

Similarity and Contrast on Conceptual Spaces for Pertinent Description Generation

Giovanni Sileno^{1,2(✉)}, Isabelle Bloch¹, Jamal Atif², and Jean-Louis Dessalles¹

¹ LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France
`giovanni.sileno@telecom-paristech.fr`

² Université Paris-Dauphine, PSL Research University,
CNRS, UMR 7243, LAMSADE, 75016 Paris, France

Abstract. Within the general objective of conceiving a cognitive architecture for image interpretation able to generate outputs relevant to several target user profiles, the paper elaborates on a set of operations that should be provided by a cognitive space to guarantee the generation of relevant descriptions. First, it attempts to define a working definition of contrast operation. Then, revisiting well-known results in cognitive studies, it sketches a definition of similarity based on contrast, distinguished from the metric defined on the conceptual space.

Keywords: Similarity · Contrast · Conceptual spaces · Relevance · Description generation · Triangle inequality · Asymmetry · Diagnosticity effect

1 Introduction

Similarity, for its fundamental role in human reasoning, occupies a central role in cognitive science. In general, it is modeled as a function of a context-dependent distance (e.g. [1]). On the other hand, most machine learning methods rely on adequate metrics to perform comparisons between inputs; certain methods even attempt to achieve (pseudo-)metric learning. This convergence towards geometric models is not without problems. First, there exist many empirical studies, starting from the famous work of Tversky [2], that show that similarity in human judgment does not satisfy fundamental metric axioms. Secondly, a good part of *reasoning* operations performed by artificial devices still relies on symbolic means (e.g. ontologies expressed in description logics), that do not have a direct geometric interpretation. Attempts to fill these gaps exist, for instance by enriching the metric model of similarity with additional elements to align it with the empirical results (e.g. [3]), or by approaching the logical structures used in artificial reasoning via geometrical notions (e.g. [4]).

This work attempts to follow an alternative path. On the one hand, we continue along the tradition of works on psychological spaces, and precisely, we build upon the theory of *conceptual spaces* [5,6], acknowledging a third cognitive level, between symbolic and associationistic, where conceptual entities are

geometric representations. In our general research project, we aim to introduce conceptual spaces as a sort of *middle-ware* for image interpretation applications. Here, assuming that conceptual spaces exist, we focus on operations required for generating *relevant descriptions*, i.e. *pertinent characterizations* of a given input. We sketch a technical solution satisfying properties of predication observed in direct experiences. As a result of this proposal, we argue that the mismatches between geometric and empirical properties of similarity might not be a consequence of the space *per se* but of neglecting part of the mechanisms behind description generation.

The paper unveils its arguments incrementally. We first give a brief overview of the theory of conceptual spaces [5,6] and of a recent variation [7], focusing on pertinent predication. We then sketch the infrastructure required for generating relevant descriptions starting from simple examples of predication, but which are already problematic for semantic approaches relying on set theory. From this base, we present an alternative definition of similarity, that predicts results observed in empirical studies (asymmetry of similarity judgments, non satisfaction of the triangle inequality, diagnosticity effect). A note on further developments ends the paper.

2 Conceptual Spaces and Predication

2.1 Overview on Conceptual Spaces

According to Gärdenfors' theory of *conceptual spaces* [5,6], the meaning of words can be faithfully represented as convex regions in a high-dimensional geometric space, in which dimensions correspond to cognitively primitive features. Technically, conceptual spaces are usually modeled as vector spaces (e.g. [8–10]). An object of the conceptual space is characterized by several *qualities* or attributes:

$$(q_1, q_2, \dots, q_n), \forall i : q_i \in Q_i$$

where Q_i are sets of possible values for each quality q_i . Quality dimensions correspond to the ways in which two *stimuli* can be considered to be similar or different, usually according to an ordering relation of the *stimulus*. In general, the Q_i are modeled as *concrete domains* on $\mathbb{R}, \mathbb{R}^2, \dots, \mathbb{N}, \mathbb{N}^2$, etc. but proposals exist to process nominal domains (e.g. [11,12]). So far, we used the term domain in the mathematical sense. However, in works on conceptual spaces—in agreement with the cognitive psychology literature—the term *domain* identifies a set of *integral dimensions*, i.e. dimensions that cannot be separated perceptually (e.g. for humans, the color dimensions *hue-luminosity-saturation*). A conceptual space consists therefore of:

