

Representing part–whole relations in conceptual spaces

Sandro Rama Fiorini · Peter Gärdenfors ·
Mara Abel

Received: 22 November 2012 / Accepted: 7 October 2013 / Published online: 22 October 2013
© Marta Olivetti Belardinelli and Springer-Verlag Berlin Heidelberg 2013

Abstract In this paper, we propose a cognitive semantic approach to represent part–whole relations. We base our proposal on the theory of conceptual spaces, focusing on prototypical structures in part–whole relations. Prototypical structures are not accounted for in traditional mereological formalisms. In our account, parts and wholes are represented in distinct conceptual spaces; parts are joined to form wholes in a structure space. The structure space allows systematic similarity judgments between wholes, taking into consideration shared parts and their configurations. A point in the structure space denotes a particular part structure; regions in the space represent different general types of part structures. We argue that the structural space can represent prototype effects: structural types are formed around typical arrangements of parts. We also show how structure space captures the variations in part structure of a given concept across different domains. In addition, we discuss how some taxonomies of part–whole relations can be understood within our framework.

Keywords Part–whole relation · Conceptual spaces · Prototype · Context · Partonomy

Introduction

There are several types of cognitive phenomena involved in how we deal with the concepts. Such phenomena are unsurprisingly reflected in theories about how we believe concepts should be *represented*. However, the interplay between two important aspects related to the concepts, namely *conceptual similarity* and *part–whole relations*, is usually overlooked when it comes to concept representation.

The intuitive idea behind the role of similarity in categorization is that two objects belong to the same concept if they are “similar enough.” The exact definition of “similar enough” depends on the actual representation framework (Tversky 1977; Edelman 1998), but the general importance of similarity in categorization is well established (Goldstone 1994). In addition to similarity, concepts are thought to show prototype effects (Rosch 1978); concepts are defined in relation to one or more individuals that are judged to be typical exemplars of that concept (for example, a robin could be seen as a prototype for the concept of bird, in contrast to a penguin). The classification of a new object is determined by measuring its similarity to the concept prototypes.

On the other side, human cognition can also represent the relationships between entities and their parts, for example, between a horse and its four legs. These relations play an important role in how humans perceive and think about concepts. One of the central questions of this paper is how the notions of similarity and prototype are reflected in an analysis of parts and wholes. The idea of a prototypical whole seems to be intuitive enough to make it significant; for example, it is easy to think about a prototypical *pen*, with its typical configuration of parts. The degree of typicality of other pens can be measured by their similarity to the prototype. Partonomical (part–whole) similarity between wholes takes into account which parts are actually similar

S. Rama Fiorini (✉) · M. Abel
Institute of Informatics, Federal University of Rio Grande do Sul
(UFRGS), CP 15064, Porto Alegre, RS CEP 91501-970, Brazil
e-mail: srfiorini@inf.ufrgs.br

P. Gärdenfors
Lund University Cognitive Science, Lund University, Box 188,
Lund 22100, Sweden

and how parts are structured. While some experiments have indicated prototype effects in part–whole relations (e.g., Chaffin et al. 1988), the specific role of prototypes in these relations is not yet well understood. Nevertheless, there are direct and indirect evidences that partonomical similarity plays a role in object recognition and concept learning (Mash 2006; Wu et al. 2010; Alexander and Zelinsky 2012).

Our goal in this paper is to describe a *computationally oriented representation framework* for part–whole relations using knowledge about cognitive mechanisms involved in representation and reasoning with parts and wholes as sources of inspiration. In particular, we emphasize the role of similarity and prototypes. We will not propose a specific formalism for knowledge representation; the aim is, instead, to present a framework that can guide the development of such formalism. Furthermore, we restrict our discussion to part relations involving physical objects only.

Our proposal follows the tradition of cognitive semantics. It differs fundamentally from the realist semantics usually employed in knowledge representation, specifically, in ontologies (Guarino 1998). Whereas realist semantics defines meaning as mappings from language to one or more “worlds,” cognitive semantics defines meaning as mappings from language to *conceptual structures* within an agent’s mind (Gärdenfors 2000). Cognitive semantics provides a more principled account of the influence of cognitive mechanisms, such as concept learning, perception, and symbol grounding. Furthermore, purely symbolic languages usually employed for representing concepts, such as the Ontology Web Language (OWL), lead to difficulties in representing part–whole relations (cf. Rector et al. 2005). This is a good argument for looking for other representational formats. In particular, the advantages of using cognitive semantics to explain meaning have been investigated in relation to the knowledge systems (e.g., the Semantic Web) (Gärdenfors 2004; Adams and Raubal 2009). For instance, Adams and Raubal (2009) proposed the Conceptual Space Markup Language (CSML). CSML is a XML-based representation proposed as a complement to the Semantic Web languages.

As it has been put by Guarino et al. (1996), there are two approaches to the problem of representing of part–whole relations. One is the *logico-philosophical approach*, which takes the perspective of formal ontology and algebraic theories of parts, such as classical mereology (the formal study about parts and wholes) and other derived theories (Simons 2003; Varzi 2011). This approach seems to be dominant among ontologists, particularly in computer science. However, Simons seems to recognize that algebra is not enough:

When it comes to the honest toil of investigating the principles governing what objects are parts of others,

and what collections of objects compose others, it appears that most ontologists have been following the paradigm of abstract algebra when it would have been better to take a lead from sciences such as geology, botany, anatomy, physiology, engineering, which deal with the real. (Simons 2006)

On the other hand, there is the *cognitive–linguistic approach*, which is the one we adopt. Here, we consider the cognitive phenomena related to the concepts that are usually ignored in other approaches, like prototype effects and similarity. Furthermore, we account for the *context* effects of using concepts that frequently turn up. We submit that a semantic model suitable for implementing intelligent computational systems should be aligned with human cognition as much as possible.

To that effect, we base our analysis on the theory of conceptual spaces (Gärdenfors 2000), which takes similarity and prototypes into account. Its main feature is that the concepts can be represented as regions in a multidimensional space. Similarity plays a central role: the dimensions of conceptual spaces provide the means for determining similarity between concepts and between objects.

More precisely, we discuss the role of similarity in part–whole structures; in other words, what it means to say that two objects (or concepts) are *similar* because they have a similar part–whole structure. We present constructive proposals for modelling the conceptual structures of parts and wholes. In brief, we argue that the conceptual space of a whole can be seen as a product space, composed of the conceptual spaces of its parts accompanied by part structural information. The notion of using product spaces to form more complex conceptual construction is not new (cf. Aisbett and Gibbon 2001); however, we innovate by adding more information to product spaces based on cognitive phenomena.

Part–whole relations can also take many forms, having different meanings. For instance, the part relation *engine-car* is of a different nature than *tree-forest*. Many authors have proposed diverse categorizations of these forms (e.g., Chaffin et al. 1988; Gerstl and Pribbenow 1995; Simons 2003; Guizzardi 2005). We analyze how some of these forms manifest themselves in the conceptual world, allowing us to account for the plasticity of part–whole relations. In addition to that, we discuss how prototype structures in a partonomical hierarchy affect object classification; and how the same whole can be seen in different ways, taking into account contextual information on the parts. We finally present a simplified model of an object recognition algorithm based on our framework.

This paper is structured as follows. In “[Parts, wholes, and cognition](#)” section, we review the basic cognitive

background on part–whole relations. In “[Conceptual spaces](#)” section, we introduce the main ideas behind conceptual spaces. In “[Representing parts and wholes in conceptual spaces](#)” section, we describe the notions of holistic and structure spaces, the core of our contribution. “[Types of structure domains](#)” section details the types of part–whole relations and how they can be accounted for by our approach. In “[Marr’s hierarchical model, revisited](#)” section, we present a reinterpretation of Marr’s hierarchical model. “[Same whole, different parts,](#)” “[Paronomies,](#)” and “[Other ontological considerations](#)” sections are devoted to the discussion of related issues, such as context and paronomies. In “[Object recognition with structure spaces](#)” section, we sketch an object recognition framework using structure spaces as basis for concept representation.

Parts, wholes, and cognition

Before proposing a cognitive semantics approach to part–whole representation, we must consider how humans perceive and think about objects and their parts. However, the research on the cognition of parts and wholes is not as extensive as one would expect, given its importance for human reasoning. For instance, early work by Tversky and Hemenway (1984, 1989) showed that parts play a central role in differentiating between basic-level concepts, and also suggested that the parts form a bridge between perceptual and functional knowledge.

