Computing Contrast on Conceptual Spaces

Giovanni Sileno[1,2], Isabelle Bloch[1], Jamal Atif[2], and Jean-Louis Dessalles[1]

[1] LTCI, Télécom ParisTech, Université Paris-Saclay, Paris, France
[2] Université Paris-Dauphine, PSL Research University, CNRS, UMR 7243, LAMSADE, 75016 Paris, France
giovanni.sileno@telecom-paristech.fr

Abstract. This paper provides an updated formalization of the operation of contrast, and shows that, by applying it on conceptual spaces, membership functions to categories as e.g. those captured by adjectives or directional relationships emerge as a natural by-product. Because the outcome of contrast depends not only on the objects contrasted (a target and a reference, as for instance a prototype), but also on the frame in which those are contained, it is argued that contrast enables a continuous contextualization, offering a basis for "on the fly" predication. This investigation is used for inferring general requirements for the application of contrast and its generalization, and for comparison with current practices in the conceptual space literature.

Keywords: Contrast, Conceptual Spaces, Categorization, Predication, Contextualization

1 Introduction

In its original formulation, the theory of conceptual spaces [1] assumes a generally working association between regions and linguistic marks; because symbols are associated with sets of points, the theory can be seen as relying on an *extensional semantics*, in continuity with symbolic approaches. An alternative proposal, introduced in [2], starts from the observation that *predication* (i.e. the production of predicates about a certain object or situation) should follow principles of relevance, which, from a descriptive point of view, means determining an object by utilizing its distinctive features. Accepting this, predicates should be the result of *contrast* operations made *on the fly* between conceptual objects. In a previous paper [3], starting from this basis, and representing conceptual objects as points in a conceptual space, we were able to explain, all in maintaining a geometric view of psychological spaces, the misalignments between theoretical properties and empirical observations of *similarity judgments* [4], concerning *symmetry*, *triangle inequality*, but also the overlooked *minimality axiom* and *diagnosticity effect*. Motivated by this result, part of our research has been directed towards a more structured formalization of contrast, whose functional components were only sketched before. This paper aims to present preliminary results of this effort and considerations about future developments.

The paper proceeds as follows. After a brief overview on the theory of conceptual spaces, the contrast mechanism is explained on a simple case issued from a mono-dimensional perceptual space (§ 2). These results are then generalized to the multidimensional case, analyzing illustrative examples (§ 3). The paper ends with a discussion, contextualizing our approach within the current practices in the conceptual space literature, highlighting some open questions (§ 4).

1.1 Overview of the Theory of Conceptual Spaces theory

In the literature [1, 5], the introduction of conceptual space is motivated by the observation that the meaning of words can be represented as *convex regions* in a *high-dimensional geometric space*, whose dimensions correspond to cognitively primitive features. In practice, conceptual spaces are usually modeled as vector spaces (e.g. [6–8]). An object of a 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, depending on an ordering relation between them; they are usually modeled on $\mathbb{R}$, $\mathbb{R}^2$, ..., $\mathbb{N}, \mathbb{N}^2$, etc. but proposals exist to process nominal dimensions (e.g. in [9, 10]). In agreement with the cognitive psychology literature, dimensions are organized in *domains* or sets 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:

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

where each D_i is a domain. As each D_i consists of a set of qualities, the resulting structure is hierarchical. This representational 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 [9]), and *concepts*, defined as regions of the space.

According to the original theory, *natural properties* emerge as *convex regions* within a domain [1] to guarantee betweenness among *similar* elements. Concepts generalize properties as *weighted* combinations of them, typically across multiple domains, and should satisfy the convexity constraint.[3] Prototypes can be seen as *centroids* of concept (including property) regions. The other way around, the division of conceptual spaces into regions can be seen as the result of a competition between prototypes, that might be captured by e.g. *Voronoï tessellations*; this approach can be used for *categorization*. Following empirical evidence, the theory suggests to measure the distance between points in the conceptual spaces (i.e. distance between conceptual objects, exemplars or prototypes) through a weighted Manhattan metric of the intra-domain distances.

[3] This is source of discussion: see for instance [11] about the consequent impossibility of capturing correlation geometrically, and the various arguments provided in [12].

1.2 Restarting from Predication

For its insistence on *lexical meaning*, approached by the association of linguistic marks to regions, the original account of conceptual spaces can be seen as an extension of the symbolic approach, in the sense that it follows an extensional semantics. In contrast, the alternative proposal introduced in [2] considers that predication should follow principles of relevance, which, from a descriptive point of view, means describing an object by utilizing its distinctive features. Consequently, predicates are hypothesized to be the result of *contrast* operations made *on the fly* between conceptual objects.

