

APPLYING INTERPRETATION-DRIVEN ASSOCIATION TO DESIGN DOMAINS

KAZJON S. GRACE¹, ROB SAUNDERS² and JOHN S. GERO³

University of Sydney, Sydney, Australia,

1. kazjon@arch.usyd.edu.au, 2. rob.saunders@sydney.edu.au

3. George Mason University, Fairfax, USA, john@johngero.com

Abstract. This paper presents a computational model of visual association-making. Our model focuses on the interaction between the processes of representation and matching in association. Re-interpretation of the objects being associated changes the landscape in which the matching process is searching for potential mappings between those objects. We call this process interpretation-driven search. We demonstrate the capabilities of our system through some examples of previous work in simple shape domains, then discuss ongoing research into applying this system to design domains.

Keywords. Visual association; interpretation; computational model; design cognition.

1. Introduction

The process of association is the construction of a mapping between shared aspects of two objects. This fundamental cognitive process underlies analogical reasoning and related processes such as metaphor, allegory and case-based reasoning. Each of these processes is of interest to design researchers, with analogy in particular being recognised as an important component of creative design (Goel, 1997). Modelling association independent of these processes allows us to focus on the way object representations are constructed, which has been identified by a number of analogy researchers as an important area of focus for computational models (French, 2002; Kokinov, 1998). This paper presents a computational model for developing associations founded on a cognitive approach. The paper demonstrates an implementation of this model and discusses the utility of such a system in understanding and assisting designing.

We understand association as being composed of three processes: *representation* of the objects, *matching* between the objects and building a *mapping* around that match. If these processes are modelled serially a chicken-or-egg problem arises. How does the representation that encodes a particular mapping arise before the mapping itself? Representation must occur in parallel with matching and mapping, so that the way objects are conceptualised evolves in conjunction with how they are mapped. We have developed a model of association, exemplified in the visual domain, that focuses on the iterative interaction between the search for mappings and the construction of representations, an interaction that we call interpretation-driven search.

This exemplary model focuses on visual associations, i.e. associations between the structure of visual elements within objects rather than the causal and functional relationships between those elements. Visual analogy has been identified as an important domain in design, creativity and problem solving (Davies and Goel, 2001). The system presented here forms associations between collections of line shapes like those that might be found in sketches, plans or visual design elements. Our intent for this ongoing research project is to apply our system to real-world design domains.

Computational models of association and analogy have potential applications in both understanding how human designers use those processes and in design support systems research. Realising this potential represents a significant challenge in developing systems with the robustness and portability necessary to produce meaningful results in a variety of design domains.

2. A computational model for interpretation-driven association

The model described in this paper can be decomposed into three interacting processes, seen in Figure 1. *Perception* is the system that describes objects it encounters, *mapping* is the system that relates those objects; and *interpretation* is the system that changes the descriptions of the objects.




Figure 1. The structure of the interpretation-driven association system.

Low-level representations of the two objects are provided to the perception process, which discovers features within those objects and categorises them into concepts, creating new conceptual categories where necessary. In the implementation described here, features are line shapes and the low-level representations provided are vector images. The perception process also extracts topological and typological relationships that exist within each object from the features and concepts it has created. Each object is represented by a graph of its features and the relationships between those features. The mapping process then searches the two graphs for identical patterns of relationships between features. Whilst this search is occurring the Interpretation process can change the structure or content of the graphs. Interpretation and mapping iterate until a suitable mapping is found.

2.1 PERCEPTION.

The perception process forms the representations of objects that our system uses. These representations form a baseline from which interpretations are applied during its search for mappings. In the implementation described in this paper representations are composed of shapes and of relationships between those shapes. Shapes are individual visual features detected, described and categorised by the system, and the relationships between them are both typological and topological. This element-and-relationship representational structure is consistent both with work in analogy-making (Gentner, 1983) and design thinking (Gero, 1990).

