Journal of Experimental & Theoretical Artificial Intelligence

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/teta20

Analogy as a search procedure: a dimensional view

Matías Osta-Vélez & Peter Gärdenfors

JOURNAL OF EXPERIMENTAL & THEORETICAL ARTIFICIAL INTELLIGENCE
https://doi.org/10.1080/0952813X.2022.2125081

Taylor & Francis
Taylor & Francis Group

ARTICLE

OPEN ACCESS Check for updates

Analogy as a search procedure: a dimensional view

Matías Osta-Vélez[a] and Peter Gärdenfors[b]

[a]Center for Mind & Cognition, Institute for Philosophy II Ruhr-Universität, Bochum, Germany; [b]Department of Philosophy and Cognitive Science, Lund University, Lund, Sweden

ABSTRACT
In this paper, we outline a comprehensive approach to composed analogies based on the theory of conceptual spaces. Our algorithmic model understands analogy as a search procedure and builds upon the idea that analogical similarity depends on a conceptual phenomena called 'dimensional salience.' We distinguish between category-based, property-based, event-based, and part-whole analogies, and propose computationally-oriented methods for explicating them in terms of conceptual spaces.

ARTICLE HISTORY
Received 27 December 2021
Accepted 11 September 2022

KEYWORDS
Analogy; categorisation; conceptual spaces; search problems; similarity

Introduction

Analogy is a cognitive mechanism that highlights a similarity relation between objects or situations that are in principle different. For instance, arteries can bear an analogical relation to highways since they share a salient feature (their transportation function) against a rich background of differences. Likewise, a woman breastfeeding her baby and a bird allofeeding a hatchling are analogical events even if the individuals and processes involved are all highly different.

The semantics of analogical statements is rather peculiar. Unlike factual statements, which convey information about states of affairs in the world, analogies seem to have a primarily epistemic goal: They enrich and structure our conceptual knowledge by pointing out similarity relations across seemingly distant fields of knowledge and thus play an organisational role in the formation of abstract categories (see Gentner & Hoyos, 2017).

A significant body of research in the cognitive sciences supports the latter idea. It has been shown that analogy mechanisms play a crucial role in language acquisition (e.g. Behrens, 2017), category learning (e.g. Gentner & Hoyos, 2017; Tomlinson & Love, 2006), teaching (e.g. Vosniadou, 1995), scientific reasoning and discovery (e.g. Gentner, 2002; Oppenheimer, 1956), and AI (e.g. Barbot et al., 2019; Jani & Levine, 2000; Mitchell, 2021).

Two different types of analogical structures have monopolised the attention in the last decades: *direct* and *composed* analogies. The former compares an individual source with an individual target, like in the examples given above. Composed analogies, on the other hand, compares two pair of objects according to some salient relation between the elements in each pair. For instance, the sentence 'the foot is to the leg as the hand is to the arm' is a composed analogy since the salient (mereological) relation between *foot* and *leg* is 'symmetrical' to the mereological relation between hand and arm. Within the cognitive sciences, most studies concern direct analogies. In this article, however, we focus on composed analogies. We will use the notation "$A : B : C : D$" for a composed analogy where the pair $A : B$ is compared to the pair $C : D$.

Our main goal is to show that the theory of conceptual spaces (Gärdenfors, 2000, p. 2014) can provide an explanatory framework for how humans judge composed analogies. Building on

CONTACT Matías Osta-Vélez matiasosta@gmail.com Center for Mind & Cognition, Institute for Philosophy II Ruhr-Universität, Bochum, Germany

an early geometrical model of analogical reasoning known as the 'parallelogram model' (Rumelhart & Abrahamson, 1973), we will show that composed analogies can be analysed in terms of similarity relations across dimensions that are salient for the elements involved in the analogies.

Our approach is based on a novel way of understanding analogies as search problems with the following general structure:

Search space: A set of concepts in a lexicon L;

Initial state: $A : B : C : X$ (with X unknown);

Goal condition: Find (at least) one element X in L such that the semantic relation in $A : B$ is replicated in $C : X$;

Search algorithm: To be defined after an analysis of the kind of semantic relation in the initial state;

Final state: A concept (or preference order of concepts) satisfying the goal condition.

The process of solving an analogy $A : B : C : X$ involves the following steps:

(1) Identifying the semantic relation in $A : B$ and restricting the search space accordingly in the light of C.

As we will show below, the semantic relation can be of varying types. We will consider (i) categorical ("dimensional") relations (e.g., *tuna* : *shark* or *hot* : *cold*), (ii) property-category relations (e.g., *yellow* : *lemon*), (iii) event-based relations (e.g., *open.door* : *closed.door*), and (iv) part-whole relations (e.g., *foot* : *leg*).

Step 1 will tell us what to look for and where to look for it. For instance, the analogy *red* : *apple* : *yellow* : X is about a fruit category and a prototypical property in the colour dimension. The search space for X will be the conceptual space of fruits and the goal condition will be satisfied by fruit categories for which the colour yellow is prototypical.

(2) Specifying the search algorithm.

In this article, we present four search algorithms for different semantic relations. The kind of relations will define the similarity space in which the algorithm works.

(3) Applying the search algorithm to identify one or more categories that meet the goal condition.

The following section discusses some theoretical aspects of analogical relations in the light of established ideas in cognitive science. In Section 3, we introduce the theory of conceptual spaces. In Section 4, we turn to category-based analogies of the type *dog* : *wolf* : *cat* : *lynx* and introduce the parallelogram model as a basis for a search algorithm. As the name suggests, such analogies are based on comparing categories. Section 5 is devoted to property-based analogies of the type *red* : *apple* : *yellow* : *banana*; we present a typicality criterion to be used in the search procedure. For these analogies, a category is compared to a characteristic property. In Section 6, we turn to event-based analogies, exemplified by *horse* : *gallop* : *man* : *run*. For this type, force and result vectors will be central. The final type of analogy is part-whole analogies of the form *hand* : *arm* : *foot* : *leg*. We argue in Section 7 that also such meronomic relations can be modelled by using conceptual spaces. Our focus in this paper is to model human reasoning, but the models we present are amenable for implementations in AI systems via the search algorithms we identify. As we will show for all four types of analogies, using conceptual spaces as a basis for the models opens up for new types of computational implementations.

Similarity and its problems

It is widely acknowledged that similarity is the fundamental relation behind analogical processes (see Cummings, 2020; Gentner & Markman, 1997; Holyoak, 2012). However, similarity is a problematic notion (see Goodman, 1972; Smith, 1989; Tversky, 1977) and many of its puzzling features permeate the study of analogy. Consider the sentence 'Whales are like sharks.' It can be correct if the comparison is in terms of perceptual similarity but wrong if we focus on biological and taxonomical features. In general, similarity judgements require extra information for specifying in which respects the things compared are similar.

This becomes particularly clear when we look at analogies, since they target similarity relations that stand out from a rich background of differences. For instance, consider the sentences (i) 'Leopards are like cheetahs' and (ii) 'Sport cars are like cheetahs.' (i) expresses a straightforward 'surface' similarity, while (ii) expresses an analogical relation in which the similarity among categories is reduced to one salient shared feature: speed. Interpreting (ii) as a meaningful statement depends on the agent's semantic knowledge and on her ability to grasp non-literal similarity relations (e.g., Gentner et al., 1995; Ortony, 1979). This shows that a theory of analogy requires a specification of what kind of similarity relation underlies the process.

The most influential approach to analogical similarity is the structural alignment view (SAV), pioneered by Gentner and colleagues (see Gentner & Holyoak, 1997; Gentner, 1983; Markman & Gentner, 2000). According to SAV, analogy depends on a mapping between two representational structures consisting of objects, attributes (unary predicates) and high-order relations. The mapping returns a set of *commonalities* (shared characteristics or objects among the structures), a set of *alignable differences* (pairs of features or objects that are in principle different but can be still put in correspondence), and a set of *non-alignable differences* (elements of some of the structures that failed to be mapped because they are not in either of the previous two sets. It is crucial to this theory that the alignable differences are more salient than the non-alignable differences (Gentner & Markman, 1997, p. 50).

A central idea in SAV is that analogical similarity is predominantly relational, that is, it focusses more on the mapping of high-order relations than on object or attribute-matches (see Gentner et al., 1995). Gentner and some other authors go as far as to assert that our ability to handle relational predicates in analogies is what marks the difference between human and non-human cognition (Gentner, 2003). To illustrate this idea, consider the scenes depicted in Figure 1.