$$C = D_1 \times D_2 \times \dots \times D_m$$

where each D_i is a domain. As each D_i consists of a set of qualities, the resulting structure is hierarchical. According to its proponents, this infrastructure enables

the distinction between *objects*, i.e. points of the space (used to represent *exemplars* and *prototypes*, i.e. *exemplar-based* and *prototypical* bodies of knowledge [11]), and *concepts*, defined as regions of the space.

To guarantee betweenness among *similar* elements, *natural properties* correspond to *convex regions* in a domain [5]. A concept is a combination of properties, typically across multiple domains (linguistically, properties are usually expressed as adjective-like attributes, while concepts as nouns or verbs). Prototypes emerge as centroids of those convex regions (properties or concepts); at the same time, the division of conceptual spaces in regions can be seen as the result of a competition between prototypes, that might be captured by *Voronoi tessellations*, useful for *categorization* applications; technically, existing implementations exploit e.g. *region connection calculus* (RCC) [8], or polytope structures [13].

Evidently, there is a strong affinity between the representation based on features used in machine learning and the idea of conceptual spaces. For instance, *word embeddings* techniques also represent the meaning of words as points on a high-dimensional Euclidean space; however, conceptual spaces offer two advantages: first, working with regions and not only with similarity, they provide an intuitive way to process subsumption, overlap and typicality; second, dimensions of conceptual spaces have (or should have) a direct relation to a domain, while word embeddings dimensions are essentially meaningless.¹

2.2 Predication and Relevance

The theory of conceptual spaces assumes a generally working association between regions and linguistic marks. For this insistence on *lexical meaning*, the proposal can be seen as an extension of the symbolic approach. A recent alternative proposal [7] considers instead that predicates are the result of *contrast* operations made *on the fly* between conceptual objects, following principles of relevance. In essence, contrast is a “difference” operator (denoted with $-$) between the vectors pointing to the exemplar (O) and to the prototype (P):

$$C = O - P$$

Being the outcome C a vector, with (theoretically) the same dimensionality of the conceptual space, it may be a conceptual object as well. However, this vector should not be interpreted extensionally, but rather as a conceptual *force* or *modifier*; we will name these objects *contrastors*.

The proposal carries interesting innovations. First, whereas practically all other works rely on a global distance, it does not necessarily require a *holistic* perspective on all available dimensions. Second, it does not have to refer to average lexical meanings emerging from usage, but it is computed contingently with C and P grounded on the agent’s own experience.² Third, working with

¹ Jameel and Schockaert [14] show that a NLP architecture based on conceptual spaces yields better results than word-embeddings and knowledge-graph embedding.

² The negotiation of the specific symbols associated to C and P (*anchoring*), here assumed as given, remains deferred at social level.

contrast would allow to overlook the convexity constraint, requiring in principle only access to the representational level of points. For these reasons, we take it as a starting point for our investigation. In the following, we start working out an implementation of contrast, missing in the original paper [7].

3 Experiments for the Specification of Contrast

3.1 First Example: “red dog”

In predicate logic, the expression “red dog” is usually written as an x such that $Red(x) \wedge Dog(x)$, that, in the set-theoretic semantics, refers to an entity included both in the set of dogs and in the set of red entities. However, a red dog is not red as would be a red face, nor as would be a red book; in the usual labeling of colors it looks actually rather brown. Being red—semantically—might depend on the type of object on which the predicate applies. Accepting this, we suppose that the description of an object is constructed in at least two steps: first, an association to the nearest prototype (categorization), and then the extraction of the characteristic features by contrast. In this work, we bypass the prototype association, assuming it as given (machine learning algorithms have been proven successful in this respect). We focus only on the contrastive component of predication. In our example, “this dog” exemplar contrasted with the “dog” prototype should return a “red” contrastor.

