Structure, Similarity and Spaces

Sandro Rama FIORINI [a,1], Mara ABEL [a] and Peter GÄRDENFORS [b]
[a] *Institute of Informatics, UFRGS, Brazil*
[b] *Lund University Cognitive Science, Lund University, Sweden*

Abstract. Much of the discussion about shape representation during the last two decades was fundamentally related to questions about the representation of parts. Inspired by the cognitive processes governing how people represent and think about parts, we give a brief summary of our framework for representing part structures. In particular, we are interested in the role of similarity and prototype effects in this context.

Keywords. Part-whole relations, similarity, conceptual spaces

Introduction

Humans seem to be prone to divide the complex shape of objects into parts. In seeing a cat, we divide its overall shape in some more-or-less well defined parts, such as the head, trunk, tail and legs. We can then use this information to recognize and think about cats. It seems that structure is intrinsically related to our everyday notion of shape. That leads us to a broader question: what are the cognitive phenomena that allow us to represent and think about parts of object as a whole, and not just parts of their shapes? In this paper, we introduce a novel way to represent the relation between parts and wholes that takes into consideration some of these phenomena.

In an influential work, Farah [4] carried out a meta-study about different kinds of agnosia in humans and proposed that two processes participate in object recognition. Object recognition can be *structural*, where recognition is achieved by the identification of parts of the object and its internal structure. On the other side, recognition can be *holistic*, based on the global characteristics of the object, such as overall shape. According to Farah, both processes work together in recognitions of broad categories of entities (e.g. faces, objects and written words). It is important to note that, as pointed by Peissig and Tarr [7], the structural *versus* holism problem is correlated, but independent of the recurrent question whether object (shape) recognition is model- or view-based (i.e. [2,3]).

If we assume that holistic and structural processes are necessary for object recognition, therefore it is reasonable to expect that both require an underlying conceptual structure conveying holistic and structural information. When it

[1] Corresponding Author: Sandro Fiorini, Institute of Informatics, UFRGS, P.O. Box 15064, CEP 91501-970, Porto Alegre-RS, Brazil; E-mail: srfiorini@inf.ufrgs.br. We thank CAPES and Petrobras project PB PRH-17 for the support in this work and Joel Carbonera for the comments.

comes to conceptual structures for object recognition, similarity effects seem to play a fundamental role (c.f. [3,1]). The general notion is that recognition could be reduced to judgements of similarity between perceptual input and internal representations. The question we are trying to answer is how a holistic/structural representation framework that supports similarity judgements should look like. We base our efforts on the Theory of Conceptual Spaces [5], a representation framework that embeds the notions of concept similarity and prototypes. Our general approach is to extend conceptual spaces so that it becomes more suitable for the representing holistic and, specially, structural information about concepts and objects. It should provide the grounds for novel computational approaches to concept representation based on holistic and structural similarity.

1. Conceptual Spaces

Gärdenfors' Theory of Conceptual Spaces [5] puts forward a new way for representing concepts using geometrical and topological structures, which complements symbolic and connexionist approaches. Given the available space, we present just a brief introduction, but it is enough to say that the Theory is based on the notions of concept similarity and prototypes. Put it simple, a conceptual space is a space in the mathematical sense, where objects are points and concepts are regions or sets of regions. If this space has a well-defined metric, then it is possible to tell the similarity between objects (and concepts) by measuring the distance between them: further objects are apart, less similar they are. The dimensions of a conceptual space have a special meaning: they denote the features — or qualities — through which entities can be compared and are frequently grounded in perception. Good examples are hue, mass, height, etc. Certain quality dimensions always co-occur, forming subspaces called quality domains. Examples of quality domains are colour, shape, taste, etc. Complex concepts are defined as set of regions in many quality domains. For instance, the concept of apple can be defined as a combination of the regions green and red in the colour domain, plus the cycloid regions in the shape domain, plus the sweet and acid regions in the taste domain and so on. An individual apple is represented by a single point (or vector) in the multi-dimensional space formed by all quality domains of apple and that is close (similar) enough to the regions that form the concept of apple. A type of apple is formed by subregions of apple.

2. Structure Spaces

We are interested in using conceptual spaces to represent holistic and structural information about concepts and objects. The holistic portion of an entity can be seen as its usual features: colour, shape, weight, etc. These can be readily represented as regions in quality domains. However, representing structural information is far from trivial. In doing so, we are fundamentally interested in describing the partonomic relations between parts and wholes. The Theory of Conceptual Spaces does not provide a complete solution for representing relations in general;

it simply suggests that they could be represented by a (Cartesian) product between the relata. We take this basic notion and develop it further to account for structural information.

