Artif Intell Rev (2006) 26:129–140
DOI 10.1007/s10462-007-9043-4

Spatial inference with geometric proportional analogies

Emma-Claire Mullally · Diarmuid P. O'Donoghue

Published online: 3 October 2007

Abstract We describe an instance-based reasoning solution to a variety of spatial reasoning problems. The solution centers on identifying an isomorphic mapping between labelled graphs that represent some problem data and a known solution instance. We describe a number of spatial reasoning problems that are solved by generating non-deductive inferences, integrating topology with area (and other) features. We report the accuracy of our algorithm on different categories of spatial reasoning tasks from the domain of Geographical Information Science. The generality of our approach is illustrated by also solving geometric proportional (IQ-test type) analogy problems.

Keywords Analogical similarity · Spatial inference · Topographic maps

1 Introduction

We present a model that offers an instance-based approach to generating inferences on discrete spatial information. Collections of geometric figures are compared to pre-stored solution instances and an identical match means that the pre-stored solution can be applied to that problem. In this paper we extend the range of problem instances that can be identified over our previous work, increasing the types of problem that can be solved. We illustrate the generality of our approach by illustrating its operation on both geometric proportional analogy (GPA) problems and by reporting results on problems related to topographic (land-cover) maps.

E.-C. Mullally (✉) · D. P. O'Donoghue
Department of Computer Science, National University of Ireland, Maynooth, Ireland
e-mail: emmaclaire.mullally@nuim.ie

D. P. O'Donoghue
Department of Computer Science/National Centre for Geocomputation, National University of Ireland, Maynooth, Ireland
e-mail: diarmuid.odonoghue@nuim.ie

This paper explores a non-obvious similarity between GPA problems and some problems related to topographic maps. We see both as typically involving small, local collections of polygons where the solution relies on some simple manipulations to these polygon neighbourhoods. In the remainder of this paper we shall examine other approaches to these and other qualitative spatial reasoning problems. We then look at a class of GPA problems that have previously been overlooked. These involve GPA problems where the figures involved contain colour or pattern information. We then describe the CSM (Contextual Structure Matching) algorithm that solves these problems, highlight similarities and differences with other models for interpreting analogical comparisons. We then discuss some problems with topographic maps and compare the performance of our algorithm with expert human map-readers. Finally, we discuss some opportunities for extending our work.

2 Qualitative spatial reasoning

Spatial reasoning involves reasoning about locations in the real world and in diagrams, maps, and schematics. Qualitative spatial reasoning is the problem of reasoning about properties and relations recorded on spatial data. Two basic approaches to qualitative spatial reasoning (QSR) have emerged. First, the formal approach aims to support deductive reasoning about spatial relations. The main spatial algebra are DE9IM (Egenhofer and Herring 1990) used by the OGC (1999) and IBM's DB2 Spatial Extender, and the RCC8 (Randell et al. 1992). These aim to reason about non-quantitative spatial information, such as adjacency, containment and disconnectedness.
The second approach to QSR explores non-deductive reasoning, typically exploring more cognitively aware inference techniques (Evans 1967; Bohan and O'Donoghue 2000; Forbus et al. 2003). Our approach is applied to problem of enriching the classification of existing topographic maps and in particular OS MasterMap® which is a large scale digital map of Great Britain. Whilst OS MasterMap provides classification for individual features such as buildings it does not yet explicitly represent complex structures such as schools. Text exists to identify the location of such features but there is no explicit association between the individual simple features and the complex feature.

The remainder of this paper is organised as follows. First, we describe the specific problems within qualitative spatial reasoning. Secondly, we describe GPA problems and give a brief explanation of how they are solved. Finally, we show how *CSM*, our model, can be used both to resolve GPA problems and to solve problems identified in topographic maps.