At the object/attribute level, images *b* and *c* are highly similar: they have many common properties and only one alignable difference: the spatial configuration of the objects. On the other hand, a and b have many alignable differences, like shape and colour in every object mapped, but also an important commonality: a similar relational structure (*shorter.object-left.of.taller.object*). This last commonality makes *a* and *b* analogues, while *b* and *c* share only a relation of *literal* similarity.

The above example shows how analogical similarity is of a special kind and cannot be captured by the intuitive idea of similarity as a function of a positive difference between shared and nonshared features (see Goldstone et al., 1991). SAV succeeds in proving this point, as well as in advancing a richer notion of similarity that distinguishes between kinds of differences. Nonetheless, the approach has some issues. First, SAV does not build on any specific theory of conceptual structure.

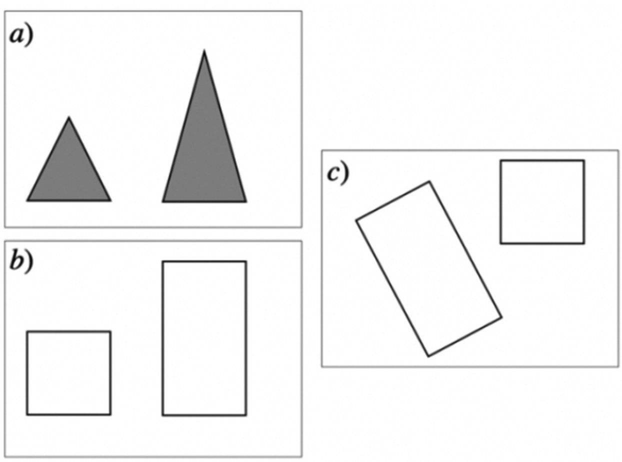

Figure 1. Three pairs of analogous objects.

As is the case in most of computer science, SAV instead analyzes analogical representations from a propositional perspective in which predicates are distinguished only according to their arity. From this perspective, adjectives and nouns are all the same kind of terms, while it is rather clear that, from a cognitive perspective, they play different conceptual roles (Gärdenfors, 2000, p. 2014). This difference can be easily spotted in analogies. Consider, for instance, 'red is to apples as yellow is to lemons' and 'dog is to wolf as cat is to lynx.' The former relates two object categories (*apple* and *lemon*) with two of their prototypical properties. In this sense, the analogy builds on the internal structure of the categories. The latter, on the other hand, compares four object categories on the same conceptual level sharing a common immediate superordinate category. As we will show later, this kind of difference can be easily accounted for if we build the theory of analogy on a theory of conceptual structure, instead of a theory of propositional form.

Second, while it is true that relational predicates play a central role in analogies, it is not clear that we can fully reduce analogical similarity to relational similarity. Analogies like 'sports cars are like cheetahs' do not build on any relational concept, but on the salience of a shared dimension. Furthermore, most relational predicates can be analysed in terms of single dimensions.[1] For instance, comparative adjectives like $\mathit{Taller}(x, y)$ or $\mathit{Younger}(x, y)$ express a difference with respect to one shared dimension of the objects (in these cases, *height* and *age* dimensions). The fundamental factor for analogies that use them is, however, not the structure of the relation but our knowledge of the dimension they compare.

Third, computational implementations of relational similarity are difficult to construct in a systematic manner. Using the distance functions of conceptual spaces, as they are used in the criteria for different kinds of analogies presented in this article, opens up new forms of implementation. We outline how this can be done by describing different search procedures. This is important for systems of artificial intelligence that aim at mimicking human reasoning, since analogies play such a central role there. Furthermore, our analysis will indicate that not only one, but several computational methods are required to handle the different types of analogies.

Dimensional salience

The approach advanced here builds on these last two observations. First, we claim that composed analogies need to be analysed in the light of a theory of conceptual structure. Second, we propose that in most cases, analogical similarity depends on dimensional salience, more precisely, on identifying one or more dimensions that will serve as a frame of comparison for the categories in the analogy. The degree of salience of these dimensions for the given categories correlates with the analogy's 'quality' or 'aptness.' This last idea is rather straightforward. Consider the following analogies:

1. *dog* : *puppy* : *cat* : *kitten*
2. *sweet* : *apple* : *sour* : *lemon*
3. *rabbit* : *lion* : *tuna* : *shark*
4. *hot* : *warm* : *cold* : *cool*

Each of these analogies consists of projecting a salient semantic relation among the categories of the first pair into the categories of the second pair. This relation depends on identifying one or more dimensions of the categories that can serve as 'analogy factors.' In (1), the analogy factor is the *age* dimension, in (2) is *taste*, in (3) *size* and *ferocity*, and in (4) the *temperature* dimension. The analogy factor is generally differential: it selects a dimension in which the categories of the first pair have significantly different values.

A challenge while evaluating an analogical relation is to identify, among the many dimensions that can constitute the categories involved, which are the ones that can better bear the analogical relation. For example, *size* can be a good candidate for analogy factor in (3), but *colour* clearly not. In our approach, the dimensions that are going to have priority as potential analogy factors are the most salient dimensions of the categories in the first pair. Such a salience factor is difficult to model in proposition-based computational implementations.

A straightforward prediction of this approach is that the mental processing speed of an analogy will be positively correlated with the degree of saliency of the analogy factors and negatively correlated with the number of dimensions that can be considered potential analogy factors. For instance, (4) is a straightforward analogy because a unique dimension relates its four categories; (3), on the other hand, offers multiple possible dimensions as potential analogy factors and, as a consequence, it has a higher degree of analogical complexity.

While classical approaches tend to look for highly general models of analogy (e.g., Gentner, 1983; Holyoak & Thagard, 1989), our view departs from the idea that analogy is concept-specific. In other words, analogies exploit properties of the representational structures associated to the word classes that appear in them. Since different word classes represent different kinds of concepts (see Gärdenfors, 2014), we need a theory that integrates different sub-models. In the remainder of the article, we propose several possible models for this theory that may be interpreted as search algorithms. In contrast to traditional approaches in logic and computer science where all predicates are treated on a par, we aim to show that dividing them into their different conceptual roles will yield more fruitful computational systems, specified as different search procedures.

Conceptual spaces

Conceptual spaces (CS) is a theoretical framework that represents concepts in terms of geometrical and topological relations (see Gärdenfors, 2000, p. 2014). CS assumes that concepts have an internal structure that is grounded in *quality dimensions* and *domains*. Many quality dimensions are elemental features of perceived stimuli. Sounds, for instance, are perceived as having three dimensions: *loudness*, *pitch*, and *timbre*; each of these dimensions can be independently used as a framework for comparing different sounds, but they are fully integrated in each individual stimulus (we cannot perceive a sound without some of these dimensions). In these cases, we say that these dimensions are 'integral.' On the other hand, when we can attribute a 'value' to a dimension independently of any other dimensions, we talk about 'separable' dimensions. For instance, the length of a desk is independent of its solidity, thus solidity and length are two separable dimensions. Some dimensions such as kinship relations, economic variables and theoretical physical variables are not perceptual. In many cases, however, the mathematical structure of the dimensions can be precisely described.

In this article, we do not discuss the origins of the dimension but we take them as given to the system that is to compute the validity of composed analogies. The origins of dimension have been investigated in Gärdenfors (2000, p. 2014; Gärdenfors, 2021).

A crucial point in CS is that dimensions can be represented as instantiating different geometrical structures. For instance, *length*, *weight*, and *loudness* can be represented as half-lines isomorphic to the non-negative real numbers. When we have a set of integral dimensions, their geometrical structures compose into what it is called a 'domain.' Domains are multi-dimensional structures able to represent individual stimulus as points with coordinates in each of their constituent dimensions. For instance, the colour domain is a geometrical structure composed by three integral dimensions: *hue*, *intensity*, and *brightness* (see Figure 2), any instance of a colour can be represented as a point in that structure with a value in each of these dimensions. Colour terms like 'red,' 'blue,' 'yellow,' etc. can be represented as convex sets of points ('regions') in the colour domain. In CS, such

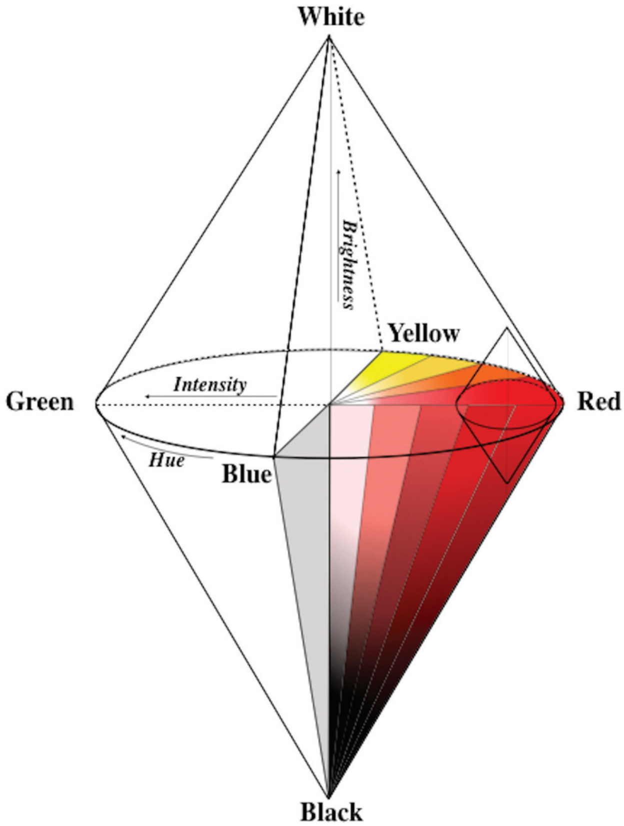

Figure 2. The color domain.

regions are called 'properties' and they can be seen as special cases of concepts due to the fact that they are restricted to a single domain.