This change of view carries interesting innovations. First, whereas practically all other works rely on a global distance over all available dimensions (which potentially are infinite), the use of contrast does not necessarily require a *holistic* perspective, but restrains the focus to (adequately) distinctive dimensions. Second, working with contrast would allow for the most to bypass the problem of maintaining definite regions, so as to require in principle only access to the representational level of points and some rough regional information. Third, it does not refer to average lexical meanings emerging from usage, but it is computed contingently with contrastors and prototypes grounded on the agent's own experience (the association of specific symbols or *anchoring* [13, 14] remains deferred at social level). The present contribution aims to develop further this approach.

2 Contrast

Consider coffees served in a bar. Intuitively, whether one of these coffees is predicated as being "hot" or "cold" depends mostly on what the locutor expects of coffees served at bars (ignoring effects due to her contingent state), rather than a specific absolute temperature. In other words, the description of an object or exemplar o (e.g. a coffee served in a bar) results from its *contrast* with a certain prototype (of coffees served in bars) p. Naming the resulting output *contrastor*, we can utilize its categorization as a basis for predication, at least for such modifiers. Then, for instance, denoting provisionally the operation of contrast with $-$, and categorization with $\rightsquigarrow$, we would have:

$$o - p = c \rightsquigarrow \text{“hot”}$$

Let us assume that conceptual spaces are vector spaces defined on $\mathbb{R}^n$; as for the moment we are focusing on the temperature dimension, we have $n = 1$. If both left and right operands were simple points, and contrast corresponded to vectorial difference, c would be a *free vector*, maintaining the same scale of the input points. However, at least this last aspect is implausible, because "modifying" categorizations resulting from contrast might be applied to very different scales (cf. "small molecule" vs "small galaxy"). We need at least an adequate scaling.

2.1 Scaling

Because prototypes are associated with a certain concept region, we might consider some regional information as well. Following models dealing with imperfect spatial information such as the *egg-yolk* model (e.g. [15]), we refer to two boundaries: an internal one (containing the *yolk*), computed as $p \pm \sigma$, and external one (containing the *egg*), computed as $p \pm \rho$, where p, the prototype point, is the centroid of the region (for simplicity, egg and yolk regions are here supposed to be symmetric). Elements falling within the yolk are *typical*, *normal* instances of the concept associated with the prototype; within the egg but not in the yolk are instances still associated with the same concept, but manifesting some distinctive characteristic; when elements are outside the egg, they are not *directly* associated with the concept.[4]

Partitioning Let us consider a perceptive space made just by one dimension (e.g. temperature), modeled as $U = [-\rho_U, \rho_U]$; this is a bounded space (plausible assumption in the context of finite cognitive resources), defined by two boundaries of opposite *polarity*. The neutral element, as well as center of the space, is indexed by 0. We can denote U also as $\odot_0^{\rho_U}$, i.e. a region centered on 0 of radius ρ_U. This space may be naturally divided into partitions such as $\mathcal{M}_2 = \{[-\rho_U, 0], [0, \rho_U]\} = \{U^-, U^+\}$, or $\mathcal{M}_3 = \{[-\rho_U, -\frac{\rho_U}{3}], [-\frac{\rho_U}{3}, \frac{\rho_U}{3}][\frac{\rho_U}{3}, \rho_U]\}$, etc. Note that these constructions are made independently of any world semantics: they are just based on the perceptual possibilities given by the space.

Centering Now, returning to our example of contrast, we aim to compare our target with *typical* exemplars of the concept, which, by construction, are contained in *yolk* region $p \pm \sigma$.[5] One way to do that is to perform a point-wise contrast. Let us denote the extensional descriptions of o and p respectively by $O = \{o\}$ and $P = \{x : |x - p| \leq \sigma\} = \odot_p^\sigma$. Performing a point-wise vectorial difference of the sets[6] (here denoted with $\ominus$) we have:

$$O \ominus P = \{x - y : x \in O, y \in P\} = \{o - y : y \in P\} = \{-(y - o) : y \in P\} = (P \ominus O)^s$$

where $(X)^s$ is the symmetric of X with respect to the origin. Observing that the inner operation is a translation, we have: $O \ominus P = (\odot_p^\sigma \ominus \{o\})^s = \odot_{o-p}^\sigma$.