Structural information can be represented in conceptual spaces through what we call *structure spaces*. A concept are represented as regions in a conceptual space; a structure space is a subspace of it, where part structures are described. A vector in a structure space denotes a particular configuration of parts of an individual. That is, a single vector encode the information about what parts compose a whole and also about how these parts are related to whole. Similar configurations of parts are close together in this space. For a single vector to convey all this information, much of it is naturally allocated to the dimensions. Given a concept C and a set of concepts $P_1, \ldots, P_n$ denoting parts of C, then we can generally define the structure space containing C as the product space of the quality domains of $P_1, \ldots, P_n$ and n quality domains denoting specific structure information about each part P_i. We call these domains *structure domains*; they represent information such as the displacement of the part in the whole, part quantification and so on. The actual structure space of C is the product space of the regions that define $P_1, \ldots, P_n$, plus regions in the structure domains. For instance, the structure space of Apple could be formed by the product of Core, Flesh, Seed and Stein, plus regions denoting the general positioning each part in an object-centred coordinate system. A vector in the structure space of C denotes a particular apple-structure: a combination of individual parts, each with a specific value for colour, shape, taste and so on. Close points in this space represent similar apple-structures. The combination of regions of each part in the product restricts what are the valid individual components of an apple. More importantly, the structure space can be further divided into specific regions defining types of apple-structures; e.g., the concept of an apple with acid flesh and short stein.

This basic formulation of our framework might raise questions, such as problem of co-determination between holistic and part qualities; the role of parts in taxonomies (c.f. Tversky [8]), the representation of the many kinds of partonomic relations; the question about dependent and essential parts; computational feasibility and so on. To all of those we have at least partial answers, but due to the limited space, we shall touch upon the issue we consider most critical: the problem of transitivity.

Wholes have parts, which can also have parts and so on. This might become an issue when we define a whole as a product of parts: given parts also can have parts, complex wholes could soon become multidimensional monsters; the structure space of the concept Universe would be impossible to describe fully. In more formal terms, we could say that the structure space of a whole might become the transitive closure of its parts. In order to solve this problem, we first assume that the part relation is not essentially transitive. This position contrast with formal theories of parts, such as Classical Mereology, but it has becoming increasingly common in recent years (c.f. [6]). Instead, we take a more cognitive stance; experience and perception are the sole determinants of which parts *directly* compose a whole. For instance, what determines that the person's heart is not part of the company where she works in is the fact that there is no use for such conceptualization in the actual context. However, if there is a change in context (e.g. parts of employee body becoming property of the company), we can easily adapt our conceptualization. We do not assume any hard *a priori* ontological distinction on parts and wholes, for it might hinder the plasticity of the representation. Instead, we provide a way in such plasticity could operate in our framework, by improving the definition of structure spaces. We introduce the notion of *dimensional filter*. A dimensional filter is a conceptual operator that projects a subset of the quality domains of a concept onto a smaller subset; it "selects" relevant domains of another concept. We can then redefine a concept as a product space of filtered parts, were just the relevant domains are selected to compose the structure space of the whole. The filtering (i.e. projection) is controlled by processes like attention and context. For instance, a combustion engine can be part of a car or part of an electricity generator. The quality domains of the engine that are relevant for car are related to its characteristics as a car mover. So, the projection of combustion engine into the structure spaces of car carries only some of its more relevant domains. This scheme solves the dimensional explosion by providing a way in such the transitive closure can be avoided; parts of parts that are not relevant for the whole can be filtered out.

3. Final Remarks

We are now developing a mathematical formulation of structure spaces based on metric spaces. This should pave the way for computational applications. We are also investigating the use of structure spaces in robotics and geology. Some auto-localization algorithms for robots employ similarity reasoning to compare its surroundings with its internal map. This comparison could benefit from a representation that allows structural similarity matching. In petroleum geology, an analogous problem of structural similarity exists in the task of matching geological structures in different exploration wells, to which no satisfactory computational solution yet exists. In the same way, structure spaces could help solve this problem by providing a principled way in which sequences of features can be compared.

References

[1] R. G. Alexander and G. J. Zelinsky. Effects of part-based similarity on visual search: the frankenbear experiment. *Vision research*, 54:20–30, Feb. 2012. PMID: 22227607.

[2] I. Biederman. Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94:115–117, 1987.

[3] S. Edelman. Representation is representation of similarities. *Behavioral and Brain Sciences*, 21(04):449–467, 1998.

[4] M. J. Farah. Is an object an object an object? Cognitive and neuropsychological investigations of domain specificity in visual object recognition. *Current Directions in Psychological Science*, 1(5):164–169, 1992.

[5] P. Gärdenfors. *Conceptual Spaces: The Geometry of Thought*. The MIT Press, Cambridge, Massachussetts, 2000.

[6] I. Johansson. Formal mereology and ordinary language - reply to varzi. *Appl. Ontol.*, 1(2):157–161, Apr. 2006.

[7] J. J. Peissig and M. J. Tarr. Visual object recognition: Do we know more now than we did 20 years ago? *Annual Review of Psychology*, 58:75–96, Jan. 2007.

[8] B. Tversky. Parts, partonomies, and taxonomies. *Dev. Psychol.*, 25(6):983–95, 1989.