More generally, concepts are structures of greater complexity than domains since they often require multiple domains and dimensions to be correctly represented. For instance, our concept *apple* requires the colour domain plus dimensions for representing qualities such as *colour*, *ripeness*, *taste*, *texture*, *size* and *shape*. One way of explicating this concept is by considering the product space of all the previous dimensions; the concept apple will be a convex sub-region of such a space, as illustrated in Figure 3.

Generalising these ideas, we define a conceptual space as a collection of one or more domains with a distance function – a metric – which represents properties, concepts, objects and their similarity relationships. Similarity between concepts and objects can then be easily estimated since it is a monotonically decreasing function of their distance within the space (Nosofsky, 1992; Shepard, 1987). The distance function can vary, the most common one is the Euclidean, but also Manhattan, other Minkowski and polar metrics may be appropriate in different contexts (see Gärdenfors, 2014; Johannesson, 2002; Shepard, 1964).

An important advantage of this framework is that it can represent the prototypical structure of concepts (see Gärdenfors, 2000; Rosch, 1983), that is, the idea that there is one instance which represents the concept better than any other. Within convex regions, one can take some specific point – or set of points in some cases – as the prototype of a category. As a result, and using the built-in metric of the space, one can measure the degree of typicality of any member of a category by estimating its distance to the prototype. For example, focal colours are often considered in cognitive

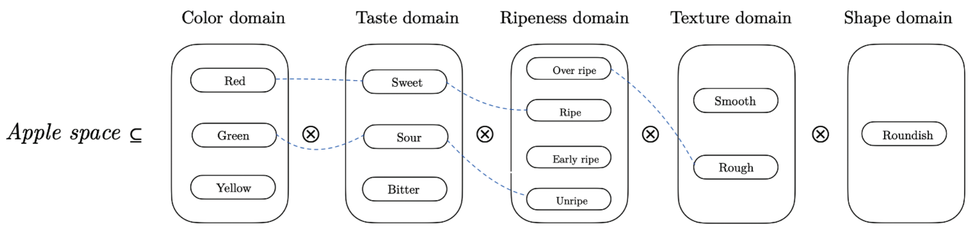

Figure 3. 'Apple space.' the dotted lines represent correlations between properties for the concept *apple*.

science and linguistics as prototypes of the colour space (see Douven, 2019; Rosch, 1975). Assuming the prototypical structure of concepts does not require that an actual object represents the prototype. Conceptual spaces can represent every possible object falling under a concept. Gärdenfors (2000) claims that a prototype can correspond to a partial vector containing only information about the values of the most relevant dimensions for the concept.

Much more can be said about this approach but, for the purposes of this paper, we will focus on two aspects of CS: that concepts have a dimensional and prototypical structure, and that there exists a distance function that allow us to compare concepts and objects in the space. In what follows, we will show how to put these ideas to work in order to analyse analogies.

The parallelogram model

One of the earliest models of analogy was developed by Rumelhart and Abrahamson (1973) under the assumption that analogical inference builds on features of the organisation of knowledge in semantic memory. According to the authors, it is possible to express analogical similarity as a function of the semantic distance between categories represented as points in a multidimensional space. In particular, they claimed that analogies of the form $A : B : C : D$, where the elements are objects classes, must follow a 'parallelogram rule' according to which the vectorial distance between categories A and B must be equal (or highly similar) to the vectorial distance between C and D (see Figure 4).

The model understands analogical inference as a choice problem in which agents first represents the categories A, B, and C as points in a multidimensional space and select for a category D that is the best option for fulfiling the parallelogram rule. Roughly, this rule consists of fixed vector addition and subtraction operations: $D = (B - A) + C$.

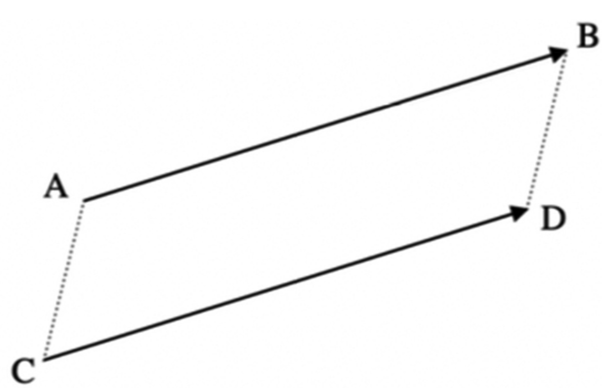

Figure 4. Vector relations among the four analogy terms represented as points in a space, according to Rumelhart and Abrahamson (1973).

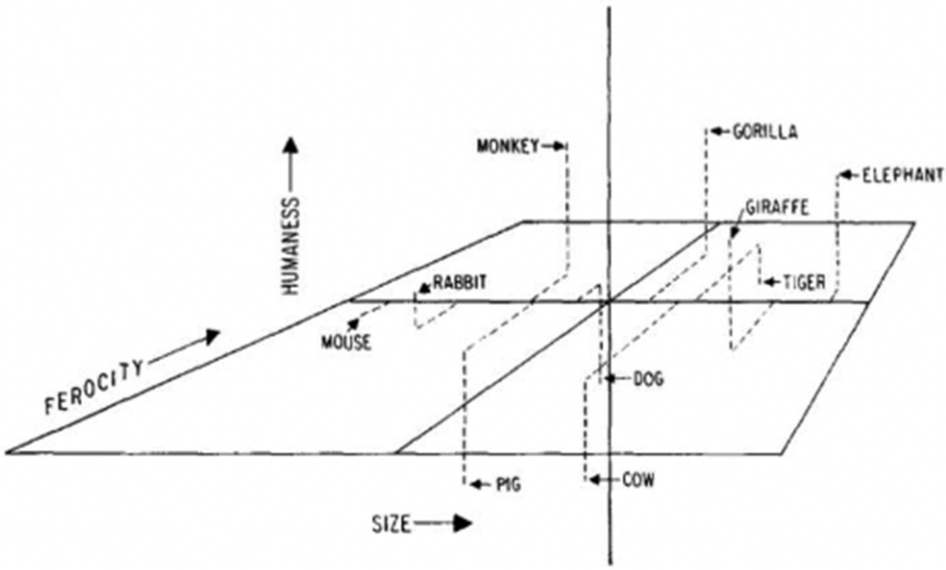

Figure 5. 'Mammal space' organized around the *humanness*, *size*, and *ferocity* dimensions. From Rumelhart and Abrahamson (1973, p. 3).

In a series of experiments using Henley's (1969), p. 3-dimensional mammal-space (see Figure 5 for some examples), Rumelhart and Abrahamson showed that when presented with analogy problems like *monkey* : *pig* : *gorilla* : *X*, with *rabbit*, *tiger*, *cow*, and *elephant* as alternatives for *X*, subjects rank the four options following the parallelogram rule. The parallelogram model predicts that *cow* is the preferred solution. Their experiment clearly supported the model.

Overall, we believe that this model provides the correct insight into category-based analogies. However, it has two important limitations. First, it lacks explanatory depth because it does not allow for the identification of the analogy factors nor does it provide a notion of analogical similarity. Second, it can only deal with analogies in which all terms are at the same conceptual level (called here 'category-based analogies'); analogies like *swim* : *fish* : *fly* : *bird* or *bear* : *mammal* : *lizard* : *reptile* cannot be analysed by this model since not all categories in them have the same vector representation and are consequently incomparable.