The perceptual processes used in the system are influenced by previous objects the system has seen. Shapes are categorised based on similarities in their descriptions, and these categories, which we call “concepts” are generated and populated at run-time. A shape that is unlike any previously encountered will generate a new conceptual category, and future shapes judged sufficiently similar will be added to that category. This concept formation is accomplished using an unsupervised clustering algorithm operating on the descriptions of shapes. The purpose in including these constructive behaviours is twofold; firstly the less the authors of the system are involved in the specific representations it uses the stronger the claim that can be made about the autonomy of its associations (Hofstadter and Mitchell, 1994), and secondly it allows for the investigation of the effect of different past experiences on associations that are made, a question of interest in design research.

The system’s ability to construct its own concepts allows the object representations to include grounded typological relationships. Relationships such as “these two shapes belong to the same conceptual category” and “these two shapes belong to similar conceptual categories” arise from the system’s classification of shapes. These can then be matched to other objects containing the

same relationships. Conceptual similarity is determined by proximity in the space of possible shape descriptions. Shapes are described by their outlines, with similar outlines defining similarity and identity between shapes. This similarity is invariant to scale, rotation and position. Four shapes that belong to each of three different concepts are shown in Figure 2. Whilst the majority of these shapes are triangular (an artefact of the generation algorithm used), none of the categories are exclusively triangular, as the number of vertices is not important, only the overall outline.

The set of shapes comprising an object is then transposed into a graph-based representation in which each shape is a node, and each relationship that exists between a pair of shapes is an edge between those nodes. These edges each carry a label identifying the relationship they encode.




Figure 2. Three different categories of shape detected by the perception system.

The exact set of relationship types that can be used is customisable, but relationships that have been used in the system include: proximity, scale, orientation, bearing, overlap, contained within, shared vertices and shared edges, in addition to the typological relationships of similarity and sameness mentioned earlier. The relationships are expressed relatively (e.g., ‘this object is half the size of that object’) to aid in generalising them to other shapes.

This graph of shapes and relationships between them is then used in mapping, one graph representing the source and the other the target.

2.2. MAPPING

The mapping process finds ways in which the shapes comprising the objects share a set of relationships. That is, given that each object is represented by a set of shapes, a mapping is when some or all of the shapes within two objects have the same relationships between them in the same pattern. It is not necessary that the shapes themselves be the same, only that the relationships

between them are. In this way the mapping process of our system searches for structural rather than purely visual similarities between two objects. When a mapping is found that includes a sufficient number of shapes in both objects in a shared set of relationships, the system has constructed a new association. What portion of the shapes in an object must be included for a mapping to be sufficient is an adjustable parameter.

The mapping process searches the two graphs for sub-graph isomorphisms in which shapes are mapped by matching edges in their respective graphs. Only some relationships on an edge need be the same. A mapping that produces two pairs of shapes that both share orientation but not bearing is a successful mapping. In the absence of the interpretation mechanisms, this search can only find associations that are present in the representations constructed by the perception system. Since the relationships are stored as relative values, mappings can be made between shapes that are apparently quite different but still share a pattern of relationships.

An example of an association that can be made without any interpretation affecting the graphs is shown in Figure 3. In this association all three shapes in both object A and object B are of the same concept and share the same progression of sizes – each shape is twice the size of the one below it. This association of relative size and relative sameness can be made even though in terms of absolute size and absolute shape they are different.

When interpretation is not used the mapping algorithm attempts to match the object graphs as they have been constructed by the perception system. It is during this search that interpretation is employed to direct the search algorithm by changing the search space. The interpretation system alters the graph representations, which in turn alters both the space the mapping algorithm is searching and the course of its search. New interpretations are produced as a result of search, influence the search process and change the mapping that that process eventually produces. Without an interactive parallel model of mapping and representation, the system would be limited to associating identical (if relative) relationships like the one in Fig. 3.




Figure 3. A simple association that can be made without an interpretation.

2.3. INTERPRETATION

Occurring concurrently with the search for a mapping between the graph representations of the objects is the model's interpretation process. Interpretation allows our model to find associations that are more abstract than when two sets of shapes share a pattern of relationships. Interpretation in association is a process that changes the representations being used. The interpretation process allows the model to find an association between two sets of different shapes that have different relationships between them, so long as the pattern of those relationships share a common structure. The system can relate, for example, a set of shapes that are all the same size to a set of shapes that are all rotated the same way, because the set of relationships between both sets of shapes shares a structure, even though they differ in content.