For simplicity, we consider dogs as defined merely by their colors. We have taken the RGB colors of 9 common furs of dogs from the Internet, and converted them in HLS dimensions (*hue, luminosity, saturation*), in accordance with the conceptual space literature. Other color spaces, such as CIELAB or CIELUV can even better match visual perception, and first experiments confirm that the proposed approach applies in these spaces as well. For the sake of the argument, we continue with HLS. The statistical properties of the set are:³

mean: [0.10 0.52 0.46], std dev: [0.02 0.22 0.27]

Simplifying, we could take the mean as a prototype of color of a dog. This computation is transparent or at least less sensible to frequency effects; we are not averaging on actual populations of dogs, but on breeds. Figure 1 shows the selected points and the centroid, including points from the HLS spectrum with a plausible label association, e.g. “red” to $(0, 0.5, 1)$.

Let us consider a dog exemplar that would go under the “red sable” label. As a first step, we see contrast as *vectorial difference* of exemplar and prototype:⁴

[0.07 0.24 0.92] - [0.10 0.52 0.46] = [-0.16 -0.28 0.45]

³ Hue is an angular dimension, so the calculation of mean and standard deviation follows *circular* (also known as *directional*) *statistics* methods.

⁴ Given two angles a and b , we have computed $a - b$ as the angle of the vectorial difference of the two normalized vectors corresponding to the input angles, which is equivalent to the circular mean of a and $b + \pi$.




Fig. 1. Colors of dog furs and standard colors (with labels), on HLS dimensions. (Color figure online)

As the contrastor aims to capture the distinctive characteristics of the subject entity (e.g. this dog), with respect to a reference entity (e.g. the dog prototype), it lies, as a vector, in the same space as the two conceptual objects. The operation in itself however gives no evident clue on how to compare the outcome with, for instance, a previously acquired contrastor.

Let us consider now a red book. Assuming that practically all colors are possible for a book cover, each quality dimension can be modeled under a uniform distribution. When normalized, their standard deviation is $1/\sqrt{12} \approx 0.29$.⁵ We may then define an *empirical principle of relevant dimension*: the more the standard deviation along a normalized quality dimension approaches 0.29 the less we expect that quality to be relevant to form the prototype. In the extreme case, we should not take it to define the centroid. Applying roughly the principle on the dog case, the only pertinent dimension for the prototype is hue.

Now, the standard mathematical tool for vectorial spaces requires that points have a value for all coordinates. However, conceptual spaces have potentially infinite dimensions, and points may have undefined dimensions. To enable the possibility to operate with any point, we introduce a *void value* * when a certain dimension is not applicable.⁶

We want to identify a method to generate, from this representation, that our dog is a red dog, and that our book is a red book. For the book, being the prototypical book color void, the color characterization mirrors the color specification given by the sensory module, i.e. there is no contextualizing effect

⁵ $\mu = 1/2$, $\text{Var} = E[(X - \mu)^2] = \int_0^1 (x - 1/2)^2 dx = 1/12$.

⁶ In a similar spirit, Aisbett and Gibbon [15] introduce the idea of *distinguishing point*.

due to the prototype.⁷ In formula, we have that:

$$(.., a, ..) - (.., *, ..) = (.., a, ..)$$

The color spectrum serves as a source of contrastors. In the case of dogs, we expect instead a contextualizing effect. Assuming that contrastors have prototypes as well, we require a method to compute to which category the contrastor we computed falls upon, in order to enable a reuse of this category in different contexts (e.g. red dog, red book, red face). In clustering algorithms (see e.g. [16]) the comparison between two numeric points is done using *distance measures* (usually Euclidean, Manhattan or Chebychev), or functions such as *cosine similarity*, *Pearson correlation measure* etc. In *typicality-based clustering* (e.g. [17]), it relies on a *typicality degree* computed using a *internal resemblance* (with other members of the cluster) and an *external dissimilarity* (with members of other clusters).