3 Geometric proportional analogy (GPA) problems

The first qualitative spatial reasoning problem we examine concerns GPA problems. GPAs are a comparison formed between two collections of geometric figures of the form A:B::C:D, where A and B form the source domain, and C and D are the target domain. In a GPA problem, A, B and C are known while the solution D, is unknown. The objective is to use the information contained in the source domain (A and B) as a basis for completing the partial description C—thereby generating the missing solution D—see Fig. 1.

The first computational model to solve GPA problems was created by Evans (1967). This model solved many of the "Miller Analogy Test" problems, which are a set of intelligence test questions. Evans computational model used visual shape matching to process the problem and select a solution from a list of five candidate solutions (D1–D5). Additionally, Evans

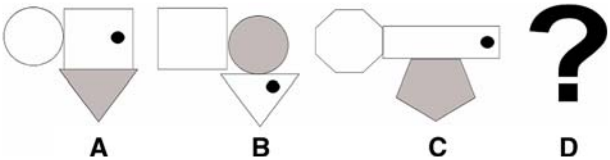

Fig. 1 A GPA

model only operated on plain GPAs that had no colour or pattern information. As we shall see in Sect. 3.1 manipulating attributes will prove central to our solution.

The Structure Matching Engine (SME) (Gentner 1983, Falkenhainer et al. 1989) has also been used to solve GPAs (Tomai et al. 2004). SME represents the source and target domains using predicates, generating the A to C mapping by finding the largest isomorphic mapping between the predicate structures of A and C. Tomai et al. (2004) also solved GPAs by processing the images representing the source and target domains, selecting the best from a list of five possible solutions. Tomai's solution looks at the overall shape of the image as well as using any of its rotation and reflection information. Like Evans model, Tomai's model is also limited to addressing plain GPA problems (without colours, patterns etc).

Three factors make these models unsuitable for our purposes. First, we need to solve GPAs that involve attributes like colour and category information. Second, in solving GPA's we need to include rotation and orientation information, whereas this information is not relevant when processing topographic maps and will thus be ignored. Therefore our model must be capable of dealing with rotation invariant data. Finally, we need to generate the solution rather then selecting it from a list of candidate solutions.

3.1 GPAs involving attributes

This paper focuses on GPA problems that include and manipulate geometric objects that include attributes, such as colour or pattern information. Solving the GPA in Fig. 1 necessitates applying the *shaded* attribute of the pentagon in part C to the rectangle in the solution (D). Failure to include the attribute would result in an incorrect solution.

Two previous models have incorporated attribute information in GPA problems. Bohan and O'Donoghue (2000) examine a variety of GPA problems involving attributes. However, this model is specific to GPAs and does not deal with topographic maps or with the complex polygon clusters we shall describe later. Mulhare et al. (2001) describe a model for classifying objects in topographic maps, but this model does not address GPAs nor does it use point information or deal with complex polygon structures. As we shall see, our general purpose solution solves GPA problems as well as problems with topographic maps.

4 CSM: contextual structure matching algorithm

The first step in our solution to these GPA problems is to represent each image in symbolic form using predicates. These predicates detail how the geometric objects in each image are spatially related and also detail any attributes that an object may have. Each part (A–C) of the problem is treated as a Voronoi diagram, from which a Delaunay diagram is created. Objects are then joined if a spatial relationship exists between these objects. Examining the edges allows the relationship between the objects in each part of the image to become clear.

Figure 2 shows the Delaunay diagram for parts A and B. Three relationships are used to describe the local topology of parts A, B and C: `line-adjacent` represents objects that