In this system an interpretation is defined as inducing an equivalency in meaning between one type of representation in the source and another in the target. This means that a particular edge-tag in the source graph is treated as a match with a different edge-tag in the target graph. These tags represent a pattern of relationships between shapes that exists both in the source and target objects, though neither the shapes matched nor their relationships need be literally similar.

The interpretation formation process takes potential mapping candidates from the search being performed by the mapping process and searches the node pairings they suggest for tags that do not match. If a coherent substitution of one tag in the source for another in the target would improve that randomly selected candidate mapping, the interpretation formation process suggests it as a new interpretation. Interpretations produced in this fashion are then evaluated against the current interpretation based on how many extra shapes they could match. If a new interpretation is superior to the current one by this metric it becomes the default way to view the objects and mapping proceeds under this new interpretation. This change of interpretations can only happen when the system has been unable to progress in its search for a period of time. This restriction ensures that the system is not too plastic in the adoption of different interpretations over a short period.

Through this interpretation process the system is able to make associations that are not based the existence of identical relationships in the source and target, but on shared structures of relationships that may literally be quite different. A simple form of interpretation-based association is shown in Figure 4. Here, in order to produce a mapping between the two objects, an interpretation has been constructed that treats changes in scale and orientation as equivalent. From the perspective "a change in size is the same as a change in orientation", A and B can then be associated.




Figure 4. An association based on the interpretation ‘+15% scale’ = ‘+15° orientation’.

In order to encourage diversity over multiple associations between the same two objects, the system will discard interpretations that it has used successfully many times in the past on the same association problem, as they will likely lead to solutions the system has seen before. Figure 4 represents the solution the system will find only after abandoning the default perspective and investigating other interpretations. With no interpretation the system finds that both objects in Figure 4 are groups of shapes all from the same concept. After a several runs finding the “same concepts” mapping the default null interpretation is discarded, other interpretations will replace it and the association presented in Figure 4 will be found.

This divergent behaviour expressed over many runs of the system allows multiple different associations to be found between two objects. Presenting the pair of objects in Figure 5 to the system, both the interpretations “conceptual similarity in A is the same as being from the same concept in B” and “being vertically adjacent in A is the same as sharing a vertex in B” are found. Which of these (and other) multiple possible interpretations is found by the system depends both on its history of making associations and on the direction that the mappings search process takes. The past experience possessed by the system at solving this and other, related association problems affects the interpretations and associations that the system produces.




Figure 5. An association based on the interpretation ‘adjacency’ = ‘shared vertex’.

3. Applying interpretation-driven search to a design domain

The association system presented in this paper can detect, describe, relate and categorise shapes within a design object, given only a simple vector-based representation of that object. The system can then use the higher-level representations it has constructed to find associations between objects, reinterpreting relationships within those shapes to find structures and patterns that exist in the representations of both objects. These associations, like the one in Fig. 5, may not be obvious to a human viewer. This ability is of interest to design researchers for both developing intelligent design support systems that exhibit associative reasoning themselves and in the modelling of how human designers use associative reasoning in their design process.

Systems designed to support designers through association could be of use at several points in the design process. The two design processes to which association is most applicable, as defined by the Function-Behaviour-Structure ontology of design (Gero, 1990), are *synthesis*, the creation of new design structure, and *reformulation*, the modification of design structure, behaviour or function based on existing design structure. Used to support the synthesis process an associative reasoning system would act as a form of brainstorming tool, providing a designer with new perspectives on the design problem and connections to other domains through which the design can be understood metaphorically. The association system could act as a flexible and interpretive visual index for design cases, similar to the way the Electronic Cocktail Napkin (Gross and Do, 1996) interfaced with a database of existing designs. Used to support the reformulation process an associative system could provide the ability to see new connections in the representations externalised by the designer and act as a form of artificially stimulated reflection-in-action (Schön, 1983).

