

Seeing Things – Inventive Reasoning with Geometric Analogies and Topographic Maps

Diarmuid P. O'DONOGHUE,
Department of Computer Science/National Center for Geocomputation
National University of Ireland Maynooth,
Co. Kildare, Ireland.
diarmuid.odonoghue@nuim.ie
Amy BOHAN, Mark T. KEANE
Department of Computer Science,
University College Dublin,
Dublin 4, Ireland.
{amy.bohan, mark.keane}@ucd.ie

Submitted 2 December 2005.

Revised 1 February 2006.

Abstract This paper examines two seemingly unrelated qualitative spatial reasoning domains; geometric proportional analogies and topographic (land-cover) maps. We present a Structure Matching algorithm that combines Gentner's structure *mapping* theory with an attribute *matching* process. We use structure matching to solve geometric analogy problems that are centered on manipulating attribute information, such as colors and patterns. Structure matching is also used to creatively interpret topographic (land-cover) maps, serving to add a wealth of semantic knowledge and providing a far richer interpretation of the raw data. We return to the geometric proportional analogy problems and identify alternate attribute matching processes that are required to solve different categories of geometric proportional analogy problems. Finally, we assess some implications for computationally creative and inventive models.

Keywords: Geometric proportional analogies, attributes, topographic maps, context, creative interpretation.

1 Introduction

Analogical comparisons play a well documented role in the creative process ¹⁾. Many of the most famous breakthroughs in science can be neatly summarised by a simple analogical comparison ¹⁰⁾. For example, Kekule's famous analogy between a chain of carbon atoms and a snake biting its own tail marks the birth of aromatic chemistry. This paper uses analogical inference as its primary means of generating new information, for both geometric proportional analogies and for topographic maps.

The focus of our work is on the domain of qualitative spatial reasoning – that is reasoning about figures in a qualitative (not a quantitative) manner. We examine two seemingly unrelated problems from the domain of qualitative spatial reasoning; geometric proportional analogies ^{2, 7, 17)} and topographic (land-cover) maps ^{14, 15)}. We will show how many problems in topographic maps can be treated as geometric proportional analogy (GPA) problems – using a similar mechanism to solve both problems..

The category of GPA problem that is described involves manipulations of attribute information, such as pattern and color. We describe a *structure matching* algorithm for solving these problems that combines Gentner’s *structure mapping* theory ⁹⁾ with an attribute *matching* process. We then describe how this algorithm can be adapted to process topographic maps, where successive analogical inferences are used to gradually generate an interpretation of the map’s raw data. This interpretation more closely resembles our understanding of a semantically rich three-dimensional world.

Towards the end of this paper, we explore an entirely different category of geometric proportional analogy problem. These new problems seem to rely more heavily on lateral thinking to identify the more complicated yet subtle patterns involved. We describe how the attribute matching processes can also identify solutions to these more challenging problems.

The remainder of this paper is structured as follows. First we examine some GPA problems that involve attributes and attribute transformations. We present a *structure matching* solution to these problems, which combines Gentner’s *structure mapping* theory with a subsequent attribute *matching* process. Secondly, we describe how structure matching was used to generate creative interpretations of topographic maps. The third part of the paper returns to GPA problems, describing how some more difficult GPA problems can be solved by creatively re-interpreting the problem to find radically different solutions. Finally, we assess the implications of our work for the area of computational creativity.

2 Geometric Proportional Analogies

Geometric, proportional analogies ^{2, 7, 17)} (GPA) are comparisons between two collections of geometric figures, called the *source* and *target* domains. A GPA is of the form A:B::C:D (read as “A is-to B as C is-to D”), where A, B & C are given.. Typically the *source* domain (A:B) identifies some transformation, which must then be applied to C, yielding D.

Geometric, proportional analogies define two key pieces of information. First, the *transformation* is defined within the source domain (i.e, the change from A to B). For example, the analogy in Fig. 1 centers on inverting the polygons of part A to produce part B.

Secondly, the *inter-domain mapping* is found between parts A and C of the two domains. We map the structure ⁹⁾ of domain A to the structure of domain C. So, the square in part A maps to the circle in part C, while the circle in A corresponds to the triangle in C. It is by combing the mapping and the transformation that the missing target (D) can be generated.




Fig. 1 A Plain Geometric Proportional Analogy (GPA)

Two models of analogy by Evans ⁷⁾ and Tomai *et al* ¹⁷⁾ have also sought to solve GPA problems that involve plain figures (these problems are often recognized as the Miller IQ analogy problems). While the underlying mechanisms in these two models are significantly different, there are number of similarities between the two models. Firstly, both models take graphic images of the GPA problems as input. Secondly, both models make use of the geometry of the objects in finding a solution and thirdly, both models choose between a number of given alternative solutions (i.e. D1, D2...D5) that are presented along with the problem parts A, B and C.

Our model, called Ludi, differs from these models in a number of important ways. Firstly the input to our models is in symbolic form, rather than in the form of a graphic image. (Later we will see that topographic maps are also stored in a symbolic format, rather than as an image or picture). Secondly, we actually *generate* the required solution (D), rather than selecting the best solution from a list of alternatives. Thirdly, we ignore the geometry of the objects and therefore make our solution applicable to a wider range of problems. This third point will prove invaluable in adapting our solution to the domain of topographic maps.