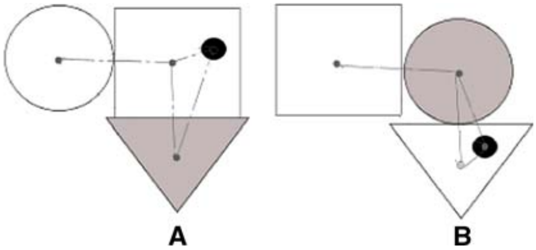

Fig. 2 Delaunay diagram

Table 1 Predicates describing parts A and B of the GPA

Predicates describing A	Predicates describing B
Line adjacent(circle, square)	Line adjacent(square, circle)
Line adjacent(square, triangle)	Line adjacent(circle, triangle)
Inside(point, square)	Inside(point, triangle)
Shaded(triangle)	Shaded(circle)
Plain(circle)	Plain(triangle)
Plain(square)	Plain(square)

share a common boundary, `point-adjacent` represents objects that share (only) a common vertex, while `inside` represents an object that exists completely within the boundaries of another object. Colour attributes are represented by either `plain` or `shaded`. Examining the changes in the predicates describing A and B highlights the transformation that was applied to A in order to produce B (see Table 1).

We do not use the relations of the RCC (Randell et al. 1992, Gerevini and Renz 2002) or DE9IM (Egenhofer and Herring 1990) for the following reasons. First, our relations are derived directly from the representation of topographic data. Second, our relations correspond to several RCC or DE9IM relations, but using the more compact representation expedites the expensive structure matching process. It is a relatively straight-forward matter to derive the RCC or DE9IM representation equivalent to our relations. Finally, we do not represent RCC's disconnected relation (DC) as adding this would slow down the later graph isomorphism problem.

In this problem the colour attributes of the circle and the triangle are changed during the transformation. Some previous methods for solving GPAs (Evans 1967; Gentner 1983; Tomai et al. 2004) do not find the correct transformation in order to solve this problem as they ignore the attributes.

Next we identify the largest consistent 1-to-1 mapping between the descriptions of the corresponding parts of the GPA, parts A and C. Like Gentner (1983) this mapping is subject to the predicate identicality constraint, whereby only identically named relationships are mapped together. The size, shape, colour and rotation of the objects are not taken into consideration at this stage. Figure 3 shows how this process identified the object correspondence between each object in A and C.

Finally, the transformation between A and B is applied to C, generating D. This method allows D, shown in Fig. 4, to be generated.

4.1 The CSM algorithm

The description of each part of the problem forms a labelled graph. Each geometric area forms a vertex and is assigned to one of 13 labels, corresponding to the 13 topographic categories.

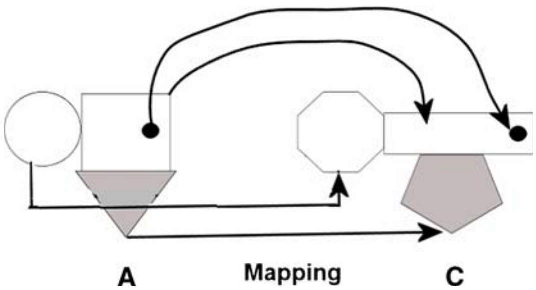

Fig. 3 Mapping A–C

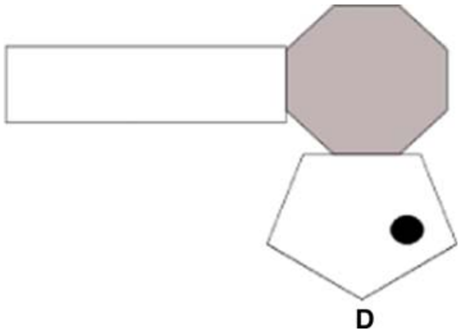

Fig. 4 The solution to the GPA in Fig. 1

Text points also correspond to a vertex forming an additional category. Relations between geometric objects form edges and are assigned to one of 3 additional labels, corresponding to relations *line-adjacent, point-adjacent* and *has-point*.

