

Interpretation-driven Visual Association

Kazjon Grace & Rob Saunders

Faculty of Architecture, Design and Planning
University of Sydney
kazjon@arch.usyd.edu.au, rob.saunders@sydney.edu.au

John Gero

Krasnow Institute for Advanced Study
George Mason University
john@johngero.com

Abstract

In this paper we outline ongoing research into a computational model of association based on the reinterpretation of a source object to fit the target. We describe the structure of the model and the concepts from which it arises. Preliminary results of visual associations made by the system in a simple shape domain are presented. We also discuss a planned application of our model to the analysis of a real-world creative design.

Introduction

Association is the construction of a mapping between source and target objects. This fundamental cognitive ability underlies analogical reasoning, metaphorical imagery and other creative processes based on constructing abstract similarities. This paper presents a computational model that focuses on construction and reinterpretation of representations during association; a particularly important process for computationally modelling creative analogy-making (French 2002; Kokinov 1998). Results of applying an implementation of the computational model to simple visual association problems is given and the application of the system to more complex problems is discussed.

Association is composed of three subprocesses: representation of the source and target objects; matching between the representations; and construction of a mapping around that match. These processes cannot be modelled serially, representation must occur in parallel with matching and mapping (Kokinov 1998). This contrasts with association as typically modelled in computational analogy-making (French 2002), where the concepts and/or the relationships between them are fixed. We have developed a model of association that focuses on the iterative interaction between the search for mappings and the construction of representations, an interaction that we call *interpretation-driven* search.

The system's ability to reinterpret objects extends mapping capability beyond matching identical features present in the provided representations. Our system's interpretation process guides and is guided by the ongoing mapping process. New interpretations are discovered through the search for mappings. Interpretation provides a capability akin to Copycat's 'conceptual slippage', except that there is no pre-defined list of conceptual equivalencies.

The following sections describe the computational model with reference to an implementation for simple visual problems. We also explore the application of the system to more complex visual problems in a design domain.

Interpretation-driven Association

The model described in this paper can be decomposed into three interacting systems, 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.

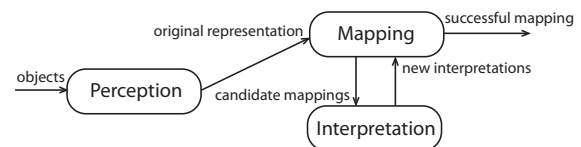


Figure 1: The structure of the interpretation-driven association system, showing how the representations produced by the Perception system are iteratively searched for mappings and changed through interpretation.

Perception

In the implementation presented here, objects are vector images composed of polygonal shapes. The perception system detects shapes and describes them using the contour of their outlines. The system's representations are constructed from these detected shapes and from relationships built between them, both typological and topological. Shapes are categorised into concepts, which are groups of shapes with similar outlines. A shape that has a representation unlike previously learned ones will generate a new conceptual category, and future shapes judged sufficiently similar will be added to that category.

We model these constructive behaviours in perception as the less the authors of a system are involved in its specific representations the stronger the claim that can be made about the autonomy of its associations (Hofstadter and Mitchell 1994). This autonomy is a necessary precursor to any claim that the system itself is capable of acting creatively.

The set of shape features for each object is translated into a graph-based representation where nodes represent shapes and edges represent relationships between those shapes. Topological relationships are based on the similarity of the conceptual categories the shapes are placed into. Topological relationships are based on geometric relationships within the object, including: proximity, scale, orientation, bearing, overlap, contained within, shared vertices and shared edges. To support the matching and mapping processes, relationships between shapes are expressed relatively, e.g., $size(A) = 0.5 * size(B)$. The result is a graph of shapes and relationships between them for each object, graphs which is then searched for mappings.

Mapping

The mapping process searches the source and target graphs for sub-graph mappings with an overlapping set of relationships between the two graphs. For example, a mapping between two pairs of shapes where both pairs share orientation would be successful, even if those shapes were connected by other relationships that did not match. As the relationships are stored as relative values, mappings can be made between quite different groups of shapes without applying any kind of interpretation to the representations.

Interpretation changes the object representations, broadening the possible mappings beyond finding identical relationships. The interpretation system alters the representations of source and target, which in turn alters the search space for the mapping process.