2.1 The Ludi Model

This section describes a computational model for solving geometric proportional analogies, called Ludi ²⁾. Ludi is a two phase model for solving GPA problems. The first phase alone can generate solutions to plain GPA problems (like the classic Miller analogy problems) such as that depicted in Fig. 1. The second phase of Ludi is required to solve GPA problems that contain and manipulate attributes, such as color and pattern information.

The first phase of Ludi relies on Gentner's Structure Mapping Theory ⁹⁾ to identify the inter-domain mapping (between A and C). Structure Mapping Theory has been used in many models of conceptual (non-geometric) analogies ^{8, 11,, 12 13)}. Accordingly, Ludi examines the topological arrangement of objects within the presented problem. Note that Tomai *et al* ¹⁷⁾ also focus on topology, but combine this with geometric information. Davies and Goel ⁶⁾ also look at analogies involving visual information, but they do not explicitly address GPA problems.

We start by describing how Ludi represents each of the GPA problems, using a symbolic representation of the source and target domains. A unique identifier is assigned to each object in the source and target domains – as illustrated in Fig. 2. A number of binary spatial relations are used to describe the topology of each part (A, B, and C) of the GPA problem – `above(x,y)` and `inside(y,z)`. Additionally, Ludi can use the `line-adjacent(x,y)` and `point-adjacent(x,y)` relations where appropriate. (We will discuss these adjacency relations later). Typically, each part of a GPA creates a number of predicates, but in this simple example only one predicate is required. (Ludi represents shape as an attribute, like `circle(1)` and `square(2)`, but we shall leave our discussion of attributes until the next section). So, the different parts of Fig. 1 might be described by the following predicates, based on the labeling shown in Fig. 2:

```
A: above(1,2) point-adjacent(1,2)
B: above(2,1) point-adjacent(1,2)
```

C: `above(i,ii) point-adjacent(i,ii)`




Fig. 2 Labeling objects in the source and target

Next Ludi uses Gentner’s Structure Mapping Theory ⁹⁾ to generate the inter-domain mapping. This step is central to all analogical comparisons, not just GPA problems. The inter-domain mapping (or *mapping*) identifies the structural isomorphism between the descriptions of parts A and C of the GPA. Ludi also uses the predicate identity constraint ⁸⁾, which ensures that only identically named relations can be mapped between A and C. In this simple example, the relations in domains A and C can be mapped directly to one another – allowing us to identify the object mapping.

Mapping: `above`-`above`, `point-adjacent` - `point-adjacent`, `1-i`, `2-ii`

Therefore object 1 maps to `i`, while object 2 maps to `ii`. Because the mapping process concentrates on the structure of parts A and C (i.e. the mapping is based on the `above` and `point-adjacent` relations), non-identically shaped polygons can be mapped to one another. In this way, square 1 of part A is mapped to circle `i` of part C. Similarly, circle 2 in A is mapped to triangle `ii` in part C.

Of course, the source domain also includes the crucial A-to-B *transformation*. The A-to-B transformation is represented implicitly within the Ludi model. That is, the transformation is represented by the difference in the collections of predicates that describe parts A and B of the GPA. Unlike Davies and Goel ⁶⁾ we do not use a separate vocabulary to represent transformations. So, the A-to-B transformation in Fig. 1 involves a reorganization of the objects in part A.

A: `above(1,2) point-adjacent(1,2)` -> B: `above(2,1) point-adjacent(1,2)`

Armed with both the mapping and the transformation, we are now ready to solve the GPA. The correct solution (D) is generated by the standard “pattern completion” algorithm for analogical inference, called Copy With Substitution and Generation ¹¹⁾ (CWSG). Additional source domain information (i.e. part B) is copied to the target domain (forming D), with all source domain labels being substituted by their target domain equivalents - obtained directly from the mapping. Therefore, we copy B to the target forming `above(2,1)`, but we now substitute each label with its mapped equivalent (`1-i` and `2-ii`). This generates the following solution, as illustrated in Fig. 3.

D: `above(ii,i) point-adjacent(i,ii)`



Fig. 3 Solution to the Geometric Proportional Analogy in Fig. 1

Larger GPA problems involve more predicates and objects but are solved by the same sequence of steps.

2.2 Adding Attributes to Geometric Proportional Analogies

We now move onto GPA's that involve attributes and attribute transformations (see Fig. 4). We point out that including attributes in the analogy process is central to solving these problems. While both Evans⁷⁾ and Tomai *et al*'s¹⁷⁾ models solve plain GPA problems, they do not address the GPA problems described in the remainder of this paper.




Fig. 4 A Geometric Proportional Analogy with Attributes

To solve these GPA problems, we add the attribute information about each polygon to the earlier predicate descriptions. For simplicity, we focus on the pattern attribute. The analogy in Fig. 4 is represented by the same predicates as in Section 2.2, plus the following attributes:

```
A: striped(1) plain(2)
B: striped(1) gray(2)
C: striped(i) plain(ii)
```

The addition of this attribute information complicates the analogy process because there are multiple ways of identifying the mapping that occurs between the source and target attribute transformations. We define *attribute matching* as the process of determining the attribute changes in the transformation and mapping process.