[4] In principle, the distinction of σ and ρ boundaries should reproduce the difference in judgments about *typicality* and *categorization* observed in empirical settings [16].

[5] In other words, σ captures a minimal distance that brings an object outside the typicality zone centered on p, making it *distinguishable*. It can be used to express a sort of *indeterminacy* of the prototype, which will be propagated to the output contrastor.

[6] Note that, in the general case, a point-wise difference of two sets of points is a multi-set. Here, as O consists of a single point, it is a simple set.

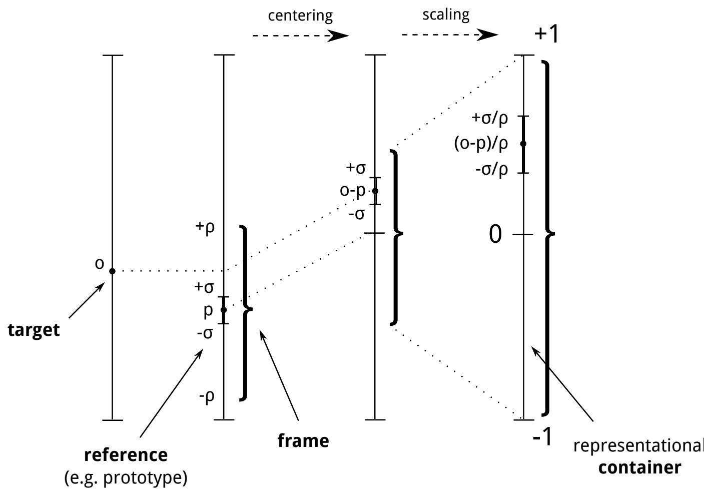

Fig. 1. The contrast operation concerns four elements: *target* and *reference* objects, which are contained in a *frame*, and a final representational *container*. It consists of two steps: *centering* (computing the difference of the target with respect to the reference), and *scaling* (standardizing the frame to the representational container).

Scaling In order to exploit the full representational domain given by the perceptual space, contextualized for objects of the same category of the target, we might consider the *egg* region $p \pm \rho$, presumably containing all the exemplars of the related concept. For doing that, given $U = [-\rho_U, \rho_U] = \odot_0^{\rho_U}$, and denoting *scaling* with "$*$" ($\odot_x^\alpha * \gamma = \odot_{\gamma x}^{\gamma\alpha}$), we need to use a scaling factor γ such that $\gamma \cdot \rho = \rho_U$. Let ρ_U be arbitrarily 1. Thus, in extensional terms, the contrastor C obtained by contrasting an object o for its prototype p is given by:

$$C = \mathsf{contrast}(o, \langle p, \sigma, \rho \rangle) = \odot_{o-p}^{\sigma} * \frac{1}{\rho} = \odot_{(o-p)/\rho}^{\sigma/\rho} \tag{1}$$

where $\langle p, \sigma, \rho \rangle$ specifies the region associated with the concept: the centroid, and the internal and external boundaries. Figure 1 illustrates its computation.

2.2 Categorization

At this point we need to settle upon a *categorization* method, i.e. how to select the category to which the contrastor belongs. When contrastors and categories are seen as extensional objects, a typical method would consist in finding the category with maximal overlap with the contrastor. Denoting the category label with r, and its extension with $M^{(r)}$, the method amounts to solve:

$$\arg\max_r |C \cap M^{(r)}| \tag{2}$$

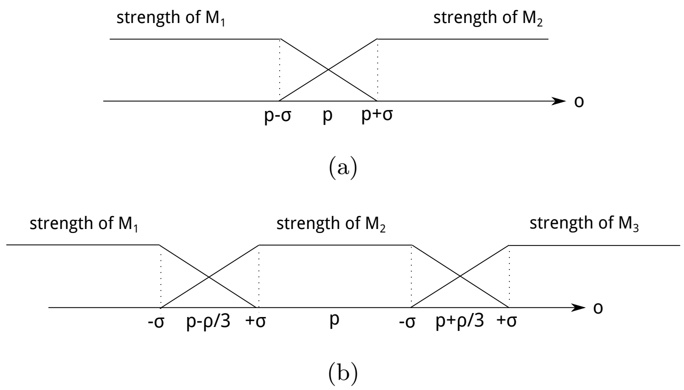

Fig. 2. Membership functions to categories captured (a) by $\mathcal{M}_2$, with $M_1 = [-1, 0]$, $M_2 = [0, 1]$; (b) by $\mathcal{M}_3$, with $M_1 = [-1, -1/3]$, $M_2 = [-1/3, 1/3]$, $M_3 = [1/3, 1]$.