A good amount of the research in the cognition of part and wholes is centered on shape recognition. During the 1990s, the discussion concerning shape recognition gravitated around two general set of theories in which the importance given to part relations was a distinguishing feature. The *view-independent* theories, mainly influenced by Marr’s computational models of vision (1982) and Biederman’s work on *geons* (Biederman 1987), postulated that objects are represented and perceived based on the configurations of visual primitives that are invariant to viewpoint changes. On the other hand, there are *view-dependent* theories, like the ones proposed by Edelman (1998) and Ullman (2000), which state that objects are represented by “snapshots” (i.e., images) of the object’s different angles, dismissing the importance of structural information. View-independent theories have the tendency to give more relevance to the part relations (between visual primitives), while this aspect is not so much emphasized in view-dependent theories. Recently, though, evidence from cognitive and neurosciences (Foster and Gilson 2002; Newell et al. 2005) supports that both processes are needed in object recognition (see Graf 2006, for a review). Reviewing the state of the art in object recognition, Peissig and Tarr (2007) argue that the discussion regarding

dependence of view in object recognition is orthogonal to the actual importance of part structure in recognition of physical objects. In this paper, we focus on the latter aspect.

There are several relevant streams of empirical research in part–whole reasoning and representation. The first stream comes from the studies of patients with integrative agnosia, which is a rare kind of impairment that makes recognition of wholes difficult, but which leaves recognition of parts unaffected. In one experiment, Behrman et al. (2006) asked one of these patients to compare objects formed by different parts. The patient could recognize dissimilarities between objects that did not share the same parts. However, the patient was unable to recognize dissimilarities when objects shared parts that were arranged in different ways. Their conclusion is that the brain seems to encode part arrangement (spatial relations) independently of part shape (part qualities).

Developmental psychology also provides some insights into this topic. It has been shown that the recognition of objects by children under 2 years of age is mostly part-based. However, children later acquire the ability to recognize objects by their full shapes. For instance, a series of experiments with 18- to 30-month-old children, conducted by Smith and her colleagues (2009, 2011), suggests that the representations of geometric structures of whole objects are built over time, and, more broadly, that shape and part relations are two distinct components of children’s judgments of shape similarity. Rackison et al. (Wu et al. 2010) showed a similar trend in the development of children’s cognition, also suggesting that *salient parts* play a role in object categorization. The perception of parts also seems to affect generalization in learning. Son et al. (2008) found that teaching children the names of simple, featureless versions of new objects (e.g., some inner parts of the objects) helps them generalize the names to similar but more complex versions of the objects. This finding supports the idea that young children focus their attention on small details (parts) when learning words for new objects. Presenting them with simple objects steers their attention to more general geometrical structures, helping them learn and generalize words for basic-level concepts. However, parts are important when differentiating objects that are similar in overall shape (e.g., cows and horses). A child normally first notices high-level part similarities, but for some concept distinctions, more attention must be given to lower levels. For instance, dogs and cows have quite similar overall parts, and children sometimes do not distinguish them in their naming. Then, they learn to differentiate on lower levels of the hierarchy, such as by noticing that dogs and cows have differently structured noses and tails.

The existence of two stages in the development of object recognition suggests that there are two distinct but correlated representation systems in the brain: one based on the parts and the other based on the whole object. Evidence for this distinction also comes from a meta-analysis carried out by Farah (1992) on the research pertaining to the patients with different types of agnosia. She suggests that the brain employs two parallel but distinct cognitive processes in object recognition. In the structural process, whole objects are recognized by recognizing its constituent parts. In the holistic process, the recognition rests on the whole object, independent of its parts. The recognition of certain categories of objects usually relies more on the one or the other. In support of this position, there is evidence of a double dissociation between impairments in word recognition (regarded as structurally based) and impairments in face recognition (regarded as holistically based). At the same time, object recognition also seems to be partially affected in both impaired conditions, suggesting that object recognition is dependent on both structural and holistic representations.

Some studies showed evidence of interaction between part-whole processing and similarity effects. For instance, Alexander and Zelinsky (2012) showed that part similarity plays an important role in visual search of real-world object. Förster (2009) performed nine studies on how people judge similar/dissimilar stimuli in global and local processings. They found that global (holistic) processing tends to focus on similarities, while local (structural) processing tends to focus on dissimilarities.

In summary, the empirical evidence unveils two important aspects of part-whole representation by humans. First, information about the part and the whole are processed differently. Second, there is evidence to support the presence of similarity effects in holistic and structural processes. Consequently, one is to expect that a cognitive-inspired framework for representing part relations should account for such phenomena in some way.

Conceptual spaces

The theory of conceptual spaces (Gärdenfors 2000) is a framework for representing concepts using geometrical and topological structures, in the tradition of other geometrical concept representation proposals, such as (Shepard 1987). It has been employed in works ranging from computer science and robotics (Chella et al. 2001, e.g. Adams and Raubal 2009; Fiorini et al. 2013) to the philosophy of science (Gärdenfors and Zenker 2011). A rationale for proposing conceptual spaces is that the *concept similarity* is essential to understand the concept formation. The theory complements two other major approaches to the concept

representation: symbolic (logical) and associationist (connectionist), supplying an intermediate representation level.¹

A conceptual space is a multidimensional space where the concepts are projected and similarities represented. It can be understood as a space in the mathematical sense, such as a Euclidean space. The concepts correspond to *regions* in a conceptual space, whereas instances (objects) correspond to *points* (or, equivalently, vectors). If the space is provided with a metric, concepts and instances can be compared. Similarity between the concepts and instances can then be defined as the inverse of their distance in the space. In the following discussion, we assume that the conceptual spaces are metric.

The representational power of conceptual spaces depends on the selection of the dimensions of the space for an application area. The quality dimensions, as they are called, represent different ways in which instances and classes in the space can be compared. A canonical example is the *color space* that contains three dimensions: *hue*, *saturation*, and *brightness*. The perceivable colors can be represented as a combination of these three dimensions. In this space, each point represents a particular color. In everyday life, we do not refer to colors with such precision; we use labels instead, like “red” and “yellow”, to refer to two distinct sets of shades in which the members look sufficiently similar to each other to be referred to by the same label. Geometrically speaking, the concept of “yellow” corresponds to a particular *convex region* of the color space. However, different languages carve up the color space in different ways. Jäger (2008) has provided strong empirical support from more than 100 languages for the convexity of color concepts. Once the convexity of the concept regions is required, it becomes natural to define *prototypical* instances as points that are central to a region. For example, the focal red that can be experimentally identified will be at the center of the region representing red.

Conceptual spaces also introduce the notion of *quality domains*. A quality domain is a group of *integral* dimensions. Quality dimensions are integral when one cannot assign an object a value on one dimension without giving it a value on the other(s) (Garner 1974; Maddox 1992; Melara 1992). A color cannot be given a hue without also giving it a brightness; a sound’s pitch always goes with a certain loudness. Dimensions that are not integral are *separable*, for example, the size and hue dimensions. Using this distinction, we define a *domain* as a set of integral dimensions separable from all other

¹ In brief, the conceptual spaces fill the gap between the lack of flexibility of logic-based constructions in relation to the concept learning and the unstructured malleability of connectionist approaches to raw data. A more detailed discussion can be found in (Gärdenfors 1997; Gärdenfors 2000).

dimensions. The three color dimensions constitute a prime example of a domain in this sense: hue, saturation, and brightness are integral dimensions separable from all other quality dimensions.

Concepts defined exclusively within a single domain are called *properties*. For example, “yellow” and “red” are the properties, since they are single regions defined in a single domain, i.e., the color space. Other concepts can be defined as a set of regions involving many quality domains. The concept of “apple” is a good example: it comprises regions in domains like color (red, green), taste, shape (cycloid), texture, smell, and nutrition.

Usually, the conceptual spaces are constructed out of many dimensions and domains that can make their depiction very challenging. We have devised a simple diagram that emphasizes the multidimensional composition of the conceptual spaces as a product of quality domains. Figure 1 exemplifies this diagram for representing the concept “apple.” The *apple space* is represented as a product space of properties (smaller ellipsoids) in the quality domains that form the conceptual space (bigger ellipsoids). This diagram is inspired on the intuitive notion that a concept in conceptual spaces can be seen as a product of regions (or subspaces) in a series of quality domains (Fig. 1a), or as a region in a multidomain space generated by the product of quality domains (Fig. 1b). The ellipsoids and domains can be drawn in different colors and sizes to convey additional information.

Representing parts and wholes in the conceptual spaces

The cognitive grounding of the relation existing between parts and wholes must be founded on a broader theory of concepts. Our aim is to show that conceptual spaces can provide the basis for such a theory. In the next sections, we describe how part relations can be founded in the conceptual spaces and discuss the consequences for the concept representation. The general idea is that the relationship between a whole and its parts is represented in a *structure space*, where structural similarity between wholes can be measured and prototypical wholes can be identified. We start by exploring the relationship between the whole and its structure.