Of course we still identify the mapping between the source and target domains, but this process initially ignores the attribute information. We use the inter-domain mapping to identify the object correspondences which then allows us to align the attributes of the objects in A with the objects in C.

2.3 Identical Attribute Matches (Object Independent)

We begin with the simplest class of GPA problem that uses the same attribute matches throughout the analogy problem. We use the object correspondence to identify a 1-to-1 matching between the attributes in A with those in C. In the GPA in Fig.4 the `striped` attribute of object 1 matches to the `striped` attribute of object i, and similarly the `plain` attribute of object 2 matches the `plain` attribute of object ii. In this case all striped attributes in A match to striped attributes in C, forming what we call an object independent attribute match.

Attribute Match: striped-striped, plain-plain

In fact, in this problem all matched attributes are identical to one another, conforming to the *identical attribute matching* constraint. These are the simplest type of GPA problem.

This generates a complete correspondence between the two domains, incorporating predicate, object and attribute information.

A: striped(1) plain(2) -> B: striped(1) gray(2)

To generate the inferences for this problem, we again use the CWSG algorithm. (First, CWSG is performed yielding the same results as in the last section). Now we include attribute information in each of the parts A, B and C. The attributes of B are copied to D, being appropriately substituted by their mapped equivalents (1-i, 2-ii). This adds the following attributes to the description of part D.

D: striped(i), gray(ii)



Fig. 5 Solution to the Geometric Analogy in Fig. 1

The identical attribute constraint will prove crucial in our efforts to automatically interpret topographic maps. This is also the type of attribute matching that is required by many geometric proportional analogy problems that involve attributes.

2.4 Non-Identical Attribute Matches

The Ludi model can also solve GPA problems where non-identical attributes are matched between the source and target domains. Fig. 6 depicts a GPA problem where the attributes of the mapped objects are *not* identical.




Fig. 6 A Geometric Analogy with Non-Identical Matched Attributes

Again we use the inter-domain mapping to identify the object correspondence. Then we align the compatible attributes of mapped objects.

So, the striped attribute of part A is matched with the hashed attribute in part C.

Attribute Match: striped-hashed, plain-plain

Again we perform inference using the CWSG algorithm. We create a copy of B to form D and then substitute all mapped and matched items as appropriate. This substitution process replaces the striped attribute in B with the hashed attribute in D. Thus our GPA can create a solution D that contains attribute information which is not found in the source.



Fig. 7 Solution of Fig. 6

Analogies, including creative analogies, rely on this same combination of mapping with inference to create new information. Our extension to the “pattern completion” inference process allows inferences to include attribute information that describes objects in the analogy. In the next section we will see how inferred attributes can transform our understanding of spatial data from a topographic map.

3 Topographic Analogies

In this section we examine another domain that also contains spatially related information. This domain concerns topographic maps that describe land-cover across a country. Topographic maps are formed from large numbers of polygons that each represents features like *roads*, *gardens* and *buildings*. While there is no obvious connection between Geometric Proportional Analogies and topographic maps, Fig. 8 depicts a GPA where the target domain is a collection of polygons selected from a topographic map^{14, 15}. We refer to analogies involving geometric figures and a topographic map as *topographic analogies*. The transformation defined in the source domain of Fig. 8 adds a polygon to the collection of polygons in part A. This polygon-adding transformation is then applied to the topographic map (C) enhancing and updating that map (D). In this case the transformation adds a segment of the *river* that was hidden by an overlying foot-bridge, thereby repairing the occlusion.




Fig. 8 A Topographic Analogy - a Geometric Proportional Analogy in a Topographic Map

3.1 Reading Topographic Maps

Before we show how the structure matching process was adapted for topographic analogies, we must first describe how information is recorded in topographic maps. At present, most maps of the world have been transformed into digital media opening up new vistas for the flexible use of their contents by business, government and private individuals. Maintaining and enhancing such topographic data presents us with a new set of challenges that could benefit from automated solutions.

The topographic map used in this paper is the Topological Map of Great Britain, called OS MasterMap¹. OS MasterMap contains over 450 million (mostly) non-overlapping polygons, describing land-cover across the country (Fig. 9). Each polygon is composed of a number of lines that define the boundary between polygons. Each recorded feature (line, polygon etc) is uniquely identified by a 16 digit Topographic-ID (ToID) number and the entire map contains over 2.5 billion ToIDs. Polygons are categorized into one of thirteen main themes, including; *road, rail, building, inland waterway, made land, unmade land, roadside etc.*



Fig. 9 A Small Segment of the Topographic Map of Port Talbot (UK)

Topographic maps are primarily devoted to recording location information, with little semantic information being added. Of course much semantic information is difficult to discern from the aerial photographs used to create maps. Topographic maps typically do not distinguish between say, a *hospital* and a *garden shed*. This lack of semantic information can result in a number of problems with these maps. We now look at some enhancements that might serve to improve the usefulness of this spatially based data.