Fig. 2. Hue, luminosity, saturation (HLS) color spectrum. (Color figure online)

For the aim of this paper, we do not need to decide a clustering algorithm (i.e. a prototype formation mechanism), nor to settle upon how a contrastor should be associated to its prototype (i.e. a categorization mechanism).⁸ In the following, we will denote the category/prototype association with the symbol ‘ \rightsquigarrow ’. For instance, in our red dog example, we have: $(0.07, 0.24, 0.92) - (0.10, *, *) = (-0.16, 0.24, 0.92) \rightsquigarrow red$. In effect, looking at the HLS spectrum (Fig. 2), the contrastor calculated for our dog exemplar is still in the gravitation of red but on the opposite side of brown and yellow.⁹

3.2 Second Example: “a above b”

Imagine we have two objects, *a* and *b*, one above the other. In predicate logic their relationship might be written as *above(a, b)*, that, from a logical point of

⁷ This contextualization can be interpreted as informational *compression*. When a prototype cannot be formed for a quality (because e.g. exemplars exhibit a uniform distribution with respect to that dimension), the sensory input cannot be compressed.

⁸ As exemplars and contrastors are vectors of the same space, it is plausible to hypothesize that they share similar prototyping/categorization mechanisms.

⁹ The fact that the contrastor is capturing magenta can be explained by an incomplete parametrization of additional semantic aspects (concerning e.g. the actual conceptual space on which contrast is applied).

view, would be the inverse of $below(b, a)$. From a natural language point of view, however, considering a an apple and b a table, we observe it is much more natural to say “the apple is on the table” rather than “the table is below the apple”. We hypothesize therefore that contrast is at stake, selecting the relation most pertinent to the situation.

Objects are extended: they occupy a certain space, that may be described as a solid shape, with a center and a rotation angle, or, when captured in images, by pixels. Intuitively, directional relations should be computed by this spatial information. Applying the principle of contrast here, we should have that a contrasted with b should return an “above” *contrastor*.

Simplifying, let us reduce objects to points. Considering their positions specified in e.g. a 2D space, we have that $(a_x, a_y) - (b_x, b_y)$ is actually a seen from b , or, equivalently, as if the origin of the space has been moved to b . Defining the *above* contrastor prototype as a vector $(0, u)$, with u positive number, and similarly *below* $(0, -u)$, *right* $(u, 0)$, *left* $(-u, 0)$, it seems we might utilize the same principle of the previous example. For instance, $(2, 4) - (1.5, 2) = (-0.5, 2) \rightsquigarrow above$ (and a bit of *left*), with $u = 2$. This interpretation brings two questions to the foreground. First, objects are never points: if they are represented as such, it is because they have been discretized at a certain granularity. Second, vectors have an intrinsic metric unit, which, mathematically, is inherited by the contrastor (when constructed through vector difference), but we expect e.g. *above* to be the same relation when applied to macro or to micro-objects. Leaving the study of such normalization mechanism to future research, we focus on the first problem.




Fig. 3. Computing directional relationships using morphological dilation [18].

We consider an existing method [18] used in image processing to compute directional relative positions of visual entities (e.g. of biomedical images). The method exploits *mathematical morphology* operators—namely *fuzzy dilation*—to create an adequate *fuzzy landscape* (a fuzzy region) from a model of the directional relationship, and the position and shape of a reference object. The *strength* of the relation between the target and the reference is quantified via a

normalized degree of intersection of the target object with the fuzzy landscape region (please refer to [18] for technical details). We have reported in Fig. 3 an example similar to the “the apple is on the table” case. The strength of b below a results to be inferior to a above b , basically because the “below a ” region, seen through the “ b ” mask, contains gray pixels, while the opposite is not true.