This labelling of the graph impacts on the central *structure matching* process, which identifies the isomorphism between two labelled graphs. One graph represents a problem situation while the other represents a solution template. The structure matching algorithm ensures that only identically labelled vertices and nodes can be placed in correspondence. Structure matching is a combination of Gentner's (1983) structure mapping process and a subsequent attribute-matching (O'Donoghue et al. 2006) processes. O'Donoghue et al. (2006) describe a number of categories of structure matching, however the specific requirements our of structure matching process are defined as follows:

1. Identify the maximal consistent isomorphic mapping of objects and relations between A and C.
2. Implement the 1-to-1 constraint ensuring that each object corresponds to one and only one object in the other domain.
3. Implement the relation identicality constraint, ensuring mapped relations are identical to one another.
4. Implement the identical attribute constraint, ensuring the attributes or categories of mapped objects are identical.

Steps 1–3 originate in Gentner's structure mapping theory of analogical comparisons (Gentner 1983). Although Gentner states that analogical comparisons involve few if any attributes, we shall show the importance of step 4 above for many GPAs and for interpreting topographic maps.

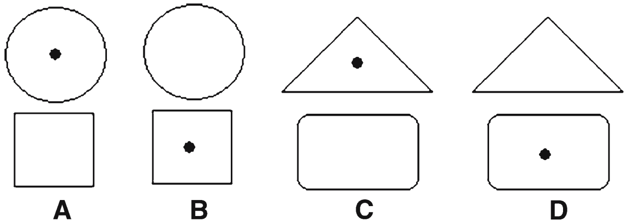

Fig. 5 GPA showing point-in-polygon classification

5 Generalised spatial problems

Next we examine how the GPA solution above is also used to solve problems in topographic maps. Like the GPA problems, topographic maps are also composed of collections of polygons, each with attribute information in the form of a category, like *road, rail, building, inland water* etc. Currently, the information that is contained in a map is generated manually and so is expensive to process. Automatic processing of this spatial data is desirable to both reduce costs and make the information available to computational processes. In the below examples the OS MasterMap data for Port Talbot (containing around 5,000 polygons) has been used in order to illustrate the solution to some spatial reasoning problems.

Both Bohan and O'Donoghue (2000) and Mulhare et al. (2001) have looked at problems involving simple collections of coloured polygons, thus we shall not concern ourselves with these problems in this paper. Our model's solution to some of these simpler problems can be found in Mullally and O'Donoghue (2005).

Instead, this paper looks at two different categories of problems that have not been previously solved by any model. The first category of novel problem that we look at concerns the use of point information in conjunction with polygons. The second category relates to complex "incremental" structures involving multiple polygons. We will now look at some common problems in topographic maps and GPAs that have not been seen before.

5.1 Point-in-polygon classification

Point-in-polygon classification is a method of using point features to classify an object (see Fig. 5).

Examples of point features in topographic maps are benchmarks, spot height, text identifying roads, names of prominent buildings and addresses. Each point feature is anchored to a point on the map and can be clearly read when looking at the map. However, text features are not directly associated with specific objects and are only useful when being read by a human.

5.1.1 Dwelling sub-categorisation using point-in-polygon

One problem associated with categorising polygons in topographic maps is finding all of the dwellings in a map. Figure 6 shows semi-detached houses and garden sheds that are all currently categorised simply as "building". One method of categorising a semi-detached house is to construct a template (Mulhare et al. 2001). Although this solution may identify every semi-detached house, it will occasionally misclassify garden sheds as they too fit the semi-detached house template.

Fig. 6 Semi-detached houses

Table 2 Confusion matrix showing the accuracy of identification strategies for Semi-D houses in Port Talbot

	P	*N*
Without address points		
T	952 (17.8%)	4388 (82%)
F	190 (3.5%)	0
With address points		
T	952 (17.8%)	4388 (82%)
F	8 (0.1%)	0

Another way of solving this problem is to look at the point features associated with dwellings. OS MasterMap has an address layer that contains an address for every dwelling (amongst some other buildings). Each address is represented as a point feature. If an address point feature is added to the original template, it is possible to remove the misclassified garden sheds from the results.