3.1 Interpreting Topographic Maps

We make a number of observations about how people *interpret* topographic maps, applying a modicum of creativity to “see” more information than can be explained by a strict reading of the map’s data. For example, map users regularly encounter the occluded-polygon problem (Fig. 8) wherever, say a *river* passes beneath a bridge. This is because the bridge ensures that part of the *river* cannot be detected from the aerial photographs that are used to

¹ Crown Copyright, Ordnance Survey, Southampton UK, 2005.

make maps. Therefore the *river* appears to be obstructed by the *bridge*. But even novice map users have no difficulty in imagining that the *river* continues beneath the *bridge*. We theorize that in the mind of a map user, the map is not segmented into collections of isolated polygons but is more likely to be composed of continuous *roads* and *rivers* that occasionally traverse one another. However, current geocomputation applications do not support such simple inferences (route-planning and GPS navigation applications typically use a separate “Transportation Network” layer which is created using human intervention).

Similarly, when examining a map that depicts polygons of the *building* category, many of these structures are easily identified as (semi-detached) *houses* – regularly structured with a *garden* and access to a *road* (see Fig. 9). So, while the exact type of *building* is not recorded on the map, the users understanding of the map is often much richer than the recorded data. Our CSM model aims to enrich the map’s recorded data, so that it more closely resembles the richer information that people perceive.

One can view this as a relatively simple form of inventive analogy, where users apply their understanding of the real world to the data presented on the map. In this scenario the real world acts as the source domain, which is used to transfer additional information to the targeted items in the topographic map. This creates an interpretation of the map that is much more useful than a strict reading of the recorded data.

Users also apply their understanding of the map to enhance their understanding of the real world, when using the map for activities like route planning. That is, a route might be planned on the map, but that plan is then executed in the real world. So, we also use the map as the source domain to supply additional information to some real-world problem. However in this paper, the source domain of the real world is used to identify a number of transformations, which enhance and enrich the topographic map.

3.2 Semantic Enrichment for Virtual Reality Maps

Generating a virtual reality representation of topographic maps has long been a goal of many researchers. Virtual reality maps may be useful in desktop virtual environments, as realistic environments for the software gaming industry or as aids to mobile navigation. However, generating virtual reality maps involves using far more information than is currently available in topographic maps. While maps contain adequate location information, it is their lack of semantic information that presents a serious challenge.

We highlight a number of requirements for inventive interpretation of topographic maps using the image in Fig. 10. This depicts a virtual reality environment containing a small portion of a typical topographic map. This image focuses on the intersection of two *roads* and a number of *buildings* and *multiple-surface* polygons. On the left we have a direct representation of the topographic map that represents a building – one flat *building* polygon surrounded by a *multiple-surface* land polygon. While this is a true depiction of the map’s data, it does not depict the semantic and visual richness contained in the real world.

A slightly more inventive interpretation of the map might allow us infer that the building is a single-story *detached-house*. We can then use this inference to estimate the house’s height and to infer that the surrounding polygon is a garden. This situation is depicted

on the left of Fig. 10 where the building is repressed as a plain 3D box. However this is only a minor advance but it does represent an improved depiction of these objects.



Fig. 10 A possible virtual reality map. The left foreground shows an unadorned building, while the right shows a more creative interpretation of a map.

The building on the right of Fig. 10 represents a more inventive interpretation of a *building* enclosed by *multiple-surface* land. Some steps in producing this interpretation are:

- Decide the number of stories and height of the building
- Determine the type of roof for the dwelling
- Decide the orientation of the dwelling and the location of the front door
- Decide the type of garden to depict
- Identify the building located in the garden as a garden shed.
- Select the type of garden shed to depict
- Decide the types of boundary to place on each side of the garden
- Decide the location of any gaps in these boundaries

If a person were presented with this challenge and produced a suitably pleasing result, it could be argued that they had displayed a measure of creativity in reaching that goal. It might be argued (by analogy) that if a program produced a similar result, it may display some traces of creative reasoning. We point out that this interpretation is useful, novel and was created in a directed manner – we shall return to these facets of creativity in section 6.

Our CSM algorithm is still not at the stage of generating such a cohesive interpretation of the world. However, we shall look at a number of specific problems that are being addressed with CSM and which assist in generating a virtual reality map. The first problem we examine concerns occluded polygons, like that depicted in part C of Fig. 8. We will also look at the problem of sub-classifying polygons into more specific sub-categories and at identifying classification errors from within topographic maps.

4 Contextual Structure Matching (CSM)

In this section we describe the Contextual Structure Matching ^{14, 15)} (CSM) algorithm for processing topographic maps. CSM is an adaptation of the structure matching algorithm described in Section 2, specifically tailored to manipulate topographic maps. CSM uses geometric analogies as a basis for identifying specific problem situations within topographic maps.

4.1 Describing the Problem Context

The key to using topographic analogies like that in Fig. 8, is to describe the local context of the polygons involved in this bridge occlusion problem. GPA's are typically composed of parts A, B and C that consist of small collections of polygons. To apply these analogies within an extensive map, we must therefore identify a small collection of topographic polygons on which to apply our GPA algorithm. We introduce a new level of resolution for dealing with topographic maps, which we call a *locality*. A *locality* is a collection of polygons consisting of one central polygon, plus all polygons that are immediately adjacent to that central polygon. Thus, the entire map is covered by numerous over-lapping localities.