Interpreting the previous operations at a higher level, we can draw interesting observations. The method enables the comparison of two descriptions, but it does not primarily make any reference to a contrast between two entities. The strength of the directional relationship is computed via its *realization* in the image space. The structuring element stands as a model of the region *above* of a *point element* located at the origin; it corresponds to the reification of all possible answers to the question “Is this point above the origin?” The dilation operation contextualizes these answers using the reference b ; it is as if the binary relation $above(a, b)$ is reduced to a unary form: $above_b(a)$. The subject entity a is then taken into account in the computation of the degree of intersection: how much a ’s extension *falls within* the virtual entity $above_b$. If we consider multiple virtual entities as *right_of_b* etc. we are recreating a problem similar to the categorization/prototype association, but where the categories have been created on the fly and concern the image space around the reference. Rather than performing a direct *contrast* between two objects ($a - b \rightsquigarrow above$), here it seems we are exploiting the dual operation *merge* ($a \rightsquigarrow b + above$).¹⁰ In effect, this mapping works also in the previous example, i.e. we are able to map $this_dog - prototype_dog \rightsquigarrow red$ to $this_dog \rightsquigarrow prototype_dog + red$.

This intuition sheds light on how \rightsquigarrow may operate. The right operand specifies the final result (the category or the nearest prototype to which the contrastor is associated); but to achieve it, several possible candidates, distributed over the conceptual space, have to be adequately compared to the left operand. If the left and right operands can be processed as points, cluster association methods will do. If they are regions, we will need to assess the *overlap* over the candidate regions, using analytical (as the intersection degree) or also random (e.g. Monte Carlo) methods. These candidate regions might need to be realized (i.e. contextualized with a reference) when they are not directly available, which is an operation computationally expensive. In future work, we will investigate how to remap the merge operations to the left part of the contrast equation.

4 Similarity

In this section we build upon the previous analysis to define *similarity*. We will start from presenting two reference models of similarity judgment presented in cognitive science, widely used in many applications (particularly Tversky’s model); we will then introduce our model of similarity as double contrast, and evaluate it with empirical results reviewed in the literature.

¹⁰ This interpretation is still compatible with the contrast formula, as, in effect, we are implicitly assuming that a has not a prototypical position.

4.1 Similarity as Feature Matching

A rich tradition of psychological and cognitive studies on similarity starts from working with sets of *features* associated with objects.¹¹ In a well-known paper, Tversky [2] argued that similarity cannot be modeled as a distance, because many empirical experiences shows that similarity judgments does not satisfy three geometric axioms:

- *minimality*: $d(a, b) \geq d(a, a) = 0$; respondents identify another object similar to the object more often than an object to itself;
- *symmetry*: $d(a, b) = d(b, a)$; for instance, “Tel Aviv is like New York” is not the same as “New York is like Tel Aviv”;
- *triangle inequality*: $d(a, b) + d(b, c) \geq d(a, c)$; for instance, Russia and Cuba are (were) similar as political systems, Cuba and Jamaica are similar geographically, but Russia and Jamaica do not share anything.

Besides, Tversky [2] observes another relevant phenomenon:

- *diagnosticity effect*: the result of similarity judgment changes when the list of possible alternative changes. For instance, participants are asked for the country most similar to Austria to be decided amongst Hungary, Poland (at those times both communist countries), and Sweden; most responses indicate Sweden. If Norway is added to the list, however, responders turn to Hungary.

Given certain assumptions, Tversky proves that similarity can be expressed with what he calls the *contrast model* (A is a set of features of a , B of b):

$$S(a, b) = \theta f(A \cap B) - \alpha f(A \setminus B) - \beta f(B \setminus A)$$

where S is the *similarity scale*, f a non-negative scale, and θ , α and β positive parameters. If $\alpha > \beta$, the model creates the asymmetry between subject a and reference b , explaining the observed lack of symmetry in similarity judgments. This account is also compatible with an observed *imperfect* complementarity of similarity and dissimilarity (*difference*, in Tversky’s terms).