The *strength* of a certain categorization can be computed as the following ratio:

$$\text{strength}(r) = \frac{|C \cap M^{(r)}|}{|C|} \tag{3}$$

Thus, the category with the maximal overlap with the contrastor (i.e. with the maximal strength), if associated with a label, may be used to describe a distinctive feature of the object (with respect to the prototype).

Natural categories Let us consider as an example the partitions issued from the perceptual space, starting from the *bipolar categorization* captured by $\mathcal{M}_2 = \{M_1, M_2\}$. Using the definition of contrastor with the strength equation given above, we obtain Figure 2a. A relevant point is $o = p$, the case in which our object is plainly prototypical, and then the contrastor is a region centered on zero, capturing the same surface on the two dual parts M_1 and M_2. If we interpret M_1 as associated with "cold" and M_2 to "hot", we could say that, according to this view, the coffee is as much cold as hot. However, this bipolar construction would fail to sufficiently capture the *graduality* of judgment, in the sense that a coffee may also be "ok", condition which is usually not defined as being hot and cold at the same time. Fortunately, to solve this we just need to add an additional category, as in Figure 2b with $\mathcal{M}_3$. Note that all partitions $\mathcal{M}_{2n+1}$ similarly built provide a neutral category. Evidently, we do not need necessarily to use such constructions, their use here is rather illustrative.

Choice of parameters In the previous sections the choice of the parameters was given qualitatively: σ captures the most typical exemplars, whereas ρ in principle covers all exemplars belonging to the same category of objects $\mathcal{O}$ (here defined on one dimension). A simple way to compute the radius ρ would be:

$$\rho = \frac{\max_{o \in \mathcal{O}} o - \min_{o \in \mathcal{O}} o}{2}$$

One way to decide σ could be by relying on the standard deviation, thus introducing information concerning the relative frequencies amongst individuals. Alternatively, it could be calculated just as ρ, but considering a core subset of $\mathcal{O}$ centered on p. Note that the definitions given so far should be modified if we relax the assumption of symmetrical regions.

Adaptation The region parameters could be modified in two ways:

- *unsupervised*: when an exemplar o, not within the current boundaries of the category of objects $\mathcal{O}$, results to be more similar to the prototype p of $\mathcal{O}$ than to prototypes of other categories, it is then labeled as belonging to p, implying a redefinition of ρ and σ;
- *supervised*: the user provides a new exemplar o explicitly labeled as $\mathcal{O}$, going beyond the current boundaries of the category; ρ and σ are then recomputed.

In both cases, the process implies an effect of *relativization*: providing more contrastive exemplars, those which were highly contrastive before become less contrastive. Conversely, if the number of maintained exemplars is bound, and pruning occurs on objects acquired more remotely in time–or, more plausibly, adaptation mostly concerns recent objects with an aggregated prototype–we can also observe a *hardening* effect: the region will recenter around the most recent elements.

General application The previous specification has been introduced to contrast an exemplar, represented as a point, with a prototype, represented with a point and two boundaries. Additional cases of application may be imagined, such as contrasting two exemplars, or two prototypes (two regions). In order to handle them, we can generalize the previous formulation, considering that four elements play a role in practice: the *target* (e.g. o) and the *reference* (e.g. $\odot_p^\sigma$) are in the foreground; the *frame* (e.g. $\odot_p^\rho$) provides background in which the first two are contextualized, and it is used to control the scaling with respect to the representational *container* (e.g. $\odot_0^1$). In principle foreground inputs may be points or regions; on the other hand, frame and container cannot be points. Note that the frame in general could not be centered on the reference region, but it is plausible to require it to contain both the target and the reference.

One way to take into account regions with the previous method is to *discretize* them, or better, to consider points as regions of minimal granularity, similarly to what happens with digital images. In effect, the assumption of limited cognitive resources implies not only the boundedness of the perceptual space, but also its *finite granularity*. In principle, we could apply a contrast as defined in (1) for each point of the target region and then aggregate the results (see § 3.1). Alternatively, to avoid to specify aggregation, we could rescale the smallest region at stake (normally the target) so that it behaves like a point. Suppose that the container $[-1, 1]$ contains $2N$ grains, and so has granularity $1/N$. Similarly, a certain frame with radius ρ, assuming it exploits all the representation, has minimal granularity ρ/N. Now suppose we have a region centered on o with radius τ; expressed in grains, the region is long ${}^{(2\tau N)}/_{\rho}$. To represent it as a single grain, all values should be divided by this value. Using this idea, we can reformulate contrast in terms of *target*, *reference* and *frame* regions as:

$$C = \mathsf{contrast}_{\mathsf{R}}(\langle t, \tau\rangle, \langle r, \sigma\rangle, \langle f, \rho\rangle) \approx \mathsf{contrast}\left(\lfloor \frac{t}{2\tau} \rceil, \left\langle \lfloor \frac{r}{2\tau} \rceil, \lfloor \frac{\sigma}{\tau} \rceil, \lfloor \frac{\rho}{\tau} \rceil \right\rangle\right) \quad (4)$$

where $\lfloor \cdot \rceil$ is an approximation to the nearest integer value. Note that f, the center of the frame, plays only an indirect role, for the condition that the frame region should contain the target and reference.

Requirements for contrast So far, we assumed that objects, prototypes, etc. are specified on (a subset of) $\mathbb{R}$, and that we can compute an algebraic difference of two points, necessary for the "centering" step. The function of *difference* here can be expressed as follows: it produces an object that, *added* (as inverse operation) to the reference point, reproduces (at least to a certain extent) the target. The "scaling" step can be seen instead as one of *decontextualization* from the magnitudes at stake for the type of objects given in inputs. These requirements could in principle be abstracted, in order to consider the contrast operator as a higher-level function over other types of representation:

- to provide a relative order between inputs, the perceptual dimension could be represented for instance as a *complete lattice*, i.e. a partially ordered set for which all subsets have a *supremum* and an *infimum*;
- the "difference" between two objects should return an object, that, used as parameter with an "addition" operator, enables an adequate reconstruction of the target object from the reference. The presence of these specific roles makes clear that we are not in front of the usual algebraic operations. Denoting them to highlight their asymmetric roles, we have:

  $$a -_{\triangleleft} b \equiv d \qquad \text{such that} \qquad b +_{\triangleright} d \approx a$$

  This should hold in particular when a, b, d are points.
- the "scaling" within the contrast operation should tune this parameter to be neutral from the frame within which the objects are placed.

The second requirement implies that most methods introduced in the literature relying on *distance* are unfortunately not sufficient for defining contrast, because we need an output that enables *reversibility*. To find back the target point from the reference we require, in addition to distance, a *sign* for a single linear dimension, or a *versor* (unitary vector) for a multidimensional Euclidean vector space. In the general lattice case, additional parameters might be needed.

3 Multi-dimensional contrast

The previous section focused on objects specified by a single linear dimension. For practical uses, the previous result should be extended to the more general case of n dimensions. Suppose these dimensions to be *perceptually independent*: the sensory input coming from one of them cannot be specified (not even partially) by another one or any composition of the others. Objects are then described by a tuple of values like $o = (o_1, \ldots, o_n)$, where o_i is the value perceived by sensory module i. Similarly, assume concepts to be formed as structures $\langle p, \sigma, \rho \rangle$, where each item is defined on multiple dimensions $p = (p_1, \ldots, p_n)$, $\sigma = (\sigma_1, \ldots, \sigma_n)$ and $\rho = (\rho_1, \ldots, \rho_n)$. We consider that a multidimensional contrast corresponds to the application of contrast on each dimension:

$$C = (C_1, \ldots, C_n) = (\mathsf{contrast}(o_1, \langle p_1, \sigma_1, \rho_1 \rangle), \ldots, \mathsf{contrast}(o_n, \langle p_n, \sigma_n, \rho_n \rangle))$$

This output furnishes a *contrastive description* of the target object; each local contrast could be associated with a category with a certain strength, and the object could be described—with respect to its prototype—by the labels of the modifying categories with the greatest strengths.

3.1 Example of perceptual contrast: directional relationships

As a first example of application of multidimensional contrast, we consider the description of *directional relationships* (e.g. "left-of") holding between visual objects. In this case, entities are perceptually defined on a single perceptual domain. Let us consider an image space $\mathcal{S} \subseteq \mathbb{N}^2$ and two binary objects A (target) and B (reference). If A is a single point, we could in principle apply the same formula of point-region contrast (1), where the frame might be e.g. a circular region or bounding box containing A and B. By drawing all vectorial differences between reference points and target point, we could visualize the contrastor as an image as well, representing in practice all translations that each point of B should perform to produce A.