Each locality is described by three types of information. Firstly, localities contain a unique identifier (ToID) for each polygon in that locality. Secondly, each locality records the category of all polygons that it contains. Finally each locality records the topology of the polygons in that locality, using the `line-adjacent` and `point-adjacent` predicates. The `line-adjacent` and `point-adjacent` relations can be seen as specializations of the `externally-connected-to` relation that is used by the Region Connection Calculus ⁴⁾ RCC-8. These two relations are used because they allow CSM to explicitly identify 2-dimensional situations where polygons share only one vertex, which are crucial to identifying many structures. Such relations are easily identified from the topographic map and allow CSM to easily identify specific problem structures. CSM also assumes that the description of each part of the problem is complete - that there are not missing relations. Thus, the RCC-8 relation `disconnected-from` is represented implicitly in CSM by the absence of a relation between two objects within the locality.

The representation of each locality creates a collection of predicates and attributes similar to that used by Ludi, described in section 2. The main difference between CSM and Ludi concerns the relations they use to represent problems. CSM uses the `line-adjacent` and `point-adjacent` relations, which are insensitive to the rotation (or orientation) of the polygons and allow it to identify structures in a rotation independent manner. In contrast, Ludi relations are sensitive to orientation (eg above) because orientation is often crucial in these problems.

4.2 System Architecture

A specific problem, such as the occluded bridge structure, is given to CSM as part A of a GPA. CSM must then find a matching C locality within the map. The topographic map is loaded into CSM which generates all required locality descriptions for the given map. These localities are then passed to the structure matching component of CSM, which compares each

locality description against targeted locality descriptions. When a structure match is detected and when some simple spatial constraints are met, then a problematic locality has been identified and is dispatched to the appropriate elaboration process.




Fig. 11 System Architecture

CSM's spatial constraints help eliminate many "false positives" that do not properly fit the locality description. These constraints augment the topological structure matching process with some geometric conditions. For example when detecting the occlusion problems, CSM uses a distance metric to ensure the occluding bridge is not wider than some threshold.

4.3 Some Occlusion Problems

We now describe how CSM repairs seven specific occlusions, like that described in Fig. 8. These problems originate in occlusions of three categories of topographic object - *roads*, *rivers* and *railways*. We create a locality description for each of the following structures and these seven localities correspond to seven distinct GPA problems. These localities form a simple case-base that identifies these specific problem structures within the topographic map.

- 1) *Road over River*
- 2) *Structure over River*
- 3) *Rail over River*
- 4) *Road over Road*
- 5) *Rail over Road*
- 6) *Path+road+path over River*
- 7) *Path+road+path over Rail*

4.4 Problem Correction

Having identified the structures of interest, we can then begin elaborating the map. Elaboration corresponds to the A-to-B transformation of the earlier GPA's. However, the corresponding part B for each of the targeted localities is formed by an inference process, rather than a predefined collection of predicates (as used by Ludi).

The elaboration needed to repair each occlusion, consist of generating and re-inserting polygons of the appropriate categories corresponding to the occluded portion of the underlying object. In this case the elaboration processes is quite straightforward, partly because CSM also identifies both the occluding and occluded polygons. Elaboration uses this information to identify the intersection between the occluding and occluded objects. By adding two additional edges we can reform the obscured polygon. Of course this process makes assumptions about the nature of the obscured polygon, using straight lines to represent the hidden edges of the occlusion. This has proven adequate for most purposes, rejoining the previously partitioned object (Fig. 12).




Fig. 12 Before and After Re-Inserting one Occluded Polygon

4.5 Empirical Results from CSM

We tested CSM’s occlusion repair process on three different subsets of OS MasterMap, from the regions of *Moffat*, *Port Talbot* and *Birmingham*. Each map represents a region of a few square kilometers, containing a total of over 43,000 polygons. The represented regions are characterized as: urban, suburban, industrial, parkland rural, and mountainous (see Table 1), presenting different challenges to the CSM application.

Table 1 Results summary for the different map regions

Region	Description	Number of polygons	Identified Obscurities
Moffat	Mountain, rural, town	11,293	47
Port Talbot	Rural, suburban, industrial	5,198	10
Birmingham	Urban, suburban, parkland	26,632	14

These maps were also presented to two human reviewers who were asked to interpret the map and manually identify all occluding bridge structures. The reviewers were given a printout of the map on which to locate the bridges, as well as having access to a map viewing application program. This allowed reviewers, for example, to view only the *road* polygons so that it might simplify the task of identifying bridges. Additionally, the reviewers were given as much time as required to identify all occluding bridges.

As can be seen from Table 2 CSM was very accurate, identifying and repairing 100% of the occluded polygons in many situations. Additionally, in each case the newly inserted polygon had the correct location, dimensions and was assigned to the correct category.

These results demonstrate the potential for qualitative spatial reasoning within topographic maps. While these seven cases generate straight-forward inferences, we now describe a number of other uses of CSM to enrich the semantics of topographic maps.