Analogy making is an extension to the association process where the association is used to transfer knowledge from the domain of the source to the domain of the target. An analogy-making system based on this model of association could suggest ways to incorporate knowledge from the target domain into the designer's work based on the association that was used to retrieve a design case.

Cha and Gero (1999) describe style in architectural design as a set of relationships by which a hierarchy of visual elements are composed. Sets of shapes with consistent relationships between them form low-level patterns, and relationships between these patterns form higher-level visual structures. Examples of this relationship-based representation of style can be seen in Figure 6. The designs of Antonio Gaudi, (Figure 6a) show repetitions of similar shapes formed together into patterns (here part of an ornate grating

in the Palacio Guell and the roof of the Casa Mila). In Figure 6b the patterns are made explicit through Cha and Gero's analysis of the relationships that constitute them. An interpretation-driven visual association system would be able to associate these (and other) Gaudi designs based on their repeating elements and could be used to investigate visual style in the domain of design ornamentation.




Figure 6. Buildings of Antonio Gaudi, before (a) and after (b) pattern analysis (after Cha and Gero (1999)).

4. Discussion

The representations used during design are highly mutable and the space of possible representations cannot be defined in advance. This means that in order for a computational model to operate in a design context it must be robust to this representational dynamism. Existing models of analogy and association either require unitary, fixed representations for their objects or can only operate within a precisely defined universe of discourse. These limitations hamper the applicability of these models to creative design. The model of association described in this paper builds and refines the conceptual categories used in its representations in response to its experiences, and uses reinterpretation to match between elements that were not defined *a priori* as able to be matched.

A number of extensions to the system described here could increase its capabilities. In this paper we describe a system that operates on specified pairs

of objects, it does not currently search for a source object to associate with a given target. This simplification allows us to build multiple interpretations of a single source object whilst searching for a match, rather than having to consider multiple interpretations of multiple objects. The capacity to associate a given object with any object from a library would broaden the design tasks that the system could be applied to.

The system's model of interpretation as induced equivalencies between relationships is one of many mechanisms by which reinterpretation could change representations and influence search. Other means of re-interpretation include changing the representations of different shape elements, excluding or focussing on different shapes or relationships and applying visual transformations to the source and target and re-perceiving them. Additional associations and interpretations could be modelled by constructing patterns out of the shapes the system currently uses as concepts and modelling these higher-level patterns as 'supernodes' in the object graphs.

The system we have developed demonstrates that a model of association where interpretation and mapping are parallel interactive processes is feasible. We have demonstrated that such a system can construct its own interpretations, be affected by its experiences, and be applied to visual design domains.

References

- Cha, M. and Gero, J. S.: 1999, Style Learning: Inductive generalisation of architectural shape patterns, *Architectural Computing from Turing to 2000*, eCAADe, University of Liverpool, Liverpool, 639-644.
- Davies, J. and Goel, A. K.: 2001, Visual analogy in problem solving, *Proceedings of the International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, 377382.
- French, R.: 2002, The computational modelling of analogy-making, *Trends in Computer Science*, **6**, 200205.
- Gentner, D.: 1983, Structure-mapping: A theoretical framework for analogy, *Cognitive Science*, **7**, 155170.
- Gero, J. S.: 1990, Design prototypes: a knowledge representation schema for design, *AI Magazine*, **11**(4), 2636.
- Goel, A. K.: 1997, Design, analogy and creativity, *IEEE Expert* **12**(3), 6370.
- Hofstadter, D. and Mitchell, M.: 1994, The Copycat Project: a model of mental fluidity and analogy-making, in K. Holyoak and J. Barnden (eds.), *Advances in Connectionist and Neural Computation Theory*, **2**, 31112.
- Gross, M. and Do, E.: 1996, Demonstrating the Electronic Cocktail Napkin, a paper-like interface for early design, *Companion to CHI'96*, ACM New York, 56.
- Kokinov, B.: 1998, Analogy is like cognition: dynamic, emergent and context sensitive, *Advances in Analogy Research*, NBU Press, Sofia.
- Schön, D.: 1983, *The Reflective Practitioner; How Professionals Think In Action*, Basic Books.