When A consists of a set of points, we might consider discretization as in § 2.2, but in general this is not possible; considering for instance the case of A being a long, thin curve. In these cases, we can still apply the point-wise contrast iteratively; each of the points of A can be used to produce a contrastor image, and an *aggregation* of these images (e.g. a normalized sum) would produce a synthetic information about the (point-wise) modifications to be performed to obtain (each point of) A from (each point of) B. In practice, this corresponds to accumulating the point-wise differences during the iteration, to count them and to normalize their counts. More formally, the accumulation set is captured by:

$$\mathcal{H}(A, B)(z) = \{a \in A, b \in B \mid a - b = z\}$$

The normalized cardinality of this set corresponds to a histogram-like function, and, noting that $\mathcal{H}$ is the set resulting from the point-wise difference of A and B, we override the previous notation $A \ominus B$:

$$\mu_{A \ominus B}(z) = \frac{|\mathcal{H}(A, B)(z)|}{\max_w |\mathcal{H}(A, B)(w)|} \qquad (5)$$

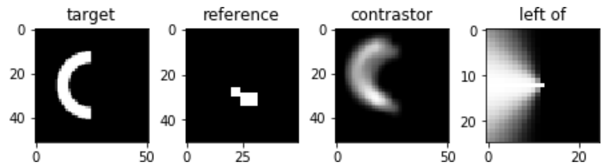

Fig. 3. Application of contrast on two binary images: the resulting contrastor can be compared with a "left of" region to evaluate the strength of the relationship.

With this reformulation $A \ominus B$ results to be a gradual representation in the general case. Contrast of regions can be redefined accordingly, for cases in which discretization cannot be used:

$$C = \mathsf{contrast}_\mathsf{R}(A, B, \langle f, \rho \rangle) = (A \ominus B) * \frac{1}{\rho} \tag{6}$$

Then, similarly to the monodimensional case, we may consider natural partitions on the perceptual space; for instance by dividing a 2D container in four regions, we could obtain rough models of "left of", "right of", "above of" and "below to". By comparing the contrastor object with these models we can decide which predicate corresponds to the directional relationship(s) holding between the objects. Figure 3 reports an example of this application: the target is contrasted by the reference, obtaining a contrastor visual object, whose overlap is maximal with the "left of" region, therefore the target is described as "left of" the reference (c.f. [17]).

3.2 Example of conceptual contrast: fruit concepts

At this point, we want to investigate the use of contrast on a scenario involving multiple perceptual domains. Let us take the simple case utilized in [18], in which some concepts about fruits are specified as regions on three perceptually independent dimensions (normalized to $[0, 1]$): *hue* (a quality dimension of the color domain), *roundness* (of shape) and *sweetness* (of taste). The resulting conceptual space consists of three mono-dimensional domains. The conceptual regions are defined as following:

concept	region			prototype (center)		
	hue	roundness	sweetness	hue	roundness	sweetness
pear	0.50–0.70	0.40–0.60	0.35–0.45	0.60	0.50	0.40
orange	0.80–0.90	0.90–1.00	0.60–0.70	0.80	0.95	0.65
lemon	0.70–0.80	0.45–0.55	0.00–0.10	0.75	0.50	0.05
granny smith	0.55–0.60	0.70–0.80	0.35–0.45	0.575	0.75	0.40
apple (green type)	0.50–0.80	0.65–0.80	0.35–0.50	0.65	0.725	0.425
apple (yellow type)	0.65–0.85	0.65–0.80	0.40–0.55	0.75	0.725	0.475
apple (red type)	0.70–1.00	0.65–0.80	0.45–0.60	0.85	0.725	0.525

Let us suppose we can obtain more abstract concepts by unifying these regions, as for instance the "apple" concept (including green, yellow and red types), or the fruit concept (or better, a partial version of it, as it includes only some fruits). Possible holes in the union are filled, implicitly assuming that it is possible for an object to be in that position. The prototype points for these new concepts could be computed in two ways: (a) as the center of the concept region (*domain-induced* prototype), or (b) as the average (i.e. weighted center or centroid) of the prototypes of the given sub-concepts (*group-induced* prototype).

concept	region			prototype (center/centroid)		
	hue	roundness	sweetness	hue	roundness	sweetness
apple	0.50–1.00	0.65–0.80	0.35–0.60	0.75/0.75	0.725/0.725	0.475/0.475
fruit	0.50–1.00	0.40–1.00	0.00–0.70	0.75/0.72	0.70/0.70	0.35/0.42