Interpretation

Interpretation in association is changing representations by taking a different perspective on one or both of the objects being associated. Interpretation is defined for the purposes of this system as inducing an equivalency in meaning between one type of representation in the source and another in the target. An interpretation states that a relationship in the source graph is to be treated as a match with a different relationship in the target graph.

The interpretation process takes an unsuccessful mapping under the current interpretation and, if a coherent substitution of one relationship in the source for another in the target would produce a better mapping, suggests it as an interpretation. Interpretations produced in this fashion are then evaluated against the current interpretation based on how many nodes they could add to a mapping if they were adopted. If a new interpretation compares favourably, it becomes the default way to view the objects and directs the mapping process accordingly. This process allows the system to make associations that are not based on identical patterns of relationships in the source and target, but on identical structures of relationships that may semantically be very different.

Figure 2 shows an association made by our prototype system. Fig. 2 A and B show the visual representations of source and target objects, while C and D show the graph representations constructed by our perception system. Both objects contain five shapes, with the shapes in the target all falling into the same concept (they have identical outlines), while the shapes in the source are similar but a different

conceptual category is created for each. Many relationships connect these shapes, but we highlight several pertinent relationships in the thick dashed lines in C and D. The lines connecting the two graphs show a mapping that was found by the system using the interpretation ‘being proximal in the source domain is the same as sharing a vertex in the target domain’. This interpretation was constructed and applied by the system during search. The system is designed to find many different associations for any problem, this is just one possible mapping with one possible interpretation.

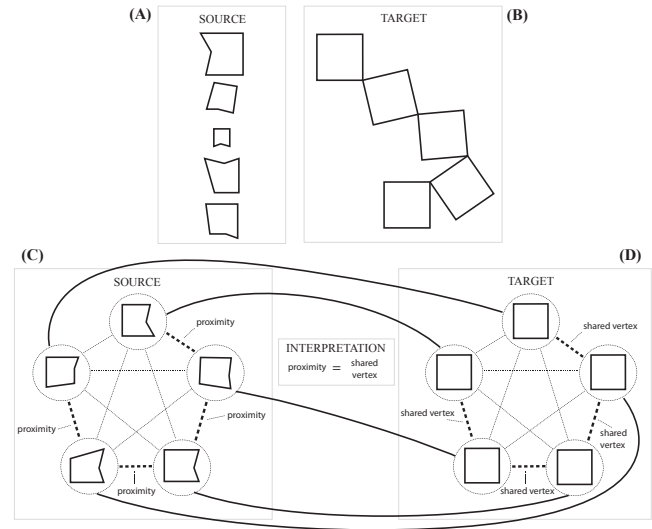


Figure 2: An example association problem: A and B are the visual representations given to the system, C and D are the graph representations constructed.

Applying Interpretation-driven Association

Preliminary association results, like those presented in Figure 2, demonstrate that the model has the capacity for finding non-obvious associations between groups of shapes. The works of a particular creator or of a particular school often share stylistic elements; reoccurring features within or between designs that are all variations of the style or theme of the work. In this ongoing research project we aim to determine whether our system can connect similar design elements with mappings that demonstrate their common style.

An example of a real world design that contains such a recurring stylistic element is Frank Lloyd Wright’s 1921 ‘Hollyhock House’. This California residence has a stone roof lined by distinctive stone friezes of a pattern of rectilinear shapes, seen in Figure 3a. The design makes strong use of geometric shapes throughout, but several other details make direct reference to the iconic frieze-work feature. Fig. 3b shows one of the custom dining chairs designed for the House, with a wood-carved back that calls to mind the design in Fig. 3a. There is also a stained-glass window design (Fig. 3c) that is clearly inspired by the frieze, with similar proportions and isometric projections of cubes representing the square blocks of the original design.

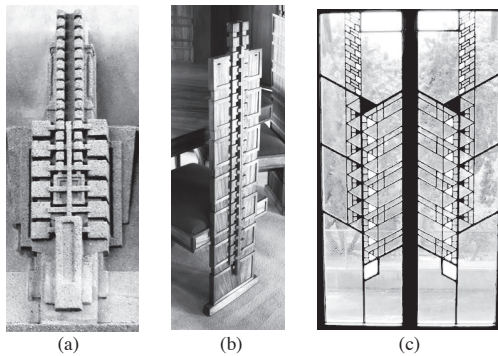


Figure 3: a) The iconic Hollyhock House frieze, b) a variation of the frieze pattern on a dining-room chair back, c) a stylistically related pattern on a window.