Category-Based analogies in conceptual spaces

This section presents a generalised version of the parallelogram model which follows the semi-algorithmic approach described in the introduction. The basic idea is that the conceptual space in which the vectorial comparison is carried out is not fixed, but rather depends on the dimensions that are taken as analogy factors in each specific analogy.

In our model, the categories in a category-based analogy $A : B : C : D$ are convex regions of a common conceptual space M (written "$C(M)$"), since they are all at the same conceptual level. For the sake of simplicity, we assume that each of these categories has a precise prototype represented by a point in the space. For category X, we refer to that point as p^X. The following describes the main steps of the search procedure.

(i) Given a composed analogy $A : B : C : X$, with X unknown, the first step in the process consists in finding the smallest conceptual space $C(M)$ such that $A, B, C \subset C(M)$. This space corresponds to the immediate superordinate category of A, B, and C. For instance, in *tiger* : *rabbit* : *eagle* : X, $C(M)$ will be the *animal-space* but, in *tiger* : *rabbit* : *truck* : X, $C(M)$ will be *thing-space*. $C(M)$ will be the search space in which the algorithm will operate. Notice that the number of dimensions apt for establishing an analogical comparison depends on the specificity of $C(M)$ (that is, its place in Rosch's (1978) vertical level of categorisation). Dimensions that are available for animals in *tiger* : *rabbit* : *eagle* : X like *diet*, *ferocity*, or *humanness* cannot be applied to things in *tiger* : *rabbit* : *truck* : X.

(ii) The second step consists in selecting from $C(M)$ a set of salient dimensions $D_1, D_2, \ldots, D_n$ where the salience is generally determined by the difference between A and B (often only one dimension is relevant). For instance, consider the relation *tiger* : *rabbit* : *eagle* : *robin*. *Ferocity* and *size* are two salient dimensions of *animal*, since an important difference can be established between *tiger* and *rabbit* across these dimensions. If these differences can be replicated for categories *eagle* and *robin*, then the analogy is sound. The choice of these dimensions as frame of comparison will generate a 'new' lower-dimensional conceptual space $C(M*)$ with a distance function $d*$.[2] This modulated distance function will be used to compute what we have called 'analogical similarity' and constitutes the main difference with the Rumelhart and Abrahamson's parallelogram model.

(iii) The last step of the search algorithm is the application of the parallelogram rule on $C(M*)$ for choosing the optimal solution to X in $A : B : C : X$. For this, we start with the prototypes p^A,p^B,p^C, and the vector $\overleftarrow{p^A p^B}$ in $C(M*)$, and we find the point $y \in C(M*)$ that is the head of a new vector $\overleftarrow{p^C y}$ that is as close as possible (same direction and magnitude) to $\overleftarrow{p^A p^B}$. The category $X \subseteq C(M)$ that gives the strongest analogical relation will be the one whose prototype p^X is closer to y than any other prototype in $C(M)$, that is, p^X such that $d*(p^X, y) < d*(p^Z, y), \;\; p^Z \in C(M)$.

Let us illustrate this procedure with a toy example. Consider the incomplete analogy *mouse* : *wolf* : *rabbit* : X and a reduced search space with categories *hippo*, *buffalo*, *elephant*, and *gorilla*. M will be the mammal space used by Rumelhart and Abrahamson (1973) (see Figure 5), and the dimensions that will serve as frame of comparison in the space $M*$ will be *size* and *ferocity*, due to the salient difference that the categories *mouse* and *wolf* maintain across them. The *humanness* dimension in Figure 5 is less salient and will not be part of $M*$. Then, in a weighted conceptual space $M*$, a point y will be determined as the head of a vector with tail in the prototype of *rabbit* that is equivalent to the vector formed by the prototypes of *mouse* and *wolf*. Assuming the positions of the prototypes as depicted in Figure 6, the prototype of *buffalo* is the optimal solution to the analogy since it is closer to y than any other prototype in $M*$.

Analogies are not all or nothing, but have degrees of aptness or soundness. For instance, categories that are very close (in the weighted conceptual space) to the optimal choice in a category-based analogy might also be good solutions. In addition to this, it is possible that different sets of dimensions are taken as analogy factors, generating multiple possible sound analogies. We believe that, for most cases, there is a particularly salient set of dimensions that will produce the strongest analogical relation. However, offering a systematic criterion for finding it is rather complicated for it is strongly dependent on the subjects' knowledge of a particular semantic domain, as well as on the semantic intuitions rooted in a community of speakers. Ultimately, finding the set of salient dimensions for a given category is an empirical question.[3]

From categories to properties

As mentioned earlier, an important limitation of Rumelhart and Abrahamson's (1973) model is that it can only deal with analogies at same conceptual level. Consider the following two examples:

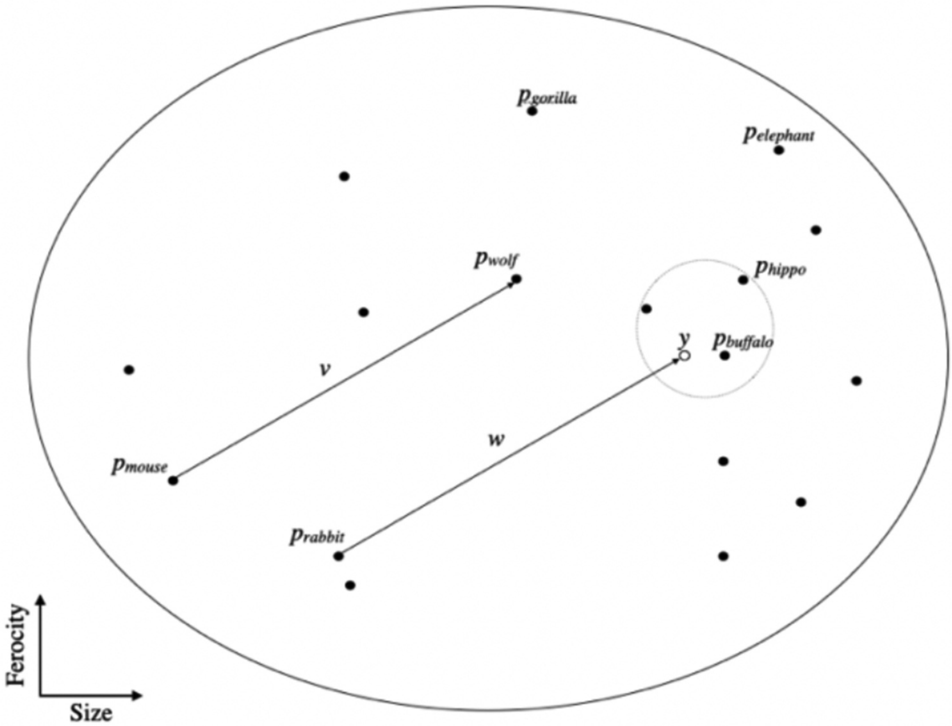

Figure 6. Mammal space organized around the *size* and *ferocity* dimensions. Based on Rumelhart and Abrahamson (1973, p. 3).

5. *apple* : *red* : *banana* : *yellow*

6. *fish* : *swim* : *bird* : *fly*

(5) and (6) are sound analogies, but they cannot be analysed in terms of the parallelogram model. How can we compare a colour with a fruit or an animal with a means of motion? From a formal perspective, there is no way of comparing two vectors from different conceptual spaces.

We call analogies like (5) and (6) 'property-based analogies.' Naturally, since the semantic relation between the pairs of terms in these analogies differs from that characterising category-based analogies, explicating them via a search algorithm requires a different approach. In particular, the search space for *X* will be the set of lexical items associated to common properties of the category in *C*.

Our proposal for property-based analogies is straightforward: We claim that the strength of an analogy depends on two factors, first, on the identification of the dimension(s) that corresponds to the property in the pair and second, on identifying the typicality degree of that property for the category in the pair. In other words, we evaluate the aptness of these analogies by checking that the properties in the pairs are from the same dimension and, with the aid of a typicality condition, that they are similarly expected for the category in the pair. In this sense, an analogy like (6) must be considered as stronger than the variant *fish* : *swim* : *bird* : *walk* because, even if birds can walk, flying is more typical than walking for that category (Osta-Vélez & Gärdenfors, 2022).

Modelling typicality

From the perspective of a search algorithm, the difficulty lies in identifying the typicality degree of properties for a given category. In Osta-Vélez and Gärdenfors (2022), we proposed a way of doing this by using distances in conceptual spaces. Let us recall that the conceptual space of a category M includes all the properties that an object falling under M might have. These properties can be divided between *prototypical* and *non-prototypical*, according to whether or not they are in the prototype of M. Non-prototypical properties can be assigned a degree of typicality by measuring their distance from the prototype (p^M) of the concept in $C(M)$.