5.1.2 Results

In order to analyse our results we will look at them in terms of *precision* and *recall* (Jurafsky and Martin 2000). Precision is the ratio of the number of dwellings identified by a human to the number of dwellings correctly identified by CSM. Recall is the ratio of the number of dwellings identified by a human to the total number of dwellings identified by CSM.

Table 2 details the identification of semi-detached (semi-D) and terraced houses in the "Port Talbot" dataset from OS MasterMap. The Port Talbot data set includes 5,430 areas of which 2155 were classified as buildings. It also includes 954 text features and the corresponding address layer has 1194 address points.

The results above show that without using the point information, 182 buildings are incorrectly classified as semi-detached houses. These misclassifications are mainly garden sheds that conform to the semi-detached house template. However, adding point information in the form of address points removes these errors as garden sheds do not have addresses.

In this case eight buildings were incorrectly identified when using the address point. These were "composite buildings" that appeared to be houses as they had addresses and appeared to fit the semi-detached house template. However, even though they were misclassified as houses, they will be correctly classified later using incremental structure matching, detailed in Sect. 5.2 below.

In the case of terraced houses, 204 buildings were incorrectly identified when the address information was not used. Similarly to the semi-detached houses, most of these misclassifications were garden sheds, with the addition of composite buildings. Some terraced houses were made to look like semi-detached houses because of extensions, which are represented as separate buildings but without an address.

Table 3 Confusion matrix showing the accuracy of identification strategies for terraced houses in Port Talbot

	P	*N*
Without address points		
T	39 (0.7%)	5301 (99%)
F	204 (3.8%)	0
With address points		
T	39 (0.7%)	5301 (99%)
F	0	0

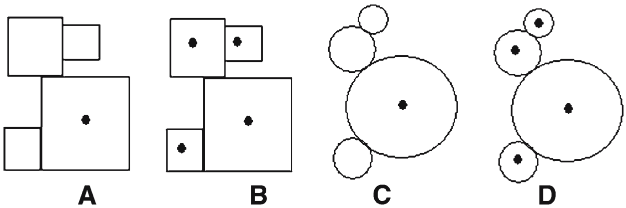

Fig. 7 GPA showing incremental structure matching

Next we examine the performance of CSM on the Basingstoke dataset, also selected from OS MasterMap. This contained just over 6,000 polygon of which 2357 were classified as buildings. It also includes almost 800 text features and 1,400 address points. As can be seen in Table 3 these results are broadly in line with the Port Talbot results.

Fifty-seven false positives were found while categorising semi-detached houses. Most of these were buildings, such as apartments, which were not specifically required to be categorised by *CSM*. Fifteen of the misclassified buildings had no address, and so could not be categorised correctly. It also means that the surrounding buildings would also be misclassified. This problem originates with the dataset and is out of our control. This problem was not evident in the "Port Talbot" dataset.

Fifty-one false positives were found when categorising the terraced houses. Twenty-nine of these false positives were caused by buildings that had no address. Similarly to the semi-D categorisations, the other false positives were caused by buildings which *CSM* was not asked to categorise.

Other possible applications of "point-in-polygon" classification include road junction identification and house number to polygon assignment. Point-in-polygon classification can also be used in order to ascertain if a road is a hill or not. By examining the height information it is possible to tell if the road has a slope and in which direction the slope is going. Similarly, this method can be used to tell the direction of flow of a river.

5.2 Incremental structure matching

Incremental structure matching is a means of assigning a better classification to a polygon once additional information is known about its surrounding area. Figure 7 shows how incremental structure matching can be used to solve a GPA problem. It consists of two parts: root identification and root elaboration (Keane et al. 1994).

An example of a problem that can be solved using incremental structure matching is identifying all of the polygons that are part of the same cluster. In Fig. 8 a cluster of buildings that form a college is described.