As mentioned in “[Parts, wholes, and cognition](#)” section, the cognition of part–whole reasoning seems to require that the descriptions of concepts take into account holistic and structural information. Additionally, similarity effects present in categorization seem to suggest that the same descriptions should also support similarity comparisons. Bringing all this together, we can redefine the concept similarity as a function of holistic and structural similarity. Intuitively,

$$\text{Concept similarity} = \text{Holistic similarity} \otimes \text{Structural similarity} \quad (1)$$

The goal is to define a representation structure that allows for such similarity calculations.

We assume that wholes and parts have their own representational units, that is, certain properties are exclusive to wholes and some properties are exclusive to parts. More specifically, we assume that the whole and each of its parts are represented in their own, distinct² conceptual spaces. For example, the concept of bird is placed in a conceptual space with its own dimensions, while the concept of beak, wings, and feet is placed in three other conceptual spaces, with independent dimensions and domains. There are of course correlations between the properties of the whole and the properties of the parts. Nonetheless, we do not assume then to be necessarily linked. For instance, the concept of black woodpecker (*Dryocopus martius*) certainly occupies the “black” region of the color domain, as would both of its wings; however, the color of its crown does not correlates to the color of the whole and is to be positioned in the “red” region.

The relationship between the conceptual spaces of wholes and parts is represented in the conceptual space of the whole. It is structured in such a way that it implements the conceptual similarity defined by Eq. (1). Accordingly, we propose a definition of the conceptual space of any whole as a *product space* of two *subspaces*: the *holistic (sub)space*, which represents the properties of the whole, allowing for holistic similarity comparisons; and a *structure (sub)space*, which represents the relation of the parts with the whole, and allow for structural similarity comparisons. Thus, we have that

$$\text{Conceptual space} = \text{Holistic space} \times \text{Structure space} \quad (2)$$

The inner form of the holistic space is, for the present purposes, reasonably unproblematic. Holistic spaces are standard conceptual spaces with dimensions and domains describing the properties about the whole. For instance, the conceptual space of apple represented in Fig. 1 can be seen as a holistic space, for it mainly describes the properties of the whole apple. In this space, whole apples can be compared regarding their similarity.

We are, however, more interested in the inner form of the structure space, which is naturally more intricate. It has to implement structural similarity, relating parts to the whole. Therefore, we have to first consider in more

² The real meaning of the word “distinct” here depends on the actual way chosen to implement the conceptual spaces. Some representations tend to aggregate all possible concepts in a unique, comprehensive conceptual space including all domains. Others may represent the concepts in really distinct conceptual spaces, formed by the combination of the relevant quality domains.

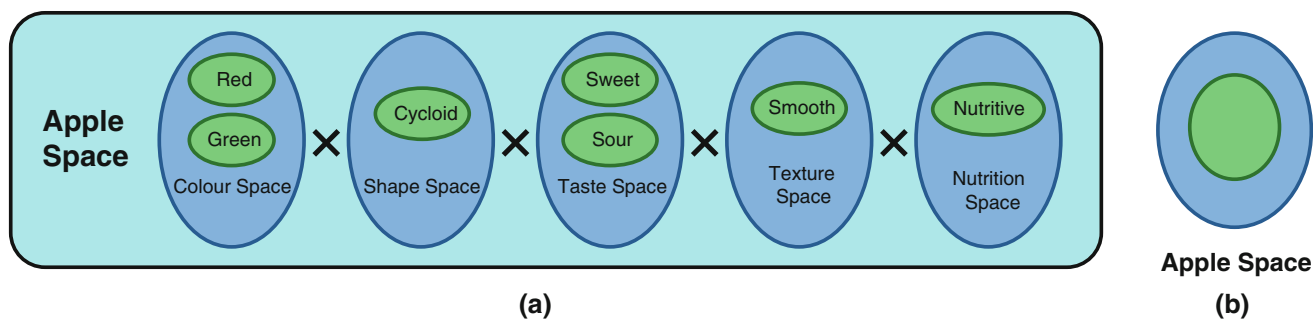


Fig. 1 Example of diagrams depicting the conceptual space of apple: **a** shows the inner form of the apple space as a product of properties (*smaller ellipsoids*) in different quality domains (*bigger ellipsoids*);

and **b** shows a compact representation of the apple space as a set of points (*smaller ellipsoid*) in a multidimensional space formed by the product of its quality domains

detail what structural similarity is, before unveiling its inner form. The intuitive notion behind structural similarity is that two wholes are structurally similar if they share a similar set of parts. This explanation is nevertheless incomplete. For instance, according to it, a pile of Lego bricks and the assembled Lego toy would be considered similar entities, since they share the same set of parts. Yet, the two entities are clearly structurally dissimilar: although the pile and the toy share the same types and number of parts, they do not share the same internal *structure*, which is essential for distinguishing between the two. Other authors have drawn attention to this separation between part and structure (Hummel and Biederman 1992). Thus, we claim that structural similarity consists of two elements: part and configuration similarity. Two individual wholes have high *part similarity* if the parts in one whole are similar to the parts in the other. It takes into consideration part categories and their number. Additionally, two individual wholes have high *configuration similarity* if parts in both wholes are also arranged in a similar way. The combination of these two kinds of similarity defines structural similarity between two wholes: similar parts placed in a similar configuration.

With the notion of structural similarity in mind, we can define in more detail what kind of construction a structure spaces have. A structural space is a conceptual space where one can represent and compare many possible part configurations. It is a high-dimensional space, where each point (or vector) corresponds to a particular part configuration and regions denote different kinds of part arrangements. It is formed by a product of a subspace of (the conceptual spaces) each constituent part, as well as a set of properties in special quality domains called *structure domains*. Structure domains modulate how parts bear to the whole, representing, for instance, displacement information. They are quality domains in the sense they also describe the qualities of a concept. Thus, following from Eq. 2, we have that

$$\begin{aligned} \text{Concept} &= \text{Holistic properties} \\ &\times (\text{Part properties} \times \text{Structure properties}) \end{aligned} \quad (3)$$

More formally, given a whole C and its parts P_1, P_2, \dots, P_n , a *structure space* is a subspace of C formed by the product of subspaces of P_1, P_2, \dots, P_n and supplemented by structure domains S_1, S_2, \dots, S_n to each P_i . The motivations for defining structure space essentially as a product of the parts is the necessity of finding a straightforward way of measuring structural similarity as a distance function between wholes. When we consider that a particular object can be represented by a single point in a space, then this single point must encode information about the specific set of parts composing the whole (encoding part similarity) and also its structural information (encoding configuration similarity). A good way of doing this is to transfer this information to the dimensions: the set of parts composing the whole is encoded by joining the quality dimensions of the parts, and the structure information is encoded by joining the structural dimensions for each part.

As an example, consider the conceptual space of apple in Fig. 2. Each part (stem, seed, skin, and flesh) is defined as a set or properties in their own conceptual spaces (Fig. 2a). In order to form the structure space of apple, subspaces of the parts are used to compose the product space that forms the structure space of apple (Fig. 2b). Notice that the use of subspaces of the parts is due to the fact that the whole relates just to a subset of the individuals described by the general part concept. For instance, the concept fruit seed describes all kinds of fruit seeds, given that only a fraction of them (a subspace) can be said to be apple seeds, i.e., oval dark seeds. Furthermore, for each part, there is a property in a structure domain that defines the configuration information of that part (depicted as lune shapes in Fig. 2b). In the context of this example, configuration information can be regarded as the displacement coordinates (e.g., position and orientation coordinates) of each part of the apple in relation

to the whole apple.³ For instance, Fig. 2d shows a two-dimensional structure domain of coordinates centered at the apple, with the region at the top corresponding to the coordinates allowed for the position of the stem. The same kind of information is added as regions in structure domains for each part in the structure space. The result is the complete apple structure space, which can be seen as a multidimensional conceptual space itself (Fig. 2c). A given point p in this space corresponds to a particular structure of apple, with a specific set of parts displaced at specific places. Points in the neighborhood of p correspond to similar apples, such as apples having a slightly different stem positioned at a slightly different place on the apple. It is important to note that structure domains allow the representation of structural information in terms of object-centered configurations. They can also be combined to describe more complex information, such as spatial relations (e.g., right of and back of).⁴

The link between structure spaces and parts is governed by what we call *dimensional filters*. A filter is a higher-order structure that defines which subspace of the part that is actually used to compose the structure space of the whole. In this context, the role of a filter is twofold. First, it can be used to filter out the sections of the structure space of a part that is not relevant to the whole, i.e., as in the *stem-apple* example above. Second, it can also be used to filter quality domains of the part (holistic or even structural) that are not relevant to the whole. For instance, a combustion engine may have quality domains describing its characteristics as a car engine and as a power generator. However, it is necessary to filter out the quality domains regarding power generation when the concept is imported in the structure space of a car. In some ways, the dimensional filter works like a context, screening out the unimportant dimensions of its parts.