4.2 Similarity with Density Effects

Soon after Tversky’s proposal, Krumhansl [3] presented an alternative model to explain the same phenomena, attempting to recover the geometrical hypothesis. She starts from the problem of minimality axiom, observing that in experiments problems increase when the subject point has many neighbors, i.e. when it is

¹¹ It is worth observing that such collections are the result of a preliminary extraction and compilation by some modeler. Approaches based on conceptual spaces, although certain implementations may be in practice very similar to feature-based works (e.g. by considering nominal dimensions), in principle insist on direct, perceptual grounds of quality dimensions.

more prototypical. She then considers similarity as a function both of an *inter-point distance*, and of a *spatial density* of stimuli points:

$$d'(a, b) = d(a, b) + \alpha\delta(a) + \beta\delta(b)$$

where δ is a density function. The similarity function is then suggested to be a composition of this distance with a monotonically decreasing function. Like Tversky, she suggests that the difference between similarity and dissimilarity in empirical tests may be due to different factors α, β or to difference in density between subject and referent points.

4.3 Similarity as Double Contrast

So far, we have not defined similarity in our account. For calculating contrast, we have used distances inherent to the integral dimensions. These distances may in effect be interpreted as quantities of how much two stimuli are *dissimilar*, but, because these two stimuli should belong to the same domain (e.g. color, image space), they cannot refer to (multi-dimensional) concepts. Similarity, instead, operates at level of concepts. This is evident in metaphors, i.e. expressions like “my love is as deep as the ocean”, “he is like a lion”. For this reason, we start from sketching a general template for metaphor generation.

For simplicity, let us consider a sentence like “he is strong”. Following the red dog example, this predicate should result from the operation:

$$this_person - prototype_person \rightsquigarrow strong, \dots$$

Saying instead “he is like a lion”, we are performing a double operation: we are matching one or more characteristic properties of that person with one or more characteristic properties of the concept of lion.

$$\begin{aligned} this_person - prototype_person &\rightsquigarrow X, \dots \\ prototype_lion - prototype_animal &\rightsquigarrow X, \dots \end{aligned}$$

However, whereas the contrastor category X is recognized as the same in the contrasts, the two contrastors have plausibly different intensities. The *asymmetry* between the subject and the referent can be seen as a consequence of the relative intensities of the contrastors (cf. *above* vs *below*). We can generalize this idea:

Proposition 1 (Comparison Ground). *A comparison ground between two conceptual objects holds when the contrasts with their prototypes results into two contrastors falling upon the same category, but possibly with different intensities. The natural reference is the object whose contrastor has greater intensity.*

The suspension dots in these equations are meant to show that in principle many contrastors could be generated, related to different domains. For instance, lion’s distinguishing features are being strong, living in the savannah, etc. Why we match the strength dimension rather than the position? A *descriptive criterion* may be because in common-sense there are few animals as strong as lions, but

there are plausibly many other animals that live in the savannah. This criterion is related to the empirical principle of relevance.

Tversky’s **asymmetry** examples, e.g. “Tel Aviv is like New York”, can be interpreted following the metaphor’s template. This sentence means that one or more distinctive characteristic of New York are used as modifiers for defining Tel Aviv. To decide whether it is nightlife, cosmopolitanism or green areas we may use the descriptive criterion to compute the distinctiveness of certain features. Saying the opposite (“New York is like Tel Aviv”) would mean to activate Tel Aviv’s more distinctive characteristics.

For the **triangle inequality**, suppose the concept of country is defined by political and geographical dimensions. A sentence like “Cuba is similar to Russia”, without defining the dimensions of the comparison, follows a scheme similar to the metaphor case, but in this case, the two objects belong to the same category. The double contrast works in the Cuba–Russia case (for the matching on the political dimension) and in the Jamaica–Cuba case (for the matching on the geographical dimension), but it does not work in the Jamaica–Russia case, as it is not able to find a common comparison ground.