Table 2 Accuracy of Occluded-polygon Insertion

Problem	Topographic Database		
	Moffat	Port Talbot	Birmingham
	Accuracy%	Accuracy%	Accuracy%
Road over River	85	66	0
Structure over River	100	100	100
Rail over River	100	100	100
Road over River	66	100	100
Rail over Road	50	100	100
Path+road+path over River	58	22	25
Path+road+path over Rail	100	70	100

4.6 Emerging Applications

Our geometric analogy solution is also being used to address several other topographic problems, which we will now outline.

4.6.1 Identifying and Correcting Mis-Classification

Correctly classifying all polygons in a topographic database is a major concern as it directly impacts on the usefulness of maps to the end user. Automated classifiers focus on the description of individual objects, focusing on their size, border length, geometry *etc.* But some polygons are ambiguously shaped and can be prone to mis-classification – *road junctions* and *semi-detached-houses* are very similar in shape (see Fig. 13). Correcting classification errors can use CSM’s contextual sensitivity to examine the class of surrounding polygons. Appropriate locality descriptions can detect many misclassifications errors.



Fig. 13 A variety of polygons representing buildings, roads and road intersections

4.6.2 Sub-Categorization

Currently OS MasterMap categorizes polygons into one of 13 different classes. We are using CSM to sub-categorize polygons into a greater number of sub-classes.

- 1) *building* -> *dwelling* -> {*detached-residence, semidetached-residence, terraced-residence*}
- 2) *building* -> *garden-shed*
- 3) *road* -> {*motorway, road-segment, cul-de-sac, T-junction, X-junction*}

Multiple hierarchical levels are being explored for these and other categories.

4.6.3 Composite Objects

Composite objects consist of collections of individual polygons. For example a *homestead* may consist of a *dwelling* with its surrounding *garden* and *out-houses*. Extensive features like *rivers* and *roads* consist of many individual polygons and propagating (say) the name of a *river* to all its segments can improve the maps usefulness.

4.7 Cascade of Inferences

The above manipulations raise a number of important issues, particularly when these tasks are taken together as a collection. Identifying a *building* polygon as a *detached-house* can itself be used as a basis for identifying its surrounding *multiple-surface* polygon as a *garden*. Thus, one inference is used as a basis for generating a subsequent inference. This second inference in turn helps identify that the *building* located in the garden is a *garden-shed*. Thus, rather than infer all our conclusions from known facts, we use some inferences as a basis for subsequent inferences. We call this a *cascade of inferences* and see it as crucial in our efforts to go from the current dearth of semantics, to the full richness required to generate a virtual reality map. Of course there is potential for early mistakes to be compounded by subsequent inferences. This places the greatest need for accuracy on the first inferences that are derived directly from the maps original data.

5 Ludi: Flexible Structure Matching

We now return to the topic of geometric proportional analogies to highlight some crucial differences between our CSM and Ludi models. The structure matching algorithm that underlies both Ludi and CSM solves “normal” GPA problems. Both Ludi and CSM generate the same solutions as required by most IQ test problems, and as also solved by models like those of Evans⁷⁾ and Tomai *et al*¹⁷⁾. We now describe some other classes of GPA problems involving attributes that can also be solved by Ludi. The two classes of problem described in this section require greater flexibility and inventiveness than we have seen so far. (Neither Evans nor Tomai *et al* address these types of problem).

Let us begin by briefly re-considering the original GPA problem depicted in Fig. 1. A more creative reading of this problem might suggest that the source domain merely replaces A by the description contained in part B – thereby foregoing any notion of an inter-domain mapping. Applying this (trivial) transformation to C would generate a D that is merely a duplicate of part B. Another (non-)solution to GPA problems might suggest that the source domain transformation can only be applied to figures that exactly match part A, and thus no solution should be given to the GPA in Fig. 1. We pass over these trivial interpretations, as do^{7, 17)}, to look at some more serious and challenging problems.

5.1 Object Dependent Attribute Matching

Moving on from these trivial cases, we present a category of GPA problem that cannot be solved by the structure matching process described earlier. This process identified a single transformation for each attribute in A. The analogy depicted in Fig. 14 requires additional flexibility in the attribute matching process, and can not be solved by the mechanisms we have described so far. The source domain (A and B) identifies two possible transformations

for the striped attribute – transforming to both striped and dotted for different objects.




Fig. 14 Local attribute matches

To generate the correct solution to this problem, Ludi must deal with the attribute transformations for each object independently from the others. So for one object in the target the striped \rightarrow striped transformation must be applied, while for the other object, the striped \rightarrow dotted transformation is applied (see Table 3).

Table 3 Different Attribute Matching Strategies

<p>Object Independent Matching (Figure 4) striped \rightarrow striped</p>	<p>Object Dependent Matching (Figure 14) striped(1) \rightarrow striped(i) striped(3) \rightarrow dotted(iii)</p>
---	--

Some of the attributes in Fig. 14 could be addressed using the object independent attribute matching algorithm. But Ludi’s object-dependent attribute matching process can also solve GPA problems where the attributes can not be matched using the simpler technique.