The formation of the concepts in the structure space is, to a large extent, determined by prototype effects. Some part structures can be seen as more typical than others. These typical individuals—not necessarily any that exists in reality—determine the focal points of the convex regions that form the concepts in structure space. Take the concept of an apple as an example. Its part structure would be determined by a prototypical exemplar of its kind, denoted by a point in the structure space of apple. In turn, this prototype determines the focal points of the convex regions, which fully define the concept of apple structure. Furthermore, regarding the relationship between the

³ An object-centered coordinate system seems to be preferable. Nonetheless, one can think of structural information based on an egocentric coordinate space, or even a retinal coordinate space (Newell et al. 2005).

⁴ More about the conceptual spaces and spatial relations can be found in Ligozat and Condotta (2005).

prototypical whole and the prototypical parts, it is tempting to say that the prototype of a whole is also composed by the prototypes of its parts. However, even if this assessment could be true in some cases, it fails in a high number of situations. For instance, our prototypical notion of grasshopper includes a subspace of the concept of wing that certainly does not include what we consider to be the common prototype of wing.

At this point, it is important to highlight that this framework does not address *what* is a part or *how* we separate parts from the whole (e.g., in perception). We are just interested in determining a way of representing the relationship between the parts and wholes so that holistic and structural similarity can be measured. We return to this point in the next sections.

Types of structure domains

Structure domains and properties are fundamental parts of our framework. They qualify the partonomical relationships between the parts and wholes. As such, it is natural to expect they themselves might have a complex form. As a matter of fact, the elementary notion of being part of something can be specialized into more specific types of part–whole relations. In recent years, several authors have proposed different taxonomies of part–whole relations, based on many different criteria (Winston et al. 1987, e.g., Gerstl and Pribbenow 1995; Guizzardi 2005; Johansson 2012). In this section, we show how different ways of constructing structure domains can explain common kinds of part–whole relations.

Complex and collective

Essentially, there are two clear-cut kinds of part relations. A part relation can be *collective*, when a class of instances of a certain type generically composes the whole. For example, the part relationship between tree and forest just defines that an instance of forest includes an arbitrary number of trees; no particular tree is conceptualized or has a specific role or place in the forest. On the other hand, *complex* relations give more specific information about how the parts relate to the whole, such as with functional or displacement information. For example, the part relationship between cat and tail specifically defines that an instance of cat has an instance of tail positioned at a particular place on the cat. These two types of relations can be found in some part–whole taxonomy proposals, such as those of Winston et al. (1987) and Gerstl and Pribbenow (1995).

Regions in structure domains, called *structure properties*, have the exact function of defining how a part fits into a whole. Different structure domains allow for different part

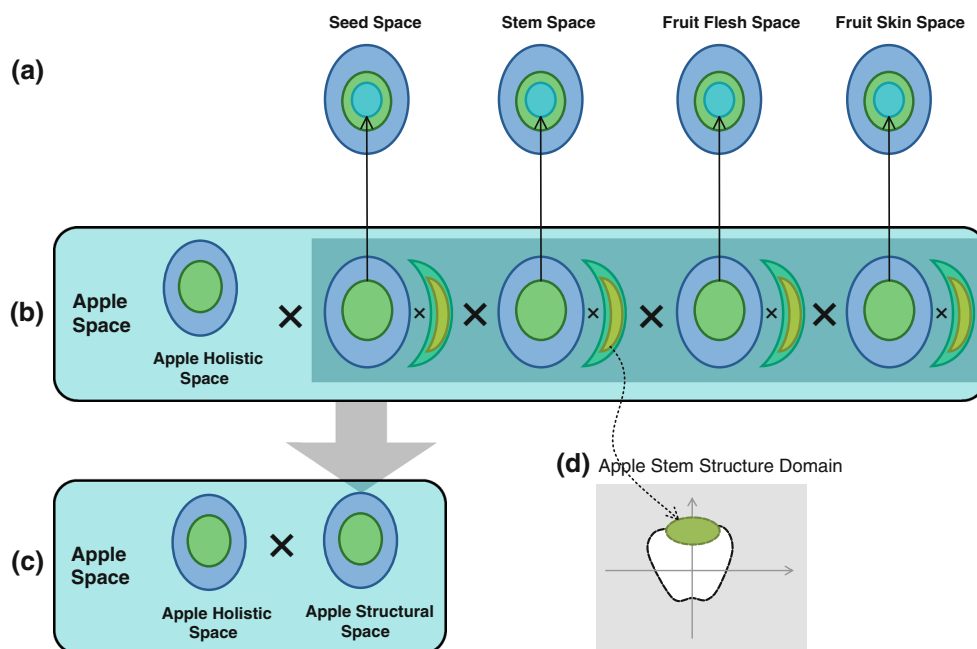


Fig. 2 Example of structure space for the apple concept: **a** the conceptual spaces of each part of apple, their inner form (*dimensions and domains*) is omitted; **b** the conceptual space of apple as a product of the holistic space and subregions of parts and also structure

information; **c** a compact representation of the holistic and structure space of apple; and **d** an graphical depiction of the regions defining the displacement for the stem in the structure space of apple

relations. In the previous section, we talked about the structure properties that represent displacement information only. These are *complex* structure properties. Collective relations can be represented in structure spaces by changing the type of structure domain used. Take the relation $R(A, B)$ as a part relationship between a part concept A and its whole B . Let S be the structure space of B . Let D_r be the structure domain added by the relation R to S . Then, based on the type of R , we can establish which kind of information is carried by D_r :

- R is a complex relation iff D_r is formed by complex structure domains;
- R is a collective relation iff D_r is formed by collective structure domains.

In complex relations, the structure domain denotes the specific configuration or role of the part, like the allowed positions and orientations. For example, the relation *part (tail, cat)* is complex because the relationship between the cat and its tail is modulated by a complex structure property, i.e., a region in a complex structure space denoting the allowed set of positions and orientation of the tail in a cat (Fig. 3a). In collective relations, the structure domain *quantifies* the part concept, such as how many instances of the part are the components of the collective whole. For instance, our general conceptualization of a forest is that it is simply composed by many trees. As such, the relation *part (tree, forest)* is a collective relation because the

structure domain that modulates it in the structure space of forest is a *collective* structure domain, that is, it does specify only a region (i.e., interval) in a *quantification space* corresponding to how many trees a forest could have and perhaps how they are packed (Fig. 3b). This scheme allows the representation of individuals as “*there are thousands/many/few trees in that forest.*”

Complex and collective structure domains can have a variety of forms, depending on the implementation. Regarding *complex structure domains*, we already described an example in the last section where an object-centered position space could work as an implementation. In such a space, the displacement of parts of an object is described as a two-dimensional coordinate of each part within the whole. However, one can devise more sophisticated implementations for complex structure. Consider displacement now defined as the geometric volume in the spatial extension of the whole where a given part can be found (or seen). For instance, if a person is asked to point where the engine is usually placed in a car, she will probably point at the front of the car, drawing an ellipsoid with her finger, while saying “around there.” This ellipsoid captures the intuition of the displacement volume we are talking about and it is naturally a function of the overall shape, position, and orientation of the part. This sort of construction can be neatly represented as a property region in a geometric volume domain, such as superquadrics (Chella et al. 2001). A point in this region represents the

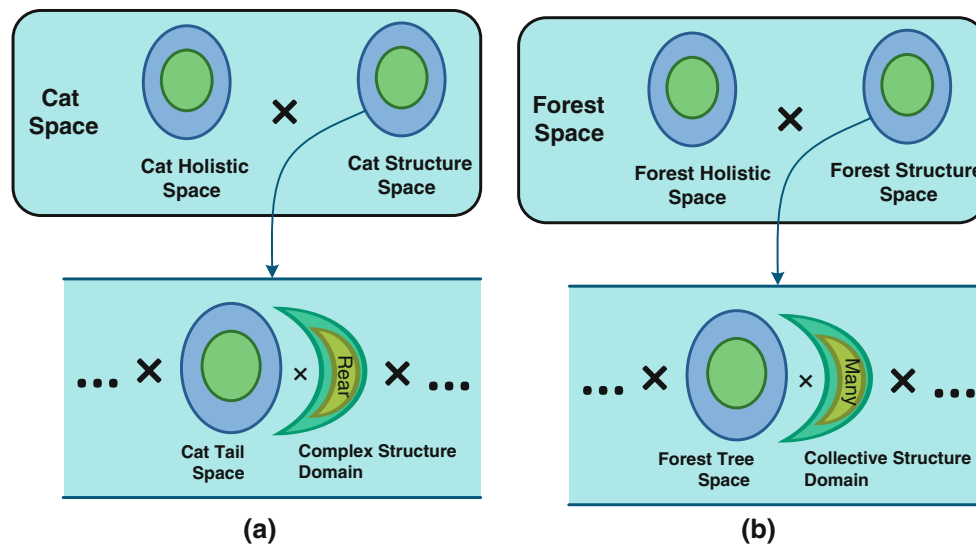


Fig. 3 Structure spaces according to partonomy types: **a** fragment of the structure space of cat showing the complex relation of tail; and **b** fragment of the structure space of tree showing the collective

relation of forest. Both “rear” and “many” are property regions on their respective structure domains

specific placed volume where the part can be found in the whole (e.g., “where the engine is placed in this car”). For example, the displacement of the skin of an apple can be represented as a region embarking the whole apple; a point in this region encodes a volume that coincides with the surface of the apple.