To explain the **diagnosticity effect**, we hypothesize that respondents construct a rough estimate of the group. Rather than a real prototype, we are dealing here with a temporary construct, bringing to the foreground the *relative* distinctive features of the given individuals. For the sake of the argument, let us consider a simplified model of Tversky’s test (*political index* 1 means communist country; *position* approximates the position on a map of the centers of the countries). Averaging these values, we can construct a sort of virtual prototype of the group, with and without Norway:

	political_index	position (x, y)
Austria	0	1, 0
Hungary	1	2, 0
Poland	1	2, 1
Sweden	0	1, 4
Norway	0	0, 4
Group prototype without Norway	0.5	1.5, 1.25
Group prototype with Norway	0.4	1.2, 1.8

Following the double contrast formula, in order to decide which country in the group is most similar to Austria, we have first to contrast the Austria object with the virtual prototype, and then apply the same operation to the other countries to form a compatible *comparison ground*. In the group without Norway, we need to specify Austria’s political dimension, as the group is split perfectly in two; for the spatial dimension, Austria is more central in the group, so geography is a bit less pertinent dimension to describe the country. Sweden is selected as the most similar country to Austria because the comparison ground lies on the political index. In the group with Norway, it is more common not to be a communist country (so the political dimension becomes a bit less pertinent), whereas the center goes further North: the geographical dimension becomes more pertinent. For this internal change, Hungary is selected as more similar to Austria.

The search for a satisfactory comparison ground can also explain the experiences in which the **minimality axiom** is not satisfied. The original tests were about recognizing, given a Morse code, the most similar code in a list of codes, including the input one [3]. According to our model, when the input is far from the group prototype it is more difficult to find a reference with which to form a comparison ground: the respondent will correctly identify the same element as the best response to the task. When the input is near to the prototype, another entity may be *satisfactorily* similar to stop the search.

Comparison. Although the models presented by Tversky [2] and Krumhansl [3] might yield as well predictions aligned with these experiences, they require the specification of adequate parameters and, more importantly, the manual selection of features (potentially infinite). Consider two objects that are almost the same, save for a detail that make them crucially different in pertinence to a task: Tversky’s contrast model would lead to implausible results, as the weight of common features would outnumber that of distinct features. On the other hand, proposals as that of Krumhansl’s have the problem of defining a coherent global distance amongst features. In principle, our proposal is not concerned by these problems: we have not used any parameters; if the conceptual spaces have been correctly constructed, they are grounded to perceptive spaces and respecting the conceptual hierarchies; contrast and similarity are computed without relying on global distances.

5 Conclusion and Further Developments

The paper presents a novel account of contrast and similarity operations to be performed on conceptual spaces. Contrast relies on distances computed along integral dimensions (belonging to independent cognitive domains), capturing dissimilarities between entities on given scales of stimuli. Similarity judgments are modeled instead as double contrasts forming a *comparison ground*, where the distinctive characteristics of a reference are used as contrastors. To our knowledge, such sequential, multi-layered nature of similarity has not been hypothesized before in the literature. The dimensions of psychological spaces, even in the theory of conceptual spaces, are usually elicited via *multi-dimensional scaling* (MDS) techniques applied on people’s similarity judgments, presupposing the existence of underlying metrics to be captured by features expressed in a verbal form. These approaches and similar dimensional reduction techniques used in machine learning conflate two aspects of similarity, respectively of perceptual and contrastively analogical nature, that our proposal attempts to distinguish.

Future work is required to complete our specification of contrast: additional semantic parameters, the analytical relationship with merge, and a definition operating on regions.¹² Prototypes were processed here merely as points; how-

¹² As the merge operation (+) seems to be captured by *dilation*, we are currently investigating methods to capture contrast (−) considering the dual morphological operator *erosion*: first experiments look promising.

ever, the empirical principle of relevance presented here relies on standard deviation, and for that, it works with a neighborhood of the mean, representative of an underlying group or population. The principle was applied in a binary fashion: either the prototype has no effect on the exemplar, or the prototype counts at par with the exemplar. The next step would be to specify a graded solution, taking into account some regional information of prototypes, captured by standard deviation or other means. We expect the problem has points of contact with the contrast on regions. Parallel to this work, we need to study the interaction of possible definitions of contrast with dependent qualities, to take into account semantically redundant information.