Fig. 15 The solution to Fig. 14

5.2 Object Independent Attribute Matching

Fig. 16 depicts a GPA problem that requires a far greater degree of lateral thinking to identify a possible solution. In solving this problem, Ludi interprets the problem in a radically different manner from any of the previously described problems. The feature that distinguishes Fig. 16 emerges when we examine the square in parts A and B which is mapped to the hexagon in part C. All three of these figures have different attributes, and thus Ludi recognizes that it can not solve this problem in the usual manner. This is because the analogy does not appear to have the necessary similarity between the source and target domains that is necessary for applying its usual approach. However, Ludi does not abandon its efforts but explores another possible interpretation of the GPA problem.

In all previous analogical comparisons, the attribute matching was derived from the object mapping. However for this problem, the attribute matching is performed independently from the object mapping generated according to Gentner’s structure mapping theory⁹⁾. The attribute transformation defined by an object pair in the source domain, can therefore be applied to objects other than the mapped equivalents of these objects in the target domain.




Fig. 16 The attribute transformation (A->B) is not applied to the mapped objects in C.

In Fig. 16, all of the attributes are treated independently from the objects to which they are attached. The source domain includes the following transformation `striped(1) -> plain(1)`. Additionally, the structure mapping process places the following objects in correspondence (1-i). But because this transformation can not be applied to its mapped equivalents (object i), Ludi finds an alternate use for this transformation by applying it to another target object. Ludi applies this transformation to object ii, because this has the same `striped` attribute as used by the transformation. Thus Ludi applies this transformation (`striped(1) -> plain(1)`) to object ii, creating a D with the attribute `plain(ii)`. The full solution generate by Ludi is depicted in Fig. 17.




Fig. 17 Solution to Fig. 16

This represents a significant extension to the structure matching capabilities introduced at the beginning, as used in the CSM algorithm.

6 Implications for Computational Creativity

Reasoning with analogies plays a well-known role in creativity ¹⁾. Typically, a novel source domain offers a new interpretation of some target problem, offering novel inferences and explanation to expand our understanding of that problem. For example, Kekule used the source domain of a snake biting its own tail to some problematic carbon based molecules. The formation of a circular structure in the source domain prompted the identification of circular structures in many carbon compounds. This creative analogy helped explain the structure of benzene, whose structure and behavior was not explained by previous theories.

This paper presented two models based on analogical inference and we shall examine each from the perspectives of inventiveness and creativity. First we look at the Contextual Structure Matching (CSM) model for reasoning about topographic maps. Maps may not appear to be a domain open to inventiveness or creativity, but Kent Becks's famous 1933 map is often cited as a creative and artistic representation of London's underground system. Interestingly, Beck achieved this by reducing the quantity of information presented to the

map reader. In contrast, the CSM model aims to enhance the information stored in the map, inventing a semantically rich interpretation of the raw map.

While this model is still in its relatively early stages, CSM offers a perspective on how a rich full-featured interpretation of the world might be invented that corresponds to a relatively simple topographic map. CSM uses the real world as a source domain to expand the information contained in the topographic map. The approach fits closely with an experienced map user imagining some virtual world, which corresponds to a given map. Results on the occlusion problem favorably compared the performance of CSM with users who were asked to *interpret* a given map.

CSM goes beyond traditional inference mechanisms, allowing new information to be “invented” for which there may be limited support. CSM’s inventiveness is based on the *cascade of inferences* technique, where the results produced by one inference process are used as input to subsequent inference activities. So categorizing one *building* as a *house* allows another *building* to be categorized as a *shed*. These inferences might then allow the front of the house to be identified, helping to identifying any graphic images to be applied within a virtual reality depiction of the world. Thus, CSM gradually constructs an inventive interpretation of a topographic map.

CSM displays some characteristics that are often associated with creativity, such as being directed, novel and useful ¹⁶⁾. CSM is directed, being driven by specific goals such as repairing occlusions. It is novel in that its inferences add new information to the sparse information previously recorded on the map. Finally, it is useful because it enables information on a map to be elaborated, adding inferences to the recorded data. Colton and Steel ⁵⁾ identify “justification of findings” as a key requirement for computational creativity.

However, CSM is limited in the types of pattern that it can detect. Ludi displays far greater flexibility in the types of patterns that can be identified between the source and target of the GPA. A complex category of GPA problem appears to rely on more lateral thinking, where the attribute transformation does not “conform” to the object correspondence identified by the inter-domain mapping. It is often argued that this facet is an important quality of creativity ³⁾. Thus, any inventiveness attributed to Ludi must relate to its ability to identify an entirely different range of patterns, which the CSM model does not detect (or require). Thus, Ludi may offer greater potential for exploring computational creativity.

7 Conclusions

Analogy plays a central role in many computational models of creative reasoning. We explore the use of geometric proportional analogies (GPA) – like those found in IQ-tests. One category of such problems can be solved using Gentner’s standard structure mapping model of analogical reasoning, which identifies a mapping between the objects in the source and target domains. However, structure mapping alone does not solve “painted” GPA’s, whose objects have attributes like colors and patterns. We solve these problems by extending the structure mapping model with a subsequent attribute matching process. One class of painted GPA’s relies on matching identical attributes between mapped objects.