Collective structure domains on the other hand represent how a *set* of parts of the same type generically related to the whole. For instance, it might be represented as simple one-dimensional spaces denoting, for instance, how many instances of that given part are expected to be found in the whole. More interestingly, we believe that collective structure domains might be able to represent the notion of *ensembles* of parts (Alvarez 2011), which are related to the capacity humans have to summarize the groups of similar objects into a cognitive compact average representation. In this context, we could see the relationship between “tree” and “Amazon Forest” as conjunction of a conceptual space describing an average Amazonian tree and a region in a collective structure property denoting how many trees this forest has (e.g., say, “billions”, or “many”).

Perhaps not surprisingly, complex and collective relations are intrinsically correlated. We can see collective relations as a generalization of many complex relations to which no displacement information is necessary. This might account to what Gerstl and Pribbenow (1995) call the *plasticity* of part–whole relations. For instance, it is possible to classify the relation *ship–fleet* in two ways. If one considers a fleet as a uniform set of ships, the relation is collective. On the other hand, if each ship has a special role in a fleet, one can consider the relation complex. We argue that there is a third case where both kinds of relations

are mixed: certain ships could have special roles while others are referred to generically (e.g., in a fleet composed of a carrier and many destroyers); then, the relation can be seen as a hybrid of collective and complex. These processes of change can be explained by the processes of folding and unfolding of complex relations into collective relations.

External partitions

In general, we talk about parts as the building blocks of objects, in the sense that an apple is formed by many parts. However, we can also talk about arbitrary partitions imposed on the objects. For instance, if we say that “the upper part of the house is blue,” we are imposing a somewhat arbitrary partitioning of the house. In some cases, these *external partitions* can be seen as part of the definition of certain concepts, thereby composing their structure space. For instance, the concept “planet” might include the notion that the planet’s polar regions are cold; the notion of “polar region” can be seen as an external partition imposed on planet; there is no actual inner structural part that corresponds to the poles of a planet. Note that the parts in a given partition are disjoint by definition, but the same entity can be subject of many simultaneous external partitions.

We can describe external partitions in the conceptual spaces by means of certain region operations (e.g., intersection) on the quality domains of wholes and their parts. For instance, the concept of “the equatorial region of a planet” translates to an intersection of a property in the shape domain of “planet” (a sphere) with the property in

the shape domain of “equatorial region” (a section of a sphere). This new concept helps to compose the structure space of “planet.”

External partitions in conceptual spaces might explain a kind of parthood that recurs in the literature, namely what Gerstl and Pribbenow (1995) call *external parthood*. They argue that certain part relations derive from the internal structure of the whole, whereas others can be said to derive from external partitions imposed on the whole. In their original proposal, Gerstl and Pribbenow (1995) define two types of external part relations: portions and segments. A *segment* is a spatiotemporal part that results from the imposition of an external scheme on the whole. This scheme distinguishes different parts of the object, indifferent to its internal structure (“the upper part of the body” and “the beginning of a story”). On the other hand, a *portion* is construed by using a property dimension to select parts out of the whole; for example, the dimension of color is used in phrases like “the red parts of a painting” or “the annoying parts of the evening television show.” We can explain portions and segments in conceptual spaces by relating them to intersection operations in particular kinds of domains. In brief, portions can be seen as the results of restrictions in spatiotemporal quality domains, whereas segments are the result of restrictions on other quality domains. For instance, the *pole–planet* example translates a typical case of segmentation: a restriction on the spatial domain of “planet.” On the other hand, a sentence like “the white patches on the planet” translates the case of portioning. The sentence refers to a specific part of planet that is formed by the other parts in which the color property regions intersect with the color property “white.”

Marr’s hierarchical model, revisited

The basic utilization of structure spaces can be exemplified by reinterpreting Marr and Nishihara’s (1978) influential hierarchical model of cylinders. In the literature, there are many attempts to model the shapes of the objects in object recognition (Marr and Nishihara 1978; Pentland 1986; Biederman 1987; Zhu and Yuille 1996; Chella et al. 2001). Many of these attempts take into consideration the visual part–whole structure of the objects, associated with some sort of shape primitive, like cylinders or more complex parametric volumes. The model by Marr and Nishihara (1978) employs sets of cylinders to approximate biological forms, as illustrated in Fig. 4. The cylinders are combined in a hierarchical manner, with the torso on the first level, and the head and limbs (arms) on the second, forearms on the third, and so on. In the following, we demonstrate how Marr’s model can be described in our framework.




Fig. 4 Part hierarchy model proposed by Marr and Nishihara. The arm is deconstructed as finer parts along the chain of part structures (adapted from Marr 1982)

We can consider each level of the Marr’s hierarchy as a single concept (“body”, “arm”, “hand”, etc.). Each concept in this hierarchy has two types of descriptions. One represents the whole (e.g., the whole body) and the other connects the parts of the hierarchy together (e.g., the limbs and head). The description of the whole is called a *model axis*: a generalized cylinder that captures the overall outline and orientation axis of a particular hierarchy level (e.g., the cylinder that represents the whole body in Fig. 4). It abstracts away the details that are usually supplied by the parts. This approach is closely related to the model of holistic versus structural descriptions: the model axis represents a holistic take on the object shape at a given hierarchy level, while the same level includes more detailed structural (part) descriptions.

The recreation of Marr’s model using our framework is quite straightforward. We can model each level in Fig. 4 (body, arm, forearm, etc.) as a concept with properties in holistic and structural domains. The elementary quality domain here is *shape*. In this formulation, we can represent the shape space by two dimensions: cylinder length and radius. Different regions in this space denote different cylinder shapes. This shape domain is used to compose all concepts in this hierarchy. Each concept is then formed by a property (a region) in the shape domain, plus a structure space formed by its parts.

The structure space is formed by quality domains, imported from subparts of the body, and structure domains. In this case, the structure domains are of the *complex* kind and encode information about the position and orientation of the parts in the coordinate space of the whole. For instance, the concept of “hand” in Fig. 4 has a structure space formed by the shape domains and the properties of each finger, plus structure domains defining the position and orientation allowed for each finger. A point in the conceptual space of “hand” describes a particular whole

shape cylinder for the hand, plus particular shapes, orientations, and positions of the fingers. In this space, we can refer to prototypical hand shape and configurations, and categories of hand shape.

Same whole, different parts

Structural similarity judgments are influenced by *context*: two apples can look very similar to a child in a supermarket, but very different to a botanist. Distinct parts have distinct importance [or “goodness,” according to Tversky (1989)] depending the context. Similar to Tversky (1977), context is represented in the conceptual spaces as different weights given to each quality dimension (or domain) in the conceptual distance measurement (Gärdenfors 2000). We can use the same method to account for the influence of context in structural similarity judgements. Domains of different parts in structure space receive different weights: parts that are less relevant receive smaller weights and parts that are more relevant receive larger weights. In the apple example, a botanist comparing apples will give more weight to internal parts when trying to decide whether a given apple is included in the “good apple” category, whereas a consumer will give more importance to external parts (Fig. 5).

If we define *context* as a vector of weights, it implies the existence of a *context space*. A context space is a higher-order space where points denote different combinations of weights of the quality domains in a conceptual space. Again, we believe that context space also includes prototype structures: some context situations are more typical than others. Prototypes denoting typical part-importance scenarios will complement the context space for a structure space. For instance, the same person can play the role of a botanist and consumer at different times; the present context “moves” through different categories in the context space. In an apple dissection situation, one can pay attention to the internal parts of the apple; nevertheless, the prototypical apple consumer context is the one where just the more external parts of the fruit are relevant for comparisons. Situations that are close to the prototype situations define the concept regions in context space; these regions can be interpreted as kinds of context, like “apple dissection” context or “apple consuming” context.

Another example in which context plays a role is a situation where the characteristics of the different parts influence the categorization of the whole. For instance, a boat with a black hull and white sail may be “the black boat” in the context of boats with white hulls, and “the white boat” in the context of boats with black sails. Context influences part saliency, which in turn affects categorization.

Partonomies

A common way of describing part structures is to use partonomies, or a tree structure of parts. A partonomy is a simple representation that highlights the common transitive nature of part relations, allowing one to visualize and navigate the different levels of deconstruction. While useful as tools for reasoning about parts, there is no general way of representing partonomies as points in a space.⁵ Thus, it is not possible to represent part hierarchies explicitly in the conceptual spaces. On the other hand, the hierarchy is implicit in a sequence of structure spaces represented as a chain of whole–part pairs. It can be derived as a symbolic construct from the chain of structure spaces, but it is not necessary for comparing the similarity of objects.