Prototypical properties have, by definition, a higher degree of typicality than non-prototypical ones but they are not all equally typical. For instance, while both roundness and redness belong to the prototype of apple, the former is more typical than the latter since a non-round apple is more surprising than a non-red one.[4] Following this idea, an ordering between prototypical properties can be established by measuring the distance to the closest point in the space where a property is not satisfied. Similarly, non-prototypical properties can be assigned a degree of typicality by measuring their distance from the prototype of the concept in $C(M)$. The following criterion systematises these ideas:

Typicality criterion

(i) For any two prototypical properties R_i, R_k in a conceptual space $C(M)$, R_i has a higher typicality degree than R_k *iff* for all $x \in C(M)$ with $\neg R_i(x)$, there exists $y \in C(M)$ such that $\neg R_k(y)$ and $d(y, p^M) < d(x, p^M)$.

(ii) For any two non-prototypical properties R_i, R_k in a conceptual space $C(M)$, R_i has a higher typicality degree than R_k *iff* $\neg R_k$ has a higher typicality degree than $\neg R_i$.

(iii) A prototypical property always has a higher typicality degree than a non-prototypical property.[5]

Now, given an analogy $A : B : C : X$ where A and C are categories and B is a property of A in dimension D, the choice for X which gives the strongest analogical relations is a different property in D whose typicality degree is closer to B's typicality degree than the typicality degree of any other property in D. We predict that if there are various properties in D with the same typicality degree as B for category C, then the analogy will be weaker than for categories for which this is not the case. For example, the analogy *lion* : *beige* : *raven* : *black* must be judged as stronger than *lion* : *beige* : *dog* : *brown* because several colours other than brown are also typical for the category dog.

From properties to categories

Notice that the search procedure changes if the terms in the analogy appear in a different order, say $A : B : C : X$ such that A and C are properties in the same dimension D, and B and X categories. An analogy like *red* : *apple* : *yellow* : X requires the agent to search for a category at the same conceptual level as apple for which yellow is highly prototypical. Instead of typicality degrees, the semantic phenomenon that is relevant here is *diagnosticity*. Diagnostic properties allow us to identify category membership with minimal information (see Chin-Parker & Ross, 2004; Tversky, 1977), for instance, *having gills* is highly diagnostic of the category *fish*. Diagnosticity is deeply related to typicality but not the same: diagnostic properties are typical, but not all typical properties are diagnostic (e.g., sweet is typical of apple but not diagnostic, since several categories in the contrast class share it).

We claim that in cases like these, the search algorithm involves three steps: First, identifying the smallest conceptual space of M such that $B \subseteq C(M)$; second, finding a set of categories in $C(M)$ for which C is typical, and third, choosing from this set the most typical category in $C(M)$. That will be the choice giving the strongest analogical relation. Strictly speaking, we cannot use the typicality criterion given above for *X*, since it is a category and not a property. However, this criterion can be adapted to categories using the same idea.

Event-Based analogies

Let us now analyse a kind of analogy that involves events in one way or another. From a cognitive and semantic perspective, the representational structures behind events are different from those behind nouns and adjectives (For a review, see Gärdenfors, 2014; Papafragou, 2015). As far as we know, they have only been studied to a limited extent (e.g., Goswami & Brown, 1990). 'Event-based analogies,' as we call them, come in two types. One concerns analogies of actions as in the following examples:

7. $woman : wave : tree : sway$

8. $horse : gallop : man : run$

The second type are about results of events. The following two examples come from the list of stimuli used by Goswami and Brown (1990):

9. $box : open.box : bottle : open.bottle$

10. $chocolate : melted.chocolate : snowman : melted.snowman$

Our aim is to show that also event-based analogies can be analysed in terms of conceptual spaces.

We build on the two-vector model of events developed by Gärdenfors and Warglien (2012) (see also, Gärdenfors et al., 2018; Warglien et al., 2012). The two-vector model states that an event is represented in terms of two components – the force (or force pattern) of an *action* that generates the event, and the *result* of its application. Both components are represented as vectors in spaces. In the special case when there is no change, that is, when the result vector is the zero vector, the event is a *state*.

A central feature of events is that they are based on causal relations: An event contains information about an agent who is the cause of an action that leads to a result related to a patient. The result of an event is modelled as a vector representing the change of properties of the patient before and after the event, as illustrated in Figure 7.

To model an event, at least two spaces are needed, an action space and a result space. The action space can be conceived as a space of forces or, more generally, force patterns, acting upon some patient. We speak of a pattern of forces since, for bodily motions, several body parts are involved; and thus, several force vectors are interacting by analogy with Marr and Vaina's (1982) differential equations (for more details, see Gärdenfors & Warglien, 2012). The result space contains dimensions that represent change in the properties of the patient. The spaces represent different types of vectors: Forces have a different nature than changes in properties.

Once the force and results vectors are available, we can use them to model event-based analogies in the same way as in previous sections. For example in the analogy (7) $woman : wave : tree : sway$ the woman is the agent of an event and *wave* can be modelled by a force vector that describes the action of the agent (Gärdenfors & Warglien, 2012; Gärdenfors, 2014). Then in the second part of the analogy, *tree* is the agent and *sway* describes a force vector that is the most similar to a woman waving. A similar account can be given for (8) $horse : gallop : man : run$. The aptness of this analogy lies in the fact that *run* is the motion verb for man whose force pattern-representation resembles the most to *gallop* when the agent

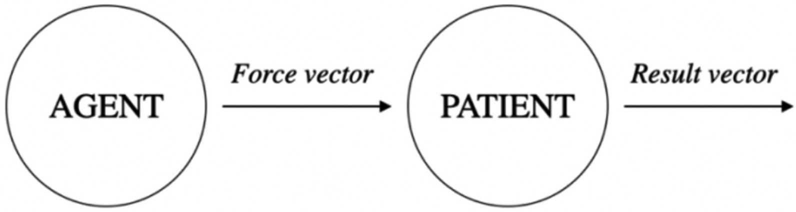

Figure 7. The main components of an event representation.

is *horse*. For similar reasons, we predict that the analogy *bird* : *wing* : *fish* : *fin* will be judged as stronger than *bird* : *wing* : *human* : *arm* since the force patterns of the movements of wings and fins are, in general, more similar than the force patterns of the movements of wings and human arms. Wings and fins are functional parts in the locomotion of the animal and thus exert substantial forces on the surrounding air or water, while arms are typically not used for locomotion.

In the analogy (9) *box* : *open.box* : *bottle* : *open.bottle*, *open box* describes a result of an action performed on the box and the result vector is simply $< box, open.box >$. Then once *bottle* is given the corresponding result vector is $< bottle, open.bottle >$. Thus, the similarity of the result vectors explains the validity of the analogy. A similar analysis can be given for (10) *chocolate* : *melted.chocolate* : *snowman* : *melted.snowman*, and for the rest of the analogies studied in Goswami and Brown (1990). Interestingly enough, Goswami (1992) and Goswami and Brown (1990) argue that children's ability to reason with this kind of analogies depend on their semantic competence, that is, on their mastery of the causal concepts involved in the analogy (*cutting*, *melting*, *opening*, etc.). As they point out, this contradicts Piaget's idea that analogical reasoning is a formal mechanism that children are only able to learn in the formal-operational stage (13–14 years old) (see Piaget et al., 1977). Our approach fully agrees with the former idea. Analogical reasoning is a semantic-based mechanism that exploits properties of conceptual representation, as a consequence no purely logical approach (that is, non-semantic) to this mechanism is likely to be successful (see also, Osta-Vélez & Gärdenfors, 2020, 2022).

It is important to note that the role of domains and dimensions is still central in event-based analogies. According to the conceptual space model of events, the meaning of a verb is 'a convex region of vectors that depend on a single domain' (or dimension) (Warglien et al., 2012, p. 172). For instance, *paint* denotes a change in the colour domain of the patient of the action and *heat* implies a change in the temperature dimension. Even if the actual actions corresponding to verbs imply changes in multiple dimensions of the patients, verbs work by directing the attention to only one of these domains/dimensions. Event-based analogies work in the same way, consider for instance *fire* : *heat* : *ice* : *cool*, the analogy is coherent because temperature is the salient dimension for both of the analogies.

Part-Whole analogies

As with properties and events, meronomic relations have their own underlying representations (see Markman, 1981; Winston et al., 1987) which cannot be fully explained in terms of the above models. Analogies that build on these relations exploit properties of these structures and must be modelled accordingly. In what follows, we sketch a proposal for doing this that uses an extension of the conceptual space theory for part-whole relations recently developed by Fiorini et al. (2014).