Finally, this work focused on the descriptive aspect of predicate generation, but relevance is not only a matter of best description. At higher level, additional factors play a role, as e.g. the rarity or the emotional response associated to the situation to be described. Understanding how this interplay works will be crucial for generating truly pertinent descriptions.

References

1. Nosofsky, R.: Similarity scaling and cognitive process models. *Annu. Rev. Psychol.* **43**(1), 25–53 (1992)
2. Tversky, A.: Features of similarity. *Am. Psychol.* (1977)
3. Krumhansl, C.L.: Concerning the applicability of geometric models to similarity data: The interrelationship between similarity and spatial density. *Psychol. Rev.* **85**(5), 445–463 (1978)
4. Distel, F., Bloch, I., Atif, J.: Concept dissimilarity based on tree edit distance and morphological dilation. In: 21st European Conference on Artificial Intelligence (ECAI) (2014)
5. Gärdenfors, P.: *Conceptual Spaces: The Geometry of Thought*. Institute of Technology, Massachusetts (2000)
6. Gärdenfors, P.: *The Geometry of Meaning: Semantics Based on Conceptual Spaces*. MIT Press, Cambridge (2014)
7. Dessalles, J.-L.: From conceptual spaces to predicates. In: Zenker, F., Gärdenfors, P. (eds.) *Applications of Conceptual Spaces*. SL, vol. 359, pp. 17–31. Springer, Cham (2015). doi:[10.1007/978-3-319-15021-5_2](https://doi.org/10.1007/978-3-319-15021-5_2)
8. Gärdenfors, P., Williams, M.A.: Reasoning about categories in conceptual spaces. In: *IJCAI International Joint Conference on Artificial Intelligence*, pp. 385–392 (2001)
9. Raubal, M.: Formalizing conceptual spaces. *Formal ontol. Inf. Syst.* **114**, 153–164 (2004)
10. Adams, B., Raubal, M.: A metric conceptual space algebra. In: Hornsby, K.S., Claramunt, C., Denis, M., Ligozat, G. (eds.) *COSIT 2009*. LNCS, vol. 5756, pp. 51–68. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-03832-7_4](https://doi.org/10.1007/978-3-642-03832-7_4)
11. Lieto, A., Radicioni, D.P., Rho, V.: A common-sense conceptual categorization system integrating heterogeneous proxytypes and the dual process of reasoning. In: *IJCAI International Joint Conference on Artificial Intelligence*, pp. 875–881 (2015)
12. Lewis, M., Lawry, J.: Hierarchical conceptual spaces for concept combination. *Artif. Intell.* **237**, 204–227 (2016)

13. Adams, B., Raubal, M.: Conceptual Space Markup Language (CSML): Towards the cognitive Semantic Web. In: 2009 IEEE International Conference on Semantic Computing ICSC, pp. 253–260 (2009)
14. Jameel, S., Schockaert, S.: Entity embeddings with conceptual subspaces as a basis for plausible reasoning. In: European Conference on Artificial Intelligence, pp. 1353–1361 (2016)
15. Aisbett, J., Gibbon, G.: A general formulation of conceptual spaces as a meso level representation. *Artif. Intell.* **133**(1–2), 189–232 (2001)
16. Rokach, L., Maimon, O.: Clustering methods. In: *The Data Mining and Knowledge Discovery Handbook*, p. 32 (2010)
17. Lesot, M.J.: Typicality-based clustering. *Int. J. Inf. Technol. Intell. Comput.* **1**(2), 279–292 (2005)
18. Bloch, I.: Spatial reasoning under imprecision using fuzzy set theory, formal logics and mathematical morphology. *Int. J. Approx. Reason.* **41**(2), 77–95 (2006)