Structure spaces import the domains of all parts. Given that parts can also be wholes, and that they have structure spaces of their own, the final structure space of a more complex whole can become a transitive closure of all its parts and subparts. This is not desirable from the standpoint of cognitive economy and, at a first glance, can be seen as a fatal limitation of our framework. The mechanism that allows us to deal with this issue relates to the transitivity problem. Any discussion of the representation of partonomies usually arrives at the problem of transitivity in part–whole relations (Simons 2003; Varzi 2006). The classical formulation of mereology holds that transitivity is one of the basic properties of the part relation. However, transitivity often breaks down. A good example of this breakdown is the following: the eye is usually regarded as part of the face and the retina as part of eye; however, the retina is hardly regarded to be part of face. There are some solutions to this problem (Varzi 2006); one of the accepted solutions is to consider ontologically distinct types of part–whole relations, which are naturally not transitive between themselves. For example, the relationship between eye and face is not of the same kind as the relationship between retina and eye. This general solution, however, does not help us to solve our representation problem, mainly because we do not assume any a priori meta-concepts in the conceptual level.

In order to give a partial solution to this problem, we deviate from the classical accounts of mereology and *take the part relation not to be transitive*. We see few reasons to consider a cognitive interpretation of part relation as necessarily transitive. First, there are plenty of examples of

⁵ A main reason is that there is no general way of measuring the similarity between three graphs. There are some attempts to do so by converting tree graphs into graph spectra (e.g., Shokoufandeh et al. 2005), but the lack of isomorphism between the two kinds of representation prevents the use of graph spectra as a framework for implementing structure similarity.

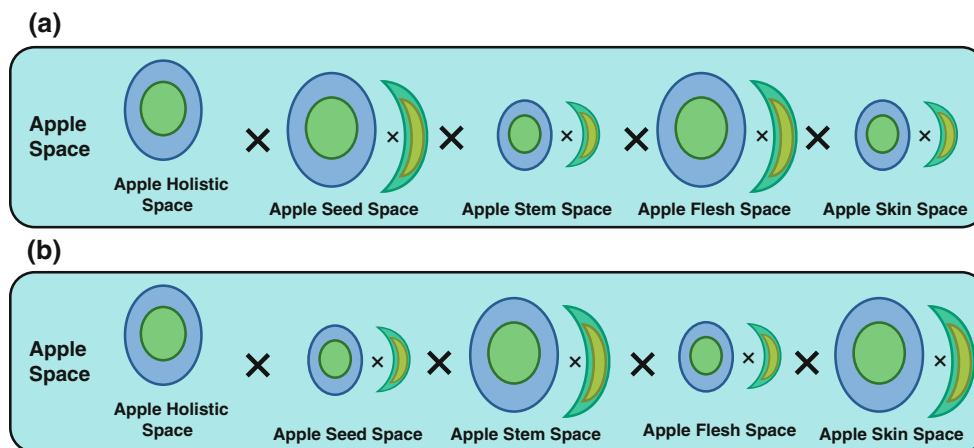


Fig. 5 Examples of different contexts in the apple structure space, where ellipses with smaller sizes denote spaces with less weight in a given context: **a** structure space of apple for a botanist; and **b** structure space of apple for a consumer

intransitive part relations. As others have pointed out (e.g., Johansson 2006), part relations might have different interpretations in different contexts. For instance, the tendency to ascribe transitivity to simple part relations possibly comes from its association with the notion of spatial inclusion, which is transitive in nature. Second, following the general notion of *situated part structures* by Moltmann (1996), we see the formation of part relations and structure spaces as a process that is generally linked to context, experience, and perception, rather than to pure deductive reasoning. We suggest that the structure space of a concept comprises all its *experienced* parts. One tenet of cognitive semantics is that conceptual structures are *embodied*. Thus, the concepts are dependent on bodily experiences and emotions (Gärdenfors 2000). In a broad sense, what dictates whether an object is a direct part of another is the experience (or perception) by the agent of a direct partonomical relationship between the two. Many factors can influence this experience, like the experience of other relations, such as causality and spatial inclusion. More specifically, considering a whole *C*, then a concept *B* that is usually experienced as direct part of *C* is also a component of the structure space of *C*. A concept *A* that is perceived as a direct part of *B* and *C* is a component of the structure spaces of *B* and *C*. In that sense, the structure space of *C* is a very sparse product of the quality domains of its associated parts and is in turn exclusively dependent on one's experience in the world.

However, even discounting transitivity, concepts can still suffer from inflation problems stemming from their legitimate direct parts. Remember that a structure space is a subspace of the conceptual space of the whole. So, the inclusion of the quality domains of “eye” in the structure space of “face” will also bring in the structure space of “eye,” which includes the quality domains of “retina.” This may lead to a situation where inflated concepts

represent themselves *and* all their possible parts. Dimensional filters act as a countermeasure for this issue. As seen in the “[Representing parts and wholes in conceptual spaces](#)” section, dimensional filters are able to select a subspace of the part's conceptual space to compose the whole. In this context, dimensional filters can filter out the aspects of the parts that are not relevant to the whole.⁶ For instance, the relationship between face and eye is mediated by a dimensional filter that blocks the fraction of the conceptual space of eye referring to retina.

Other ontological considerations

Besides partonomies and part relation kinds, there are other ontological aspects of part relations, often raised in the studies of formal theory of parts (Simons 2003; Guizzardi 2005). For instance, wholes can have *mandatory* or *essential* parts. For instance, it is usually said that brain is an essential part of human, while heart is just a mandatory part, i.e., a particular human has to have a particular instance brain, but any instance of heart. Parts can also be *optional*; in the sense, they might or might not appear in the whole to which they relate to. This sort of construction has no parallel in structure spaces in the sense we do not specify any *intensional* mechanism in the representation that allows us to tag parts as essential, mandatory, etc. The main reason for not having it is that such mechanism would remove the plasticity of our representation scheme. For instance, it is not difficult to find counterexamples for many typical illustration cases

⁶ Just as the botanist and the customer put different weights on different apple properties (Fig. 4), a filter can be seen as an *attentional* mechanism that picks out the aspects of the structural complex that are of interest to the user in a particular context.

of essential parts; there are tragic cases of living people without a brain, such as newborns, or one can even conceive the invention of a successful lobotomy procedure, making brain just a mandatory part of the human. Humans can easily adapt their conceptualizations for such situations in a non-monotonic way, where the parts stop being necessary or mandatory. Thus, it is reasonable to expect the existence of a place in the concept representation where these adaptations are possible. We believe this place is the conceptual level. Structure spaces can capture at least a portion of that non-monotonicity in part reasoning by defining levels in which a given part is entrenched in a whole by means of context changes (see “[Same whole, different parts](#)” section). Parts that bear more weight are more entrenched in a given context than others. For instance, the brain might be less important when comparing humans in a crowd.

Furthermore, the concept similarity models have been linked to the notion of *family resemblance* (Rosch and Mervis 1975). Family resemblance is an effect present in categorization, where the instances of a given concept share a great number of similar properties, but not a single property is shared by them all. Conceptual spaces tend to align more with this perspective. Humans have in common they all have a particular brain; however, we can still categorize individuals with no brain as humans because they have many other properties in common (such as other parts). This plays against an essentialist view on the concept representation. From a conceptualistic point of view, the idea of that some parts seem essential or mandatory might make more sense as symbolic level constructions *reflecting* common set of occurrences in the conceptual level. This might as well reveal itself as a point of connection between the conceptual and symbolic frameworks in the future.

Object recognition with structure spaces

Holistic and structure spaces possibly fit in alternative computer implementations of different cognitive tasks involving conceptual representations. In this section, we sketch how an artificial agent (e.g., a robot) could employ holistic and structure spaces for carrying out object recognition.

Object recognition with conceptual spaces

On an abstract level, one can see object recognition as a cyclic process of perception–action. A bottom-up process interprets raw perception stimuli into high-level abstract structures; and a top-down process converts partial high-level interpretations into directed attention in order to

clarify the missing information (i.e., visual search). We can say that a cognitive agent implementing this system achieves interpretation once its internal state stabilizes in a particular set of high-level structures abstracting the perceived stimuli.