Structure space

The main idea behind this proposal is that meronomic information about a concept *M* cannot be represented in a standard conceptual space but needs a complementary space called *structure space*. This space encodes two different kinds of information: (i) information about the parts that objects falling under *M* are supposed to have, and (ii) information about how these parts are arranged in relation to the whole (this is called *configurational* information). The former is stored in the individual conceptual spaces of each of the parts. The latter, is stored in what is called *structure domains*. In this richer notion of a conceptual space, overall conceptual similarity is a function of the similarity measure of the standard conceptual space plus *structural similarity*, which is based on a distance function operating on the structure space.

Our idea is that part-whole analogies depend primarily on structural similarity and thus on the information stored in structure space. However, conceptual information can also be used as a reinforcing analogy factor. Consider the following two analogies:

11. *horse* : *legs* : *car* : *wheels*

12. *leg* : *foot* : *arm* : *hand*

The analogy factors behind (11) are two: One is configurational and concerns a similarity in the spatial arrangement of the pairs *legs* : *horse* and *wheels* : *car*; the second is conceptual and concerns a *transportation function* that both legs and wheels have in relation to the respective wholes. In (12), the only analogy factor is configurational: *Foot* and *hand* occupy the same position in the part-whole structure of *leg* and *arm* respectively. Besides the great overall similarity in the categories in the analogy, no other salient dimension seems to be relevant.

Shape analysis

We will now sketch a possible way of modelling the structural factor in part-whole analogies using ideas from shape analysis. A central aim of shape analysis is to identify structures that remain invariant across the multiple ways in which the shape of an object presents itself to perception. Two things that these structures must capture are the length and directions of the parts of an object which, as is well known in object perception, are often sufficient to categorise it. Marr and Nishihara's (1978) classic approach to describing biological forms accomplishes this by representing the part-hierarchy of an object through cylinder-like modelling primitives. The vector coordinates for the cylinders and their connecting points will generate a multidimensional space of shapes and each particular cylinder-shape corresponds to a point in this space. A metric can be introduced to the space by using a weighted sum over all the dimensions on the different levels. A point in this structure space represents a combination of shape forms and displacements for a given decomposition level, denoting a particular limb configuration.

Another important component in shape analysis is the connectivity of the parts with respect to the whole. One approach that places special emphasis on this is presented by Zhu and Yuille (1996). Their model uses two fundamental primitive shapes: *worms*, which are rectangles with circles at the ends, and *circles* with angles that represent hinges and their flexibility. Like a cylinder, a worm can be described by the two coordinates length and width. A *hinge* between parts can be described by the radius and the angular values. Different shapes can be represented as a composition of instances of these two primitive structures (see Figure 8). As in Marr and Nishihara's model, this structure can be represented by a multidimensional vector.

Given a shape represented in the aforementioned manner, a further structure can be computed by identifying the locus of points equidistant to two or more points in the shape's boundary.[6] This

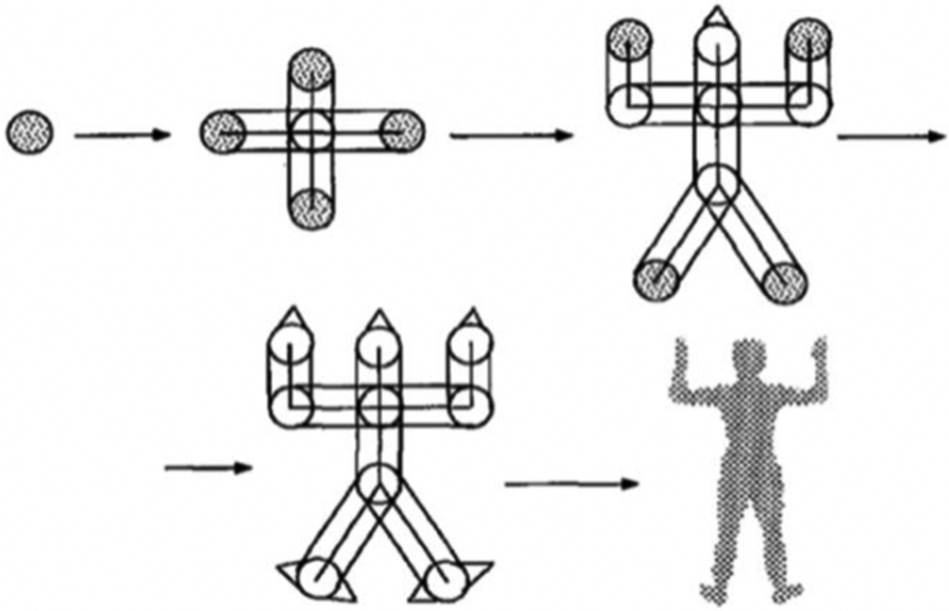

Figure 8. The composition of the shape of the human body based on worms and circles. From Zhu and Yuille (1996, p. 201).

kind of structures are called 'topological skeletons' (see Figure 9); they are invariant representations of the connectivity of the parts of an object and they seem to play an important role in human vision (Firestone & Scholl, 2014; Lowet et al., 2018).

Zhu and Yuille (1996, p. 203–204) present an algorithm that maps skeletons of different objects onto each other (for an example, see Figure 10). The algorithm generates an explicit measure of similarity for comparing different shapes, or for comparing a shape to a prototypical model. An interesting feature is that their similarity measure works even if the parts of two objects cannot be matched one by one. For example, a hand with a missing finger is still judged to have a high similarity with a typical hand and a dog with three legs is still seen as a dog.

These outlines of Marr and Nishihara's's (1978) and Zhu and Yuille's (1996) models of meronomic shape relations show that, in principle, structural shape relations can be represented in conceptual spaces, albeit high-dimensional and more complicated than for other perceptual properties.

The approach developed by Zhu and Yuille can be used by a search algorithm for part-whole analogies in the following way: Given an analogy $A : B : C : X$ such that B and X are parts of A and C respectively, the optimal choice for X will be the part of C that is mapped onto B while matching the topological skeletons of A and C. In other words, X will be the most similar part to B in C.

Note that if the order of terms in the analogy is $A : B : C : X$ such that A and C are parts of B and X respectively, the search procedure will be different. In this case, it will be necessary to specify a search space S of lexical categories such that $S = \{Y : Y$ satisfies the condition 'has. part.C'$\}$. Then, Zhu and Yuille's algorithm can be applied to find the most similar skeletal shape to B in S such that the part a in B is mapped onto the part C in X. Since it is very likely that there are various categories in S that are good candidates for completing $A : B : C : X$, the above algorithm can be used to find a set of skeletal shapes in S that meet or exceed a similarity threshold with regard to B.

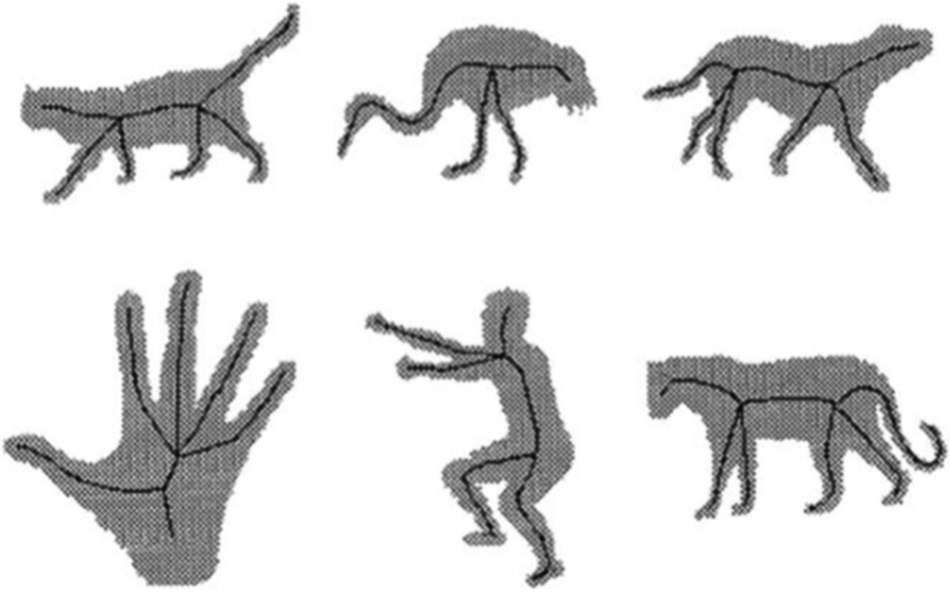

Figure 9. To the left, shape skeletons of typical objects. From Zhu and Yuille (1996, p. 194).