At this point, we apply contrast to identify the most pertinent (here in the sense of discriminatory) features. As the application of contrast requires a frame, we can choose a more abstract concept containing all inputs to be contrasted. We could consider for instance the "fruit" concept for all concepts, or the "apple" concept for the three types of apple. These regions provide both the reference (a prototype point, e.g. the center, in this case) and the frame (the concept region). The following table reports for instance the centers of the contrastors obtained by contrasting fruit concepts as "apple" with respect to the "fruit" concept (using discretization and taking $\sigma = 0.5\rho$):

concept	hue	roundness	sweetness	*red*	*green*	*blue*
pear	−0.6	−0.7	0.1	−0.3	1.0	0.0
orange	0.4	0.8	0.4	1.0	0.0	−1.0
lemon	0.0	−0.7	−0.4	0.8	0.8	−1.0
apple	0.0	0.1	0.2	0.0	0.0	0.0

The "roundness" and "sweetness" columns are easy to be interpreted. According to the given conceptual space, the orange is the sweetest fruit, the lemon the least. Pear and lemon are the least round, while orange is the most. Apples occupy almost neutral categories in all dimensions. The interpretation of "hue" is more complicated, because, by computing the difference as an algebraic difference (i.e. interpreting angles in terms of rotation), we require to utilize the reference to identify a specific color. In [3] we have considered an angular interpretation that in principle would solve this issue, but we have also acknowledged a problem with the scaling phase that, unfortunately, does not have a solution yet. Alternatively, we might consider to re-transform the hue dimension into the *red-green-blue* dimensions, which form a cartesian space, obtaining the three additional columns reported in the table above.[7] From these we are able to infer that the pears are characterized by being the greenest fruits *with respect to the color spectrum of the given fruits*; oranges by being the reddest (and least blue), lemons the most yellow (and least blue), whereas the concept of "apple" is not

[7] To obtain color values that were plausible from those given as inputs, we computed $h = [(1 - hue) \cdot 360° - 30°]/360°$ and then converted the *hls* tuple $\langle h, 0.5, 1\rangle$ to *rgb*.

distinctive with respect to color, as there are red, green and yellow apples (its distinctiveness would increase if e.g. a truly blue fruit was included in the "fruit" concept.)

Intuitively, starting from a similar table, possibly including some regional information, we could define some *weights* reifying the *salience* of these qualities for the formation of a concept, by settling an adequate measure of their *discriminatory power*.

Finally, we observe that the results of contrast may be used also for associating the object to a category. For instance, *within the frame of fruits*, one can easily compute that "granny smith" is nearest to the concept of apple, or even more to the green type of apple. Intuitively, the perceptual independence hypothesis supports the use of the Manhattan distance. However, an additional filtering may be needed when many dimensions are present to consider only adequately distinctive dimensions.

3.3 Examples of conceptual merge: "red brick" vs "brick red"

To complete the operational cycle, we present a small example concerning the *interpretation* of simple linguistic expressions. Despite the richer spectrum illustrated in Gärdenfors' books [1, 5], most works on conceptual spaces focus on the *intersective* type of concept composition, i.e. relying on *conjunction* of concepts, as with symbolic approaches; e.g. a "red brick" refers to an object that belongs to the class of "red" objects and to the class of "brick" objects.

An approach to predication based on contrast, instead, naturally implies an asymmetry between the roles of the modifier concept ("red") and the modified or reference concept ("brick") in the formation of the composed concept ("red brick"), in alignment with the *modifier-head* phenomenon observed in cognitive psychology [16]. We call *merge*, here provisionally denoted with +, the operation inverse to contrast, and revisit some of the examples brought by Gärdenfors. Let us consider the two merges associated to "red brick" and "brick red". Suppose "red" is a label anchored to a concept defined on the color dimension, and "brick" anchors to a multidimensional concept including a color dimension. As in [3], we consider a *void value* • when a certain dimension is not applicable (an object does not have a certain quality, or a prototype cannot be formed on that dimension). Two intuitive properties of *merge* can be identified (note how merge is not commutative):