In this context, conceptual spaces provide a way describing the concepts in terms of sets of possible observations. Consider an artificial agent equipped with a visual perception system and a conceptual system described in terms of conceptual spaces. Simply put, the general strategy for visual interpretation consists in converting raw visual stimuli into vectors in these conceptual spaces and then checking whether the vectors are similar (close) enough to the concept prototypes; the closest prototype indicates the concept describing the perceived object. For example, consider this agent is equipped with the conceptual space of the (holistic) apple as presented in Fig. 1, as well as a similar space for pear. When this agent is presented with visual image of an *apple*, the visual stimuli are converted into a vector in a conceptual space formed by quality domains related to the visual system, such as the shape space, the color space, the texture space, and so on. This vector is a representation of the perceived object. In order to recognize the object as an instance of “apple,” it is sufficient to check whether the perceived vector is *located inside the properties regions* that define “apple” in each domain. In this case, high-level classification is reduced to a relatively simple verification of geometric inclusion of a point in a region. We can further simplify this process by reducing geometric inclusion calculation to a distance measurement from the concept prototype. However, given the noise inherent to visual perception, it is more likely that more than one concept will be activated by the input vector, e.g., the agent might not be able to tell whether the object was an apple or a pear. If more information is necessary in order to achieve classification, then it is possible to use the candidate concepts to redirect perceptual attention in order to gather better information about the perceived object, i.e., in a top-down process. For instance, the mismatch between the texture of apple and pear (i.e., empty intersection in the texture domain) might be used to guide attention to get closer view of the surface of the object being observed. The new perceptual information complements the previous observation by refining its vector in the conceptual space, restarting the bottom-up processing and closing the agent’s recognition loop.

Using structure information

Structure spaces can improve the previous scheme by allowing the definition of independent holistic and structural processing strategies, which might bring improvements in how different stimuli are classified in

certain situations. A possible implementation is to consider the previous algorithm as *the holistic strategy* for object recognition and link it to a parallel *structural strategy*. Consider an agent is now equipped with the conceptual space of apple (and pear) as in the Fig. 2, i.e., with holistic and structural descriptions. As we have seen in “[Parts, wholes, and cognition](#)” section, holistic and structural processes occur in parallel in object recognition of humans, with one or other having some speed advantage depending on the context (Love et al. 1999). For the sake of the following example, consider the holistic processes have a slight speed advantage. Likewise, when this agent is presented with a *whole* apple, the holistic strategy is triggered, using the holistic space of apple (and other concepts) as basis for classification. As soon as the input stimulus is recognized as being holistically an apple or a pear, then the structural strategy can start in parallel trying to disambiguate between these candidate concepts. This strategy takes structural information encoded in the parts of apple and pear in order to refocus perceptual attention from the whole toward specific parts of the object being perceived. Let us assume pears have slightly longer *stems* than apples. The agent can use *stem* to shift attention to the appropriate locus of the stem on the perceived object. When the perception is shifted to a part, visual processing is primed to process a part, instead of a whole. The perception of the part result in the augmentation holistic vector perceived earlier, receiving values in other domains corresponding to the perceived part, such as part shape, texture and color, and configuration. In our example, the vector now encodes information about the whole apple being observed, as well its stem. This vector is then matched against structural fragment of the candidate concepts. If they are close enough to these concepts in the structure space (i.e., to their structural prototypes), they are kept as candidates. In this case, the concept of “pear” could be discarded as the stem part will now fail to fall in the correct property regions. The perception/action loop will keep running until no more distinguishing properties or parts can be found and concepts disambiguated; then classification is said to be achieved.

It is important to mention again that the stimuli segmentation is mainly governed by the mechanisms outside structure spaces. From a cognitive point of view, the processes that define what a part is are possibly governed by Gestalt grouping principles (Love et al. 1999), part boundary rules (Hoffman and Singh 1997), part saliency, and so on and so forth. In the computational approach described here, these principles would be implemented in lower processing levels, such as in blob extraction algorithms. Nevertheless, they still might use high-level information to tune perception (such as expected shape information).

Context effects can also play a role in this process through external systems. If we can keep a short-term

memory with recently perceived objects, then these can prime the contextual space, influencing the weights of the vector components being compared during the high-level processing. Consider an agent is trying to distinguish between many similar sailing ships with triangular sails. If a novel ship now appears with rectangular sails, the agent will tend to shift more weight to the “sail” part, in order to better discriminate between the objects being observed.

The mention to short-term memory highlights the fact that more complex cognitive architectures can be combined with holistic and structure spaces in order to specify more complex behavior. For instance, computer-oriented cognitive architectures such as the one proposed by Chella et al. (2001) could benefit from our framework, allowing them to deal with parts in the conceptual level. Furthermore, implementations of the structural analogy models such as BRIDGES (Tomlinson and Love 2006) and DORA (Dumas and Hummel 2010) could benefit of having structure spaces as an underlying conceptual structure. For instance, while BRIDGES provides a good processing strategy for measuring structural similarity between exemplars and new stimuli, it ignores holistic similarities, which is accounted by the holistic space in our framework. DORA, on the other hand, depends heavily on low-level symbolic feature descriptors. These could be replaced by regions in quality domains, allowing fine similarity comparisons between the object characteristics (i.e., *geons*).

Conclusion

In this article, we have presented a cognitive approach to represent part–whole relations, founded on the theory of conceptual spaces. Parts are associated with the whole in a structure space, where structural similarity can be measured between wholes and types of wholes. The structure space can capture many aspects of part relations. We have discussed different types of part deconstruction for the same whole, prototypical part deconstruction, variations in part structure caused by context, and part hierarchies. We also showed how different constructions of structure spaces could explain some types of part–whole relations.

The framework presented here contributes to the discussion of whether cognitive semantics is necessary for knowledge representation in computation. As it has been argued by Gärdenfors (2004), technologies based on symbolic theories, such as the Semantic Web, should also include representations that take into account cognitive phenomena. Indeed, the difficulties in representing part–whole relations using ontology representation languages (cf. Rector et al. 2005) serve as good arguments for approaching the problem using a cognitive semantics framework like the one presented here.

While the presented study focuses on the concept representation from a point of view of computer implementations, it still might be useful as a source of insights into the inner workings of the human cognition. Mainly, it calls the attention of the combination of two important cognitive phenomena: part–whole processing and similarity effects. While there is already evidence indicating the existence of interplay between these two phenomena in cognition, it would be interesting to test more thoroughly whether it exists also in other contexts, such as in events and in abstract entities, and how they relate to other aspects of cognition.

We showed that parts, wholes, and their relations can be represented in conceptual spaces, albeit the spaces are high-dimensional and more complicated than for other perceptual properties. However, a lot of empirical and mathematical work remains to turn these sketches into practical working models. The next clear step is to investigate ways in which holistic/structure spaces can be formalized as a mathematical model better suited for implementation in computer system. Also, while the high dimensionality may present some challenges, approaches such as the one from Edelman (1998) might serve as a good starting point for implementation.

Acknowledgments Sandro Fiorini is grateful for the support from CAPES Foundation, process number 1444-11-5, and ANP-Petrobras PRH-217. Mara Abel acknowledges the Brazilian Research Council (CNPq) support. Peter Gärdenfors gratefully acknowledges support from the Swedish Research Council for the Linnaeus environment *Thinking in Time: Cognition, Communication and Learning*. Also, we would like to thank Joel Carbonera, Ingvar Johansson, Joel Parthemore, Luan Fonseca Garcia, and anonymous reviewers for their comments.