Table 4 Precision and recall results for Port Talbot

	# Correct	# Total	Accuracy
Precision (Semi-d)	952	960	99%
Recall (Semi-d)	952	952	100%
Precision (Terraced)	39	39	100%
Recall (Terraced)	39	39	100%

Table 5 Confusion matrix showing the accuracy of identification strategies for semi-D houses in Basingstoke

	P	*N*
Without address points		
T	566 (9.4%)	5438 (90%)
F	496 (8.2%)	0
With address points		
T	566 (9.4%)	5438 (90%)
F	57 (0.9%)	29 (0.4%)

Fig. 8 An irregular cluster of polygons forming a college

5.2.1 Incremental structure matching in topographic maps

Incremental structure matching in topographic maps allows objects such as schools, universities, hospitals and depots to be identified. These objects are irregular structures that are made up of many varied clusters of polygons. They do not have set sizes or shapes and therefore a normal template cannot be used in their identification.

5.2.2 Root identification

The first part of solving the problem is finding a polygon that is part of the cluster, known as the *root* polygon. Identifying the root means taking a closer look at point features, or more specifically, text point features. For example, the text feature may contain the word school, hospital, university, depot, church or, in this example, college. However, the text point is not always positioned inside the correct polygon, but is located in a neighbouring polygon. For example, a text point associated with a college may appear somewhere on the college grounds rather than on one of the college buildings. In this case it is necessary to assign the text to the nearest building and to ensure that the text is sufficiently close to the building, i.e. the text should be within 500 m of the building.

Table 6 Confusion matrix showing the accuracy of identification strategies for terraced houses in Basingstoke

	P	*N*
Without address points		
T	597 (9.9%)	5408 (90%)
F	594 (9.8%)	0
With address points		
T	597 (9.9%)	5408 (90%)
F	51 (0.8%)	15 (0.2%)

Table 7 Precision and recall results for Basingstoke

	# Correct	# Total	Accuracy
Precision (Semi-d)	566	623	91%
Recall (Semi-d)	566	595	95%
Precision (Terraced)	597	648	92%
Recall (Terraced)	597	612	97%

5.2.3 Root elaboration

Once the root polygon has been identified, all other building polygons belonging to that structure must be identified. However, some constraints must be satisfied during root elaboration. The first constraint is that no building being considered for inclusion in the cluster should be more than 2 km away from the root. The second constraint is that all buildings should be contiguously linked topologically with each other. If both of these constraints are satisfied then we infer that the building is part of the same composite structure as the root, and should be classified as such.

In Fig. 8 there are a total of eight buildings in that structure. Both of the above constraints are satisfied. *CSM* initially identifies all buildings that are directly adjacent to the root and assigns them to the composite structure. Then *CSM* reconsiders other buildings in the area for addition to the structure. Buildings that are adjacent to the newly classified buildings can also be re-classified. This continues until all of the buildings in the cluster have been identified.

When the address layer of a topographic map is available, this provides additional information to identify composite structures. If the text point has been assigned to the correct group of buildings, then at least one of the buildings should have the relevant text in its address.

5.2.4 Results

In the "Port Talbot" dataset, there are a total of five composite objects to be considered, made up of three schools and two depots, which were identified by a human. *CSM* successfully identifies all five of the building clusters roots. The number of buildings in each cluster ranges from 1 to 8. For every root, all of the additional buildings belonging to the cluster were found.

In the "Basingstoke" dataset there was also a total of five composite objects identified by a human. In this case, all of the composite objects were schools. The number of buildings in each cluster ranged from 1 to 15. Each of the roots was correctly identified by *CSM*. In addition to this, all of the buildings in each cluster were also identified successfully.

For the datasets above both the precision and recall are 100%. We are currently testing *CSM* on larger datasets. Preliminary results suggest that this method is effective in identifying over 90% of building clusters.