We use this painted GPA approach to interpret 2D topographic land-cover maps, creatively envisioning possible 3D worlds that correspond to some given map. Our analogy-based solution (called CSM) relies heavily on its attribute matching process, which is central in solving these problems. CSM also allows attributes to be manipulated within these geometric analogies, as required to solve many geometric proportional analogy problems. Several specific limitations in land-cover maps that require such creative interpretation are discussed. One such problem relates to the identification and restoration of bridge-related occlusions. Results for this bridge identification problem are presented.

CSM uses a *cascade of inferences* technique to gradually create a possible world corresponding to the given map. But this still relies on matching objects with identical attributes. Not all painted GPA problems can be solved with identical attribute matching – requiring non-identical attribute matching. These more complex problems require a greater degree of insight and creative interpretation than the identical attribute problems. The Ludi model solves both simple and complex geometric proportional analogies, by also aligning non-identical attributes between the source and target domains. While identical attribute matching supports some forms of creative reasoning, non-identical attribute matching seems to open up new vistas for creative analogizing. Even more creative modes of analogical reasoning may become possible, using this more flexible approach to using analogies that involve attribute information.

Acknowledgements

The authors would like to thank the following for their assistance and support: Glen Hart, Ordnance Survey Research Center, UK; The British Council; Dr. Adam Winstanley, Alan Loughlin, and Leo Mulhare, Dept. Computer Science, N. U. I., Maynooth, Ireland.

References

- 1) Boden, M. A. “The Creative Mind”, Abacus, 1992.
- 2) Bohan, A. O'Donoghue D. “A Model for Geometric Analogies using Attribute Matching”, 11th Artificial Intelligence and Cognitive Science Conference, Aug. 23-25, NUI Galway, Ireland, pp 110-119, 2000.
- 3) Buchannan. B.G. “Creativity at the Metalevel”, AI Magazine, 22, 3, pp 13-28, 2000.
- 4) Cohn, A.G. Bennett, B. Gooday, J. Gotts, N.M. “Qualitative Spatial Representation and Reasoning with the Region Connection Calculus”, GeoInformatica, 1, 3, pp 275-319, 1997.
- 5) Colton, S. and Steel, G. “Artificial Intelligence and Scientific Creativity”, Artificial Intelligence and the Study of Behaviour Quarterly, Vol. 102, 1999.
- 6) Davies, J. Goel, A. “Visual Analogy in Problem Solving”, Proc. of IJCAI, August 4th - 10th, Seattle, Washington, USA, pp 377-382, 2001.
- 7) Evans, T.G “A Program for the Solution of a Class of Geometric Analogy Intelligence-Test Questions”, in Semantic Information Processing (Ed.), M. Minsky, MIT Press, 1967.

- 8) Falkenhainer, B. Forbus, K. Gentner, D. "The Structure Mapping Engine: Algorithm and Examples", *Artificial Intelligence*, 41, 1-63, 1989.
- 9) Gentner, D. "Structure-Mapping: A Theoretical Framework for Analogy", *Cognitive Science*, 7, 155-170, 1983.
- 10) Hoffman, R. "Monster Analogies", *AI-Magazine*, 3, 11-35, 1995.
- 11) Holyoak, K.J. Novick, L. Melz E. "Component Processes in Analogical Transfer", in *Analogy, Metaphor and Reminding*, Barnden, Holyoak (Eds.) Ablex Norwood, N.J. 113-180, 1994.
- 12) Keane, M.T. "Incremental Analogising: Theory and Model" in *Lines of Thinking* (Vol. 1), Gilhooly, K.T. Keane, M.T. Logie R. & Erdos G. (Eds.) New York: Wiley, 1990.
- 13) Keane, M.T. Ledgeway, T. Duff, S. "Constraints on Analogical Mapping: A Comparison of Three Models", *Cognitive Science*, 18, 387-438, 1994.
- 14) O'Donoghue, D. A.C. Winstanley, "Finding Analogous Structures in Cartographic Data", 4th AGILE Conference on GIS in Europe, pp 553-560, Brno, Czech Republic, April, 2001.
- 15) O'Donoghue, D. Winstanley, A.C. Mulhare, L. and Keyes L. "Applications of Cartographic Structure Matching", *IEEE - IGARSS Conf.*, July 21-25, Toulouse, France, 2003.
- 16) Ritchie, G. "Assessing Creativity", *Proceedings of AISB Symposium on AI and Creativity*, pp 3-11, York, England, March, 2001.
- 17) Tomai, E. Forbus K. Usher, J. "Qualitative Spatial Reasoning for Geometric Analogies", *Proc. 18th International Qualitative Reasoning Workshop*, Evaston Illinois, USA, August 2 - 4, 2004.



Diarmuid O'Donoghue He received his B.Sc. and M.Sc. from University College Cork in 1988 and 1990, and his PhD from University College Dublin. He has been a lecturer at the Department of Computer Science NUI Maynooth since 1996 and is also an associate of the National Centre for Geocomputation. His interests are in artificial intelligence, analogical reasoning, topology, and qualitative spatial reasoning.