Given a simplified vector representation of each of these design, we will test whether the interpretation-driven association system presented here is able to find associations between these variations of a visual theme, e.g., finding associations between the chair-back design and the stonework that inspired it. We expect our system to construct interpretations that equate the of the frieze with those of the stained glass window. In doing so it will have generated an association that encodes part of the designer's visual style.

The system as it exists currently, with its closed shape based perception system, is well-suited to working with ornamental features and other details found in a variety of domains including architecture, industrial design, textiles, iconography and graphic design. Initial tests with alternative perceptual systems, e.g., based on SURF and SIFT descriptions of shapes (Bay et al. 2008; Rowe 1999), have shown that it is also possible to work with photographs, although more work is required to identify the most salient features detected for the construction of graph representations.

Discussion

This paper presents a model of association based on the principle of re-interpreting two objects so that there is a new relationship between them. We demonstrate a proof-of-concept implementation capable of making non-obvious associations between groups of shapes. We intend to apply this prototype to vector-image representations of real-world creative design artefacts to see if it is capable of finding common stylistic elements, both within a design as with the Hollyhock example and between different designs.

Cha and Gero (1999) describe style in design using a formal grammar 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 patterns form higher-level visual structures. The system presented here will be extended to support the construction of similarly sophisticated associations between patterns by allowing higher-level concepts to be formed from groups of existing shape elements. These meta-concepts will be treated as a 'super-node' in the object

graphs, composed of a number of other concepts but able to be related to singularly. This would remove the requirement for ordinality in associations (ie. five features in the target must always map to five features in the source).

The ability to treat groups of concepts with particular relationships between them as a single entity relaxes the restrictions on possible mappings and opens up new kinds of associations. This meta-concept formation could be implemented using of algorithms for finding cliques in graphs (Moon and Moser 1965) and by learning from previously known groups. By adding the ability to construct hierarchies of thematic or stylistic features, interpretation-driven association will be able to construct mappings to relate complex creative works. Our aim is to determine whether our system can build associations that demonstrate commonalities of style and structure between creative works.

The system described here implements one simple form of interpretation; induced equivalencies between relationships. Many other forms of re-interpretation are possible in our model, such as changing the definitions of shape elements, excluding or focussing on different elements and relationships within the representations or applying a variety of transformations to the objects or their representations. Our model is extensible to multiple forms of interpretation and the kind presented here is just one example.

We have demonstrated the feasibility of our interpretation-based model of association. Our system re-represents objects in parallel with the search for mappings between those objects. Our system constructs its own representations using conceptual categories that have been developed through its experiences. This system can produce associations based on interpretations of objects in simple visual domains. Research into applying this model to more complex domains is ongoing.

References

- Bay, H.; Ess, A.; Tuytelaars, T.; and Gool, L. V. 2008. Surf: Speeded up robust features. *Computer Vision and Image Understanding (CVIU)* 3(110):346–359.
- Cha, M., and Gero, J. S. 1999. Style learning: Inductive generalisation of architectural shape patterns. In *Architectural Computing from Turing to 2000*, 639–644. University of Liverpool, Liverpool: eCAADe.
- French, R. 2002. The computational modelling of analogy-making. *Trends in Computer Science* 6:200–205.
- Hofstadter, D., and Mitchell, M. 1994. The copycat project: a model of mental fluidity and analogy-making. In Holyoak, K., and Barnden, J., eds., *Advances in Connectionist and Neural Computation Theory*, volume 2. 31–112.
- Kokinov, B. 1998. Analogy is like cognition: dynamic, emergent and context sensitive. In *Advances in Analogy Research*. Sofia: NBU Press.
- Moon, J. W., and Moser, L. 1965. On cliques in graphs. *Israel Journal of Mathematics* 3:23–28.
- Rowe, D. G. 1999. Object recognition from local scale-invariant features. In *Proceedings of the International Conference on Computer Vision*, volume 2, 1150–1157.