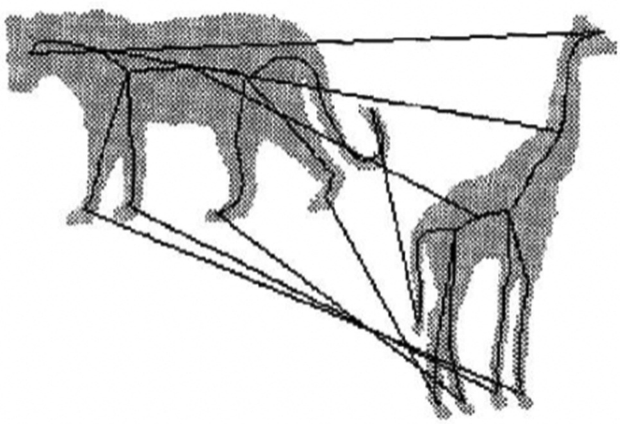

Figure 10. Mapping the skeleton of a leopard to that of a giraffe (from Zhu & Yuille, 1996, p. 207).

Conclusions

We have now analysed four types of analogies in terms of conceptual spaces: category-based, property-based, event-based, and part-whole analogies.[7] When we adopted the theory of conceptual spaces for explaining analogies, this classification was generated naturally. The four types correspond to different relational structures. (i) Category-based analogies, exemplified by *dog* : *wolf* : *cat* : *lynx*, take four concepts within the same domain and say that the vector between the first two elements of the analogy is parallel to the vector between the last two elements. Typically, all four concepts are expressed as nouns in this case. (ii) Property-based analogies of the type *pear* : *green* : *banana* : *yellow* take two categories of objects as first and third elements and two properties of the categories as second and fourth elements. The analogy states that the property (green) of the first category is *typical* for that domain of the category and that the property of the second category is typical for the same domain. In this case, the second and the fourth elements are typically expressed by adjectives. (iii) Event-based analogies, exemplified by *horse* : *gallop* : *man* : *run*, take two categories of objects as first and third elements and two actions of the categories as second and fourth elements. The analogy states that the action (gallop) of the first category involves a similar force pattern as the action of the second category (run). In this case, the second and the fourth elements are typically expressed by verbs. (iv) Finally, part-whole analogies of the form *hand* : *arm* : *foot* : *leg* exploit information stored in the structure domains of a conceptual space. They depend on a form of structural similarity that compares the parts in the analogues regarding their role in the hierarchy of parts of their respective wholes. We suggested a mapping algorithm that could work as a model of this mechanism, but many others could work as well (e.g., Macrini et al., 2011; Siddiqi et al., 1999). In this case, the second and the fourth elements are typically expressed by nouns that are higher in the meronomic hierarchy than the first and third elements, respectively.

An advantage of using conceptual spaces as a modelling tool is, as we already have noted, that the relative *strengths* of different analogies can be compared. This holds for category-based analogies, for event-based analogies and for some cases of part-whole analogies. In this way, our models yield predictions that other accounts of analogies cannot make. Future work should focus on evaluating the empirical adequacy of these predictions.

An important contribution of our approach is its detailed analysis of the role of semantic similarity in analogy. Propositional-based views, such as Thagard's (Thagard et al., 1990) or Gentner's (Gentner, 1983), also rely on semantic similarity but ignore the modulating role of dimensions and salience. We show that focusing on conceptual structure rather than propositional structure has clear advantages for explaining the diversity of analogies and for designing modelling algorithms.

These search algorithms are amenable for implementations in artificial systems. Since there already exist computational models of conceptual spaces (Adams & Raubal, 2009; Chella et al., 2001; Gärdenfors, 2014; Lieto, 2021) these algorithms could be extended by implementing the search procedures proposed here in order to account for different types of analogical reasoning. In addition, conceptual spaces can work as an interface between propositional models and subsymbolic models of cognition (Lieto et al., 2017). This opens up the possibility that our approach can be adapted in algorithm based on neural networks (see Jani & Levine, 2000).

Notes

1. Gentner (1983) defines relational predicates as n-ary predicates with $n > 1$.
2. Notice that $C(M)$ and $C(M*)$ are the same search space since they include the same set of subcategories.
3. Some empirical methods for determining dimensional salience in natural categories can be found in (Sloman et al., 1998) and (Rein et al., 2007).
4. 'More surprising' should not be given a probabilistic reading, but is the same as 'less expected' in terms of the expectation orderings presented in Gärdenfors and Makinson (1994) and Osta-Vélez and Gärdenfors (2022). As a matter of fact, the postulates for expectations orderings are incompatible with an interpretation in terms of probabilities.
5. It is shown in Osta-Vélez and Gärdenfors (2022) that the typicality criterion generates a total ordering of the properties in $C(M)$.
6. Alternatively, the skeleton can be defined as the sets of centres of discs that 'touch' the shape's boundary in at least two points.
7. A fifth type, which we do not study in this paper, are functional (causal) relations (e.g. *dog* : *leash* or *student* : *notebook*). For some ideas about how functions can be analysed with the aid of conceptual spaces, see Gärdenfors (2007).

Acknowledgments

Many thanks to Corina Strößner, Peter Brössel, and Matteo De Benedetto for their feedback on previous versions of this paper. We would also like to thank the anonymous reviewers for their helpful comments.

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

Adams, B., & Raubal, M. (2009). Conceptual space markup language (CSML): Towards the cognitive semantic web. In *2009 ieee international conference on semantic computing* (Berkeley, CA, USA) (pp. 253–260).

Barbot, N., Miclet, L., & Prade, H. (2019). Analogy between concepts. *Artificial Intelligence*, *275*, 487–539. https://doi.org/10.1016/j.artint.2019.06.008

Behrens, H. (2017). The role of analogy in language processing and acquisition. In M. Hundt, S. Mollin, & S. E. Pfenninger (Eds.), *The changing English language: Psycholinguistic perspectives* (pp. 215–239). Cambridge University Press.

Chella, A., Frixione, M., & Gaglio, S. (2001). Conceptual spaces for computer vision representations. *Artificial Intelligence Review*, *16*(2), 137–152. https://doi.org/10.1023/A:1011658027344

Chin-Parker, S., & Ross, B. H. (2004). Diagnosticity and prototypicality in category learning: A comparison of inference learning and classification learning. *Journal of Experimental Psychology Learning, Memory, and Cognition*, *30*(1), 216. https://doi.org/10.1037/0278-7393.30.1.216

Cummings, L. (2020). Arguments from analogy. In L. Cummings (Ed.), *Fallacies in medicine and health* (pp. 191–229). Springer.

Douven, I. (2019). Putting prototypes in place. *Cognition*, *193*, 104007. https://doi.org/10.1016/j.cognition.2019.104007

Fiorini, S.R., Gärdenfors, P., & Abel, M. (2014). Representing part–whole relations in conceptual spaces. *Cognitive Processing*, *15*(2), 127–142. https://doi.org/10.1007/s10339-013-0585-x

Firestone, C., & Scholl, B. J. (2014). "Please tap the shape, anywhere you like" shape skeletons in human vision revealed by an exceedingly simple measure. *Psychological Science*, *25*(2), 377–386. https://doi.org/10.1177/0956797613507584

Gärdenfors, P. (2000). *Conceptual spaces: The geometry of thought*. MIT press.

Gärdenfors, P. (2007). Representing actions and functional properties in conceptual spaces. In T. Ziemke & J. Zlatev (Eds.), *Body, language and mind* (pp. 167–195). Mouton de Gruyter.

Gärdenfors, P. (2014). *The geometry of meaning: Semantics based on conceptual spaces*. MIT Press.

Gärdenfors, P. (2021). Reasoning with expectations about causal relations. Unpublished manuscript..

Gärdenfors, P., Jost, J., & Warglien, M. (2018). From actions to effects: Three constraints on event mappings. *Frontiers in Psychology*, *9*, 1391. https://doi.org/10.3389/fpsyg.2018.01391

Gärdenfors, P., & Makinson, D. (1994). Nonmonotonic inference based on expectations. *Artificial Intelligence*, *65*(2), 197–245. https://doi.org/10.1016/0004-3702(94)90017-5

Gärdenfors, P., & Warglien, M. (2012). Using conceptual spaces to model actions and events. *Journal of Semantics*, *29*(4), 487–519. https://doi.org/10.1093/jos/ffs007

Gentner, D. (1983). Structure-Mapping: A theoretical framework for analogy. *Cognitive Science*, *7*(2), 155–170. https://doi.org/10.1207/s15516709cog0702_3

Gentner, D. (2002). Analogy in scientific discovery: The case of johannes kepler. In L. Magnani & N. Nersessian (Eds.), *Model-Based reasoning* (pp. 21–39). Springer.

Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 195–235). MIT Press.

Gentner, D., & Holyoak, K. J. (1997). Reasoning and learning by analogy: Introduction. *The American Psychologist*, *52*(1), 32. https://doi.org/10.1037/0003-066X.52.1.32

Gentner, D., & Hoyos, C. (2017). Analogy and abstraction. *Topics in Cognitive Science*, *9*(3), 672–693. https://doi.org/10.1111/tops.12278

Gentner, D., & Markman, A. (1997). Structure mapping in analogy and similarity. *The American Psychologist*, *52*(1), 45–56. https://doi.org/10.1037/0003-066X.52.1.45

Gentner, D., Rattermann, M. J., Markman, A., & Kotovsky, L. (1995). Two forces in the development of relational similarity. In T. J. Simon & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling* (pp. 263–313). Lawrence Erlbaum.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology*, *23*(2), 222–262. https://doi.org/10.1016/0010-0285(91)90010-L

Goodman, N. (1972). Seven strictures on similarity Goodman , N. . In *Problems and projects* (pp. 437–450). Bobbs-Merrill.

Goswami, U. (1992). *Analogical reasoning in children*. Psychology Press.
Goswami, U., & Brown, A. L. (1990). Melting chocolate and melting snowmen: Analogical reasoning and causal relations. *Cognition*, *35*(1), 69–95. https://doi.org/10.1016/0010-0277(90)90037-K
Henley, N. M. (1969). A psychological study of the semantics of animal terms. *Journal of Verbal Learning and Verbal Behavior*, *8*(2), 176–184. https://doi.org/10.1016/S0022-5371(69)80058-7
Holyoak, K. J. (2012). Analogy and relational reasoning Holyoak, K. J., Morrison , R. G. In *The oxford handbook of thinking and reasoning* (pp. 234–259). Oxford University Press.
Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, *13*(3), 295–355. https://doi.org/10.1207/s15516709cog1303_1
Jani, N. G., & Levine, D. S. (2000). A neural network theory of proportional analogy- making. *Neural Networks*, *13*(2), 149–183. https://doi.org/10.1016/S0893-6080(99)00106-9
Johannesson, M. (2002 Geometric Models of Similarity 90 Lund University Cognitive Studies). (Lund, Sweden.: Lund University) 171 (). , .
Lieto, A. (2021). *Cognitive design for artificial minds*. Routledge.
Lieto, A., Chella, A., & Frixione, M. (2017). Conceptual spaces for cognitive architectures: A lingua franca for different levels of representation. *Biologically Inspired Cognitive Architectures*, *19*, 1–9. https://doi.org/10.1016/j.bica.2016.10.005
Lowet, A. S., Firestone, C., & Scholl, B. J. (2018). Seeing structure: Shape skeletons modulate perceived similarity. *Attention, Perception, & Psychophysics*, *80*(5), 1278–1289. https://doi.org/10.3758/s13414-017-1457-8
Macrini, D., Dickinson, S., Fleet, D., & Siddiqi, K. (2011). Bone graphs: Medial shape parsing and abstraction. *Computer Vision and Image Understanding*, *115*(7), 1044–1061. https://doi.org/10.1016/j.cviu.2010.12.011
Markman, E. M. (1981). Two different principles of conceptual organization. In M. Lamb & A. Brown (Eds.), *Advances in developmental psychology* (pp. 199–222). Psychology Press.
Markman, A. B., & Gentner, D. (2000). Structure mapping in the comparison process. American Journal of Psychology, 3 (4), 501–538. https://doi.org/10.2307/1423470.
Marr, D., & Nishihara, H. K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London. Series B. Biological Sciences* London, England, *200*, 269–294.
Marr, D., & Vaina, L. (1982). Representation and recognition of the movements of shapes. *Proceedings of the Royal Society of London. Series B. Biological Sciences* London, England, *214*, 501–524.
Mitchell, M. (2021). Abstraction and analogy-making in artificial intelligence. *Annals of the New York Academy of Sciences*, *1505*(1), 79–101. https://doi.org/10.1111/nyas.14619
Nosofsky, R. M. (1992). Similarity scaling and cognitive process models. *Annual Review of Psychology*, *43*(1), 25–53. https://doi.org/10.1146/annurev.ps.43.020192.000325
Oppenheimer, R. (1956). Analogy in science. *The American Psychologist*, *11*(3), 127–135. https://doi.org/10.1037/h0046760
Ortony, A. (1979). Beyond literal similarity. *Psychological Review*, *86*(3), 161. https://doi.org/10.1037/0033-295X.86.3.161
Osta-Vélez, M., & Gärdenfors, P. (2020). Category-Based induction in conceptual spaces. *Journal of Mathematical Psychology*, *96*(102357). https://doi.org/10.1016/j.jmp.2020.102357
Osta-Vélez, M., & Gärdenfors, P. (2022). Nonmonotonic reasoning, expectations orderings, and conceptual spaces. *Journal of Logic, Language and Information*, *31*(1), 77–97. https://doi.org/10.1007/s10849-021-09347-6
Papafragou, A. (2015). The representation of events in language and cognition. In E. Margolis (Ed.), *The conceptual mind: New directions in the study of concepts* (pp. 327–346). MIT Press.
Piaget, J., Montangero, J., & Billeter, J. (1977). La Formation des correlate. In J. Piaget (Ed.), *L'abstraction reflechissante* (pp. 115–129). Presses Universitaires de France.
Rein, J. R., Love, B. C., & Markman, A. B. (2007). Feature relations and feature salience in natural categories. In *Proceedings of the annual meeting of the cognitive science society* Amsterdam, Netherlands (Vol. 29).
Rosch, E. (1975). Cognitive reference points. *Cognitive Psychology*, *7*(4), 532–547. https://doi.org/10.1016/0010-0285(75)90021-3
Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Lawrence Erlbaum.
Rosch, E. (1983). Prototype classification and logical classification: The two systems. In E. Scholnick (Ed.), *New trends in conceptual representation: Challenges to piaget's theory* (pp. 73–86). Lawrence Erlbaum Associates.
Rumelhart, D. E., & Abrahamson, A. A. (1973). A model for analogical reasoning. *Cognitive Psychology*, *5*(1), 1–28. https://doi.org/10.1016/0010-0285(73)90023-6
Shepard, R. N. (1964). Attention and the metric structure of the stimulus space. *Journal of Mathematical Psychology*, *1*(1), 54–87. https://doi.org/10.1016/0022-2496(64)90017-3
Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, *237*(4820), 1317–1323. https://doi.org/10.1126/science.3629243
Siddiqi, K., Shokoufandeh, A., Dickinson, S. J., & Zucker, S. W. (1999). Shock graphs and shape matching. *International Journal of Computer Vision*, *35*(1), 13–32. https://doi.org/10.1023/A:1008102926703

Sloman, S. A., Love, B. C., & Ahn, W.-K. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, *22*(2), 189–228. https://doi.org/10.1207/s15516709cog2202_2

Smith, L. B. (1989). From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 146–178). Cambridge University Press.

Thagard, P., Holyoak, K. J., Nelson, G., & Gochfeld, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence*, *46*(3), 259–310. https://doi.org/10.1016/0004-3702(90)90018-U

Tomlinson, M., & Love, B. C. (2006). Learning abstract relations through analogy to concrete exemplars. In *Proceedings of the 28th annual conference of the cognitive science society* Vancouver, Canada (pp. 2269–2274).

Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*(4), 327. https://doi.org/10.1037/0033-295X.84.4.327

Vosniadou, S. (1995). Analogical reasoning in cognitive development. *Metaphor and Symbol*, *10*(4), 297–308. https://doi.org/10.1207/s15327868ms1004_4

Warglien, M., Gärdenfors, P., & Westera, M. (2012). Event structure, conceptual spaces and the semantics of verbs. *Theoretical Linguistics*, *38*(3–4), 159–193. https://doi.org/10.1515/tl-2012-0010

Winston, M. E., Chaffin, R., & Herrmann, D. (1987). A taxonomy of part-whole relations. *Cognitive Science*, *11*(4), 417–444. https://doi.org/10.1207/s15516709cog1104_2

Zhu, S. C., & Yuille, A. L. (1996). Forms: A flexible object recognition and modelling system. *International Journal of Computer Vision*, *20*(3), 187–212. https://doi.org/10.1007/BF00208719