There are a number of other uses of the incremental structure matching process. Identifying all polygons that form a river can make use of incremental matching and support activities like propagating a rivers name to each individual polygon. It may also be useful in processing other extensive features like segments of road and rail networks. Incremental structure matching could also be applied to some of the categories identified above. A series of semi-detached houses (possibly with accompanying gardens and garden sheds) might be coalesced together to form a "suburbia" structure for use in generalised large-scale maps.

6 Conclusion

In the beginning of this paper it was stated that qualitative spatial reasoning could be used to solve GPAs. This was achieved by examining the spatial relationships between objects in GPAs, along with the attributes that each object has. *CSM* allows the solution to each analogy to be inferred, independent of any set of possible solutions. The solution generated is independent of size, shape and orientation.

Although *CSM* can be used to solve regular GPA problems, topographic maps were focused on in order to show how this technique could work successfully with real-world applications. Point in polygon classification and incremental structure matching were looked at in detail. Classifying houses using point in polygon classification yields results that are more accurate than template matching alone.

Incremental structure matching allows composite structures, such as schools and hospitals, to be identified. These structures could not be identified using previous methods as they are irregular and thus do not conform to a fixed template.

The use of point information in both examples is essential in order to obtain accurate results. Ignoring point information reduces the ability to successfully categorise objects in topographic maps. Thus, it appears that *CSM* provides a successful method for solving both regular GPAs and GPAs in topographic maps.

Acknowledgements We would like to thank Ordnance Survey (Great Britain), and especially Glen Hart, for their help with this project.

References

Bohan A, O'Donoghue D (2000) A model for geometric analogies using attribute matching. AICS-2000 11th artificial intelligence and cognitive science conference, NUI Galway, Ireland, 23–25 August 2000

Egenhofer MJ, Herring J (1990) Categorizing binary topological relations between regions, lines, and points in geographic databases. Technical report, national center for geographic information and analysis. University of California, Santa Barbara, 1990

Evans TG (1967) A Program for the Solution of a Class of Geometric Analogy Intelligence-Test Questions. In: Minsky M (ed) Semantic information processing. MIT Press

Falkenhainer B, Forbus KD, Gentner D (1989) The structure mapping engine: algorithm and examples. Artif Intell 41:1–63

Forbus K, Chapman V, Usher J (2003) Qualitative Spatial Reasoning about Sketch Maps. In: Proceedings of the fifteenth annual conference on innovative applications of artificial intelligence, Acapulco, Mexico, 2003

Gentner D (1983) Structure-mapping: a theoretical framework for analogy. Cognit Sci 7:155–170

Gerevini A, Renz J (2002) Combining topological and size information for spatial reasoning. Artif Intell 137:1–42

Jurafsky D, Martin J (2000) Speech and Language Processing. Prentice Hall.

Keane MT, Ledgeway T, Duff S (1994) Constraints on analogical mapping: a comparison of three models. Cognit Sci 18:287–334

Mulhare L, O'Donoghue D, Winstanley AC (2001) Analogical Structure Matching on Cartographic Data. 12th Artificial intelligence and cognitive science AICS-2001, NUI Maynooth, Ireland, 5–7 September 2001, pp. 43–53

Mullally E, O'Donoghue D (2005) Geometric proportional analogies in topographic maps. 25th SGAI international conference on innovative techniques and applications of artificial intelligence, Cambridge, England, 2005

O'Donoghue DP, Bohan A, Keane MT (2006) Seeing things: inventive reasoning with geometric analogies and topographic maps. New Generation Computing 24(3):267–288

OGC - Open GIS Consortium (1999) OpenGIS Simple Features Specification

Randell DA, Cui Z, Cohn AG (1992) A spatial logic based on regions and connection. In: Proceedings 3rd international conference on knowledge representation and reasoning, 1992

Tomai E, Forbus Kenneth D, Usher J (2004) Qualitative spatial reasoning for geometric analogies. 18th International Workshop on Qualitative Reasoning USA (during 26th Annual Meeting of the Cognitive Science Society), Northwestern University, Evanston, Illinois, 2–4 August 1992