- *category consistency*: dimensions which are not in the reference should not be added (for instance, we want a "brick red" to be a color):

$$(.., \bullet, ..) + (.., a, ..) = (.., \bullet, ..)$$

- *locality*: merge applies modifier concepts locally, i.e. not modifying dimensions irrelevant to the modifier (e.g. a "red brick" is merely a brick more red than the prototype brick):

$$(.., a, ..) + (.., \bullet, ..) = (.., a, ..)$$

The general application of these properties makes however clear that additional cognitive mechanisms are at stake. For the first property, consider predicates as "stone lion", or "stuffed gorilla". These are examples of merges that "break" the reference category. As there is no lion animal which is made of stone, a plausible repair would be "a stone object similar to a lion". Then, because the similarity between a stone object and a lion can be only constructed along the shape dimension (the only available *comparison ground* [3] between "stone object" and "lion"), we conceptualize a lion-like statue. For the second property, implicit correlations within dimensions do provide additional information; e.g. the "young" in "young man" does not seem to simply modify the age, leaving intact the other values: it seems that the merge on one dimension causes a recalibration of the prototype with this new information. Alternatively, this could be explained by assuming that a (sub)-concept "young man" is already available, and its prototype is activated by the merge (cf. the notion of *lexical compounds* [16]). Consider again the "red brick" case. For a layman, a red brick is simply a brick more red than the others; for an expert person, a red brick might be also e.g. more isolating than the average. An operationalization of this mechanism could explain part of the process of alignment of *linguistic semantics* with *conceptual semantics*.

4 Discussion

The present paper reports on ongoing research stemming from an alternative view on conceptual spaces, rooted on *relevant predication* [2]. This account insists on the importance of discriminatory aspects not only for individuation, but also for the formation of concepts. Recent, additional support for this hypothesis comes from cognitive studies in image recognition [19] proposing (predictive) recognition models based on internal discriminatory features such as the spatial organization of visual elements (cf. § 3.1). Not less importantly, as shown by the preliminary results presented here, by using contrast, the membership of an exemplar to a certain category—and therefore the possible consequent predication—is effectively *contextualized* on the fly, because the application of contrast requires always the intervention of background elements (conceptual frame and representational container), in addition to foreground elements (target and reference objects). By changing of pragmatic context, certain conceptual frames might become more accessible than others, and this would determine changes of interpretation for the same linguistic marks.

Additionally, the present proposal offers a computational model naturally implementing *modifier-head* concept combination [16], which, in our model, can be seen as including the *intersective* case (cf. "red book" vs "red dog" [3]). Gärdenfors [1, 5] presents an informal solution by referring to *contrast classes*, summarized by the *concept combination rule* [1]: "The combination CD of two concepts C and D is determined by letting the regions for the domains of C, confined to the contrast class defined by D, replace the values of the corresponding regions for D." This means that the domains of the modifier (e.g. temperature for "hot") are scaled to compatible domains of the reference (e.g. temperature for "coffee"), determining a new region identifying the compound concept ("hot coffee"). This is very similar to the operations suggested here, but there are differences too. First, instead of maintaining definite regions, we give priority to points and to rough regional information, in order to capture a sort of "intensity" of the modification. Second, by building upon contrast, *conceptualization* and *conceptual composition* become two sides of the same function, rather than a separate cognitive machinery. In short, the introduction of the operation of contrast might provide alternative foundations to the theory of conceptual spaces, and could be used to make explicit some of the internal mechanisms taken as given and to rely as less as possible to external parameters.

Evidently, many points still remain to be investigated. For instance, we aim to derive the salience of quality dimensions with respect to the formation of a certain concept from the conceptual structure, rather than being captured by externally given *weights*, and precisely by capturing the strength of their *contrastive* or *discriminatory power*. Thus, in addition to a refinement of the formalization and more detailed specifications of the methods (e.g. for directional dimensions), we aim to better understand the intertwining between perceptual independence and statistical independence. Perceptual independence is an important assumption for the distribution of contrast along independent dimensions. Because we cannot express something about a certain dimension in any other way than saying something about that dimension, a measure to quantify an aggregate contrast seems to be intuitively captured by a Manhattan distance. However, in the case of artificial devices, the internal configuration of sensors is not a consequence of evolutionary adaptation and may be the most various. For instance, cloned temperature sensors may receive the same information while being perceptually independent. Discovering a strong (or perfect) correlation has important consequences, in a way opposite to the introduction of new dimensions as e.g. children learn to do by separating volume from height of a container.

Furthermore, we would like to investigate in detail the *convexity* constraint, or more plausibly, the *star-shaped* constraint on concept regions; we expect it should naturally emerge from a categorization based on contrast. At a more fundamental level, while working on this paper, we acknowledged that by settling the functioning of contrast, conceptual spaces can be seen as *emerging* from contrastive functions. Stated differently, by defining contrast, we need also to define the inverse operation merge, the application of *merge* produces *order* relations between concepts, and the resulting lattice is in practice a conceptual space. Future investigations will target the formalization of this idea.

Finally, an additional track that we are currently studying concerns the use of the morphological