References

- Adams B, Raubal M (2009) Conceptual space markup language (CSML): towards the cognitive semantic web. In: IEEE international conference on semantic computing, 2009. ICSC'09. IEEE, pp 253–260
- Aisbett J, Gibbon G (2001) A general formulation of conceptual spaces as a meso level representation. *Artif Intell* 133:189–232. doi:10.1016/S0004-3702(01)00144-8
- Alexander RG, Zelinsky GJ (2012) Effects of part-based similarity on visual search: the Frankenbear experiment. *Vis Res* 54:20–30. doi:10.1016/j.visres.2011.12.004
- Alvarez GA (2011) Representing multiple objects as an ensemble enhances visual cognition. *Trends Cogn Sci* 15:122–131. doi:10.1016/j.tics.2011.01.003
- Augustine E, Smith LB, Jones SS (2011) Parts and relations in young children's shape-based object recognition. *J Cogn Dev* 12:556–572. doi:10.1080/15248372.2011.560586
- Behrmann M, Peterson MA, Moscovitch M, Suzuki S (2006) Independent representation of parts and the relations between them: evidence from integrative agnosia. *J Exp Psychol Hum Percept Perform* 32:1169–1184. doi:10.1037/0096-1523.32.5.1169
- Biederman I (1987) Recognition-by-components: a theory of human image understanding. *Psychol Rev* 94:115–117. doi:10.1037/0033-295X.94.2.115
- Chaffin R, Herrmann DJ, Winston M (1988) An empirical taxonomy of part-whole relations: effects of part-whole relation type on relation identification. *Lang Cogn Process* 3:17–48. doi:10.1080/01690968808402080
- Chella A, Frixione M, Gaglio S (2001) Conceptual spaces for computer vision representations. *Artif Intell Rev* 16:137–152. doi:10.1023/A:1011658027344
- Doumas LAA, Hummel JE (2010) A computational account of the development of the generalization of shape information. *Cogn Sci* 34:698–712. doi:10.1111/j.1551-6709.2010.01103.x
- Edelman S (1998) Representation is representation of similarities. *Behav Brain Sci* 21:449–467
- Farah MJ (1992) Is an object an object an object? Cognitive and neuropsychological investigations of domain specificity in visual object recognition. *Curr Dir Psychol Sci* 1:164–169
- Fiorini SR, Abel M, Scherer CMS (2013) An approach for grounding ontologies in raw data using foundational ontology. *Inf Syst* 38:784–799. doi:10.1016/j.is.2012.11.013
- Förster J (2009) Relations between perceptual and conceptual scope: how global versus local processing fits a focus on similarity versus dissimilarity. *J Exp Psychol Gen* 138:88–111. doi:10.1037/a0014484
- Foster DH, Gilson SJ (2002) Recognizing novel three-dimensional objects by summing signals from parts and views. *Proc R Soc Lond Ser B Biol Sci* 269:1939–1947. doi:10.1098/rspb.2002.2119
- Gärdenfors P (1997) Symbolic, conceptual and subconceptual representations. In: Cantoni V, Gesù VD, Setti A, Tegolo D (eds) *Human and machine perception*. Springer, Berlin, pp 255–270
- Gärdenfors P (2000) *Conceptual spaces: the geometry of thought*. The MIT Press, Cambridge
- Gärdenfors P (2004) How to make the semantic web more semantic. Formal ontology in information systems. In: Proceedings of the third international conference FOIS-2004. IOS Press, pp 17–34
- Gärdenfors P, Zenker F (2011) Using conceptual spaces to model the dynamics of empirical theories. In: Olsson EJ, Engqvist S (eds) *Belief revision meets philosophy of science*. Springer, Berlin, pp 137–153
- Garner WR (1974) *The processing of information and structure*. Lawrence Erlbaum, Oxford
- Gerstl P, Pribbenow S (1995) Midwinters, end games, and body parts: a classification of part-whole relations. *Int J Hum Comput Stud* 43:865–889. doi:10.1006/ijhc.1995.1079
- Goldstone RL (1994) The role of similarity in categorization: providing a groundwork. *Cognition* 52:125–157
- Graf M (2006) Coordinate transformations in object recognition. *Psychol Bull* 132:920–945. doi:10.1037/0033-2909.132.6.920
- Guarino N (1998) *Formal ontology and information systems*. Formal ontology in information systems. IOS press, Amsterdam, pp 3–15
- Guarino N, Pribbenow S, Vieu L (1996) Modeling parts and wholes. *Data Knowl Eng* 20:257–258. doi:10.1016/S0169-023X(96)00009-2
- Guizzardi G (2005) *Ontological foundations for structural conceptual models*. CTIT, The Netherlands
- Hoffman DD, Singh M (1997) Saliency of visual parts. *Cognition* 63:29–78. doi:10.1016/S0010-0277(96)00791-3
- Hummel JE, Biederman I (1992) Dynamic binding in a neural network for shape recognition. *Psychol Rev* 99:480–517. doi:10.1037/0033-295X.99.3.480
- Jäger G (2008) The evolution of convex categories. *Linguist Philos* 30:551–564. doi:10.1007/s10988-008-9024-3
- Johansson I (2006) Formal mereology and ordinary language - Reply to Varzi. *Appl Ontol* 1:157–161

- Johansson I (2012) Natural science and mereology. In: Hans B, Johanna S, Guido I (eds) *Handbook of Mereology*. Philosophia Verlag, München, Germany
- Ligozat G, Condotta J-F (2005) On the relevance of conceptual spaces for spatial and temporal reasoning. *Spat Cogn Comput* 5:1–27. doi:10.1207/s15427633scc0501_1
- Love BC, Roudner JN, Wisniewski EJ (1999) A structural account of global and local processing. *Cogn Psychol* 38:291–316. doi:10.1006/cogp.1998.0697
- Maddox WT (1992) Perceptual and decisional separability. Multidimensional models of perception and cognition. Lawrence Erlbaum Associates Inc., Hillsdale, pp 147–180
- Marr D (1982) *Vision: a computational investigation into the human representation and processing of visual information*. W. H. Freeman, San Francisco
- Marr D, Nishihara HK (1978) Representation and recognition of the spatial organization of three-dimensional shapes. *Proc R Soc Lon Ser B Biol Sci* 200:269–294. doi:10.1098/rspb.1978.0020
- Mash C (2006) Multidimensional shape similarity in the development of visual object classification. *J Exp Child Psychol* 95:128–152. doi:10.1016/j.jecp.2006.04.002
- Melara RD (1992) The concept of perceptual similarity: from psychophysics to cognitive psychology. *Psychophysical approaches to cognition*. North-Holland, pp 303–388
- Moltmann F (1996) A new notion of part structure for natural language. *Data Knowl Eng* 20:323–345. doi:10.1016/S0169-023X(96)00011-0
- Newell FN, Sheppard DM, Edelman S, Shapiro KL (2005) The interaction of shape- and location-based priming in object categorisation: evidence for a hybrid “what + where” representation stage. *Vis Res* 45:2065–2080. doi:10.1016/j.visres.2005.02.021
- Peissig JJ, Tarr MJ (2007) Visual object recognition: do we know more now than we did 20 years ago? *Annu Rev Psychol* 58:75–96. doi:10.1146/annurev.psych.58.102904.190114
- Pentland AP (1986) Perceptual organization and the representation of natural form. *Artif Intell* 28:293–331. doi:10.1016/0004-3702(86)90052-4
- Rector A, Welty C, Noy N, Wallace E (2005) Simple part-whole relations in OWL ontologies. <http://www.w3.org/2001/sw/BestPractices/OEP/SimplePartWhole/index.html>. Accessed 3 Feb 2012
- Rosch E (1978) Principles of categorization. In: Margolis E, Laurence S (eds) *Concepts: core readings*. The MIT Press, Cambridge, pp 189–206
- Rosch E, Mervis CB (1975) Family resemblances: studies in the internal structure of categories. *Cogn Psychol* 7:573–605. doi:10.1016/0010-0285(75)90024-9
- Shepard RN (1987) Toward a universal law of generalization for psychological science. *Science* 237:1317–1323. doi:10.1126/science.3629243
- Shokoufandeh A, Macrini D, Dickinson S et al (2005) Indexing hierarchical structures using graph spectra. *IEEE Trans Patt Anal Mach Intell* 27:1125–1140. doi:10.1109/TPAMI.2005.142
- Simons P (2003) *Parts: a study in ontology*, [Repr.]. Clarendon Press, Oxford
- Simons P (2006) Real wholes, real parts: mereology without algebra. *J Philos* 103:597–613
- Smith LB (2009) From fragments to geometric shape: changes in visual object recognition between 18 and 24 months. *Curr Dir Psychol Sci* 18:290–294. doi:10.1111/j.1467-8721.2009.01654.x
- Son JY, Smith LB, Goldstone RL (2008) Simplicity and generalization: short-cutting abstraction in children’s object categorizations. *Cognition* 108:626–638. doi:10.1016/j.cognition.2008.05.002
- Tomlinson M, Love BC (2006) Learning abstract relations through analogy to concrete exemplars. In: *Proceedings of the cognitive science society*. Lawrence Erlbaum Associates, Mahwah, pp 2269–2274
- Tversky A (1977) Features of similarity. *Psychol Rev* 84:327–352. doi:10.1037/0033-295X.84.4.327
- Tversky B (1989) Parts, partonomies, and taxonomies. *Dev Psychol* 25:983–995
- Tversky B, Hemenway K (1984) Objects, parts, and categories. *J Exp Psychol Gen* 113:169–193. doi:10.1037/0096-3445.113.2.169
- Ullman S (2000) *High-level vision: object recognition and visual cognition*. MIT Press, Cambridge
- Varzi AC (2006) A note on the transitivity of parthood. *Appl Ontol* 1:141–146
- Varzi A (2011) Mereology. In: Edward NZ (ed) *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition). <http://plato.stanford.edu/archives/spr2011/entries/mereology/>
- Winston ME, Chaffin R, Herrmann D (1987) A taxonomy of part-whole relations. *Cogn Sci* 11:417–444. doi:10.1016/S0364-0213(87)80015-0
- Wu R, Mareschal D, Rakison DH (2010) Attention to multiple cues during spontaneous object labeling. *Infancy* 16:545–556. doi:10.1111/j.1532-7078.2010.00061.x
- Zhu SC, Yuille AL (1996) FORMS: a flexible object recognition and modelling system. *Int J Comput Vis*. doi:10.1007/BF00208719