent clause with respect to the examples above, this procedure returns the following set of clauses:

$downtown(X) ? onthesea(X), high_business_activity(X).$
 $downtown(X) ? onthesea(X), adjacent(X, Y).$
 $downtown(X) ? adjacent(X, Y), onthesea(Y).$

The first of them is selected according to the preference criterion (procedure *find_best_clause*). Actually, the hypothesis space of the concept *residential* has been simultaneously explored, but at the time in which the three consistent clauses for the concept *downtown* have been found, no consistent clause for *residential* has been discovered yet. Thus the parallel conquer step stops since the number of consistent clauses is greater than one.

Since the addition of a consistent clause to the partially learned theory may lead to an augmented, inconsistent theory, it is necessary to verify the *global consistence* of the learned theory and eventually reformulate the theory in order to recover the consistency property without repeating the learning process from scratch. The learned clause is used to saturate again the observation. Continuing the previous example, the two literals added to O_1 are $downtown(zone_1)$ and $downtown(zone_7)$. This operation enables ATRE to generate also definitions of the concept *residential* that depend on the concept *downtown*. Indeed, at the second iteration of the separate-conquer cycle, the parallel conquer step returns the clause:

$residential(X) ? close_to(X, Y), downtown(Y), low_business_activity(X).$

and by saturating again the observation with both learned clauses, it becomes possible to generate a recursive clause at the third iteration, namely

$residential(X) ? close_to(X, Y), residential(Y), low_business_activity(X).$

The separate step consists of tagging positive examples explained by the current learned theory, so that they are no longer considered for the generation of new clauses. The separate-conquer loop terminates when all positive examples are tagged, meaning that the learned theory is complete as well as consistent.

5 The Recognition of Morphological Patterns in Topographic Maps: A Case Study

The first-order rule induction algorithm ATRE has been applied to the recognition of four morphological patterns in topographic maps of the Apulia region, Italy, namely *regular grid system of farms*, *fluvial landscape*, *system of cliffs* and *royal cattle track*. Such patterns are deemed relevant for the environmental protection, and are of interest to town planners. A regular grid system of farms is a particular model of rural space organization that originated from the process of rural transformation. The fluvial landscape is characterized by the presence of waterways, fluvial islands and embankments. The system of cliffs presents a number of terrace slopes with the emergence of blocks of limestone. A royal cattle track is a road for transhumance that can be found exclusively in the South-Eastern part of Italy.

The territory considered in this application covers 131 km² in the surroundings of the Ofanto River, spanning from the zone of Canosa to the Ofanto mouth. More precisely, the examined area is covered by five map sheets on a scale of 1:25000 produced by the IGMI (Ofanto mouth – 165 II S.W., Barletta 176 I N.W., Canne della Battaglia – 176 IV N.E., Montegrosso 176 IV S.E., Canosa 176 IV S.W.).

The maps have been segmented into square observation units of 1 Km² each. The choice of the gridding step, which is crucial for the recognition task, has been made using the advice of a team of fifteen geomorphologists and experts in environmental planning, giving rise to a one-to-one mapping between observation units of the map and cells in the database.

Thus, the problem of recognizing the four morphological patterns can be reformulated as the problem of labeling each cell with at most one of four labels. Unlabelled cells are considered uninteresting for environmental protection.

As previously mentioned, ATRE extends the system INGENS with a training functionality and an inductive learning capability in order to overcome the difficulties related to the acquisition of operational definitions for the recognition task. ATRE was trained according to the experimental design briefly presented below. One hundred and thirty-one cells were selected, each of which was described in the symbolic language illustrated in the previous Section and assigned to one of the following five classes: system of farms, fluvial landscape, system of cliffs, royal cattle track and other. The last class simply represents “the rest of the world,” and no classification rule is generated for it. Indeed, its assigned cells are not interesting for the problem of environmental protection being studied, and they are always used as negative examples when ATRE learns classification rules for the remaining classes. Forty-five cells from the map of Canosa were selected to train the system, while the remaining eighty-six cells were randomly selected from the four maps of the Ofanto mouth, Barletta, Canne della Battaglia and Montegrosso. Training observations represent about 35% of the total experimental data set. An example of partial logical description of a training cell is shown in Figure 2.

A fragment of the logical theory induced by ATRE is reported below:

```
class(X1) = fluvial_landscape ?contain(X1,X2), color(X2)=blue,
            type_of(X2)=river,trend(X2)=curvilinear, extension(X2)?[325.00..818.00].
class(X1) = fluvial_landscape ?contain(X1,X2), type_of(X2)=river, color(X2)=blue,
```

relation(X3,X2)=almost_perpendicular,
extension(X2)?[615.16..712.37],trend(X3)=straight.
class(X1)=system_of_farms ?contain(X1,X2), color(X2)=black,
relation(X2,X3)=almost_perpendicular,
relation(X3,X4)=almost_parallel,type_of(X4)=interfarm_road,
geographic_direction(X4)=north_est,
extension(X2)?[362.34 .. 712.25], color(X3)=black, type_of(X3)=farm_road,
color(X4)=black.

The first two clauses explain all training observations of fluvial landscape. In particular, the first states that cells labeled as *fluvial_landscape* contain a long, curvilinear, blue object of type river, while the second clause states that cells concerning a fluvial landscape may also present a long, straight, blue object that is perpendicular to another object (presumably, a bridge). The third clause refers to the system of farms. From the training observations, the machine learning system induced the following definition: "There are two black objects, namely an interfarm road (*X4*) and a farm road (*X3*), which run almost parallel to the north-east, and are both perpendicular to a long black object". This definition of system of farms is not complete since it includes other clauses that ATRE actually generated but are not reported in this paper. It is easy to see that the classification rules are intelligible and meaningful. Some experimental results obtained in a previous work are reported in [8].

By matching these rules with logical descriptions of other map cells it is possible to automate the recognition of complex geographical objects or geographical patterns that have not been explicitly modeled by a set of symbols.

6 Conclusions

Automated map interpretation is a challenging application domain for pattern recognition. Knowledge of the meaning of symbols reported in the map legends is not generally sufficient to recognize interesting geographical complex objects or patterns on a map. Moreover, it is quite difficult to describe such patterns in a machine-readable format. That would be tantamount to providing GIS with an operational definition of abstract concepts often reported in texts and specialist handbooks. In order to enable the automation of map interpretation tasks in GIS, a new approach has been proposed in this paper. The idea is to ask GIS users for a set of classified instances of the patterns that interest them, and then apply a first-order rule induction algorithm to generate the operational definitions for such patterns. These definitions can be either used to recognize new occurrences of the patterns at hand in the Map Repository. An application to the problem of Apulian map interpretation has been reported in this paper in order to illustrate the advantages of the proposed approach.

This work is still in progress and many problems have to be solved. As for the data model for topographic maps, the segmentation of a map in a grid of suitably sized cells is a critical factor, since over-segmentation leads to a loss of recognition of global effects, while under-segmentation leads to large cells with an unmanageable number of components. To cope with the first problem, it is necessary to consider the *context* of a cell, that is the neighboring cells, both in the training phase and in the

recognition phase. To solve problems caused by under-segmentation it is crucial to provide users with appropriate tools that hide irrelevant information in the cell description. Indeed, a set of generalization and abstraction operators will be implemented in order to simplify the complex descriptions currently produced by the *Map Descriptor*.

As for the algorithm ATRE, we plan to further investigate the influence of both the representation and the content of observations in the training set on experimental results. Case studies stressing the capability of autonomously discovering concept dependencies should also be faced.

Acknowledgements

This work is in partial fulfillment of the research objectives set by the IST project SPIN! (Spatial Mining for Data of Public Interest) funded by the European Union (<http://www.ccg.leeds.ac.uk/spin/>).

References

1. Barbanente, A., Borri, D., Esposito, F., Leo, P., Maciocco, G., Selicato, F.: Automatically acquiring knowledge by digital maps in artificial intelligence planning techniques. In: Frank, A.U., Campari, I., Formentini, U. (eds.): *Theories and Methods of Spatio-Temporal Reasoning*. Lecture Notes in Artificial Intelligence, Vol. 482. Springer-verlag, Berlin (1992) 89-100.
2. Borri D. et al.: Green: Building an Operational Prototype of Expert System for Planning Control in Urban Environment. *Proceedings of the European Conference of the Regional Science Association*, Istanbul (1990).
3. Centeno J.S.: Segmentation of Thematic Maps Using Colour and Spatial Attributes. In: Tombre, K., Chhabra, A.K. (eds.): *Graphics Recognition Algorithms and Systems*. Lecture Notes in Computer Science, Vol. 1389, Springer-Verlag, Berlin (1998) 221-230.
4. de Raedt, L.: *Interactive theory revision: An inductive logic programming approach*. Academic Press, London (1992).
5. den Hartog, J., Holtrop, B.T., de Gunst, M.E., Oosterbroek E.P.: Interpretation of Geographic Vector-Data in Practice. In Chhabra, A.K., Dori, D. (eds.): *Graphics Recognition Recent Advances*. Lecture Notes in Computer Science, Vol. 1941, Springer-Verlag, Berlin (1999) 50-57
6. Dupon, F., Deseilligny, M.P., Gondran, M.: Automatic Interpretation of Scanned maps: Reconstruction of Contour Lines. In: Tombre, K., Chhabra, A.K. (eds.): *Graphics Recognition: Algorithms and Systems*. Lecture Notes in Computer Science, Vol. 1389, Springer-Verlag, Berlin (1998) pp. 1-8.
7. Egenhofer, M. J., Herring, J.R.: Categorising topological spatial relations between point, line, and area objects. In: Egenhofer, M.J., Mark, D.M., Herring, J.R. (eds.): *The 9-intersection: formalism and its use for natural language spatial predicates*. Technical Report NCGIA 94-1, Santa Barbara, (1994).
8. Esposito, F., Lanza, A., Malerba, D., Semeraro, G.: Machine learning for map interpretation: an intelligent tool for environmental planning. *Applied Artificial Intelligence* 11(7-8) (1997) 673-695.

9. Esposito, F., Malerba, D., Lisi, F.A.: Induction of Recursive Theories in the Normal ILP Setting: Issues and Solutions. In: Cussens, J., Frish, A. (eds.): *Inductive Logic Programming*, Vol. 1866, Springer, Berlin (2000) 93-111.
10. Gaede, V., Günther O.: *Multidimensional Access Methods*, ACM Computing Surveys, 30(2) (1998) 170-231.
11. Haralick, R.M., Shapiro, L.G.: *Computer and Robot Vision*, Addison-Wesley, Reading, MA (1992).
12. Keates, J. S.: *Map understanding*. Second Edition. Longman, Edinburgh (1996).
13. Malerba, D., Esposito, F., Semeraro, G., Caggese, S.: Handling Continuous Data in Top-down Induction of First-order Rules. In: Lenzerini, M. (ed.): *AI*IA 97: Advances in Artificial Intelligence*, Lecture Notes in Artificial Intelligence, Vol. 1321. Springer, Berlin (1997) 24-35.
14. Malerba, D., Esposito, F., Lisi, F.A.: Learning Recursive Theories with ATRE. In: Prade, H. (ed.): *Proceedings of the 13th European Conf. on Artificial Intelligence*, Wiley, Chichester (1998) 435-439.
15. Malerba, D., Esposito, F., Lanza, A., Lisi, F.A.: Discovering Geographic Knowledge: The INGENS System. In: Ras, Z.W., Ohsuga, S. (eds.): *Foundations of Intelligent Information Systems*, Lecture Notes in Artificial Intelligence, 1321, Springer, Berlin (2000) 40-48.
16. Mayer, H.: Is the knowledge in map-legends and GIS-models suitable for image understanding? *International Archives of Photogrammetry and Remote Sensing* 30(4) (1994) 52-59.
17. Open GIS Consortium: *The OpenGIS Abstract Specification* (1996), <http://www.opengis.org/public/abstract.html>.
18. Pierrot Deseilligny, M., Le Men, H., Stamon, G.: Characters String Recognition on Maps: A Method for High Level Reconstruction. *Proceedings of the Third International Conference on Document Analysis and Recognition*, Vol. 1. (1995) 249-252.
19. Rouveirol, C.: Flattening and saturation: Two representation changes for generalization. *Machine Learning* 14(2) (1994) 219-232.
20. Smith, T., Donna, P., Sudhakar, M., Pankaj, A.: KBGIS-II: A Knowledge-Based Geographic Information System. *International Journal of Geographic Information Systems* 1(2) (1987) 149-172.
21. Sondheim, M., Gardels, K., Buehler, K.: GIS Interoperability. In: Longley, P.A., Goodchild, M.F., Maguire, D.J., Rhinds, D.W. (eds.): *Geographical Information Systems, Principles and Technical Issues*, Vol. 1. John Wiley & Sons (1999) 347-358.
22. Tou, J. T., Gonzales, R. C.: *Pattern Recognition Principles*. Addison-Wesley, Reading, MA (1974).
23. Yamada, H., Yamamoto, K., Hosokawa, K.: Directional Mathematical Morphology and Reformalized Hough Transformation for the Analysis of Topographic Maps, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(4) (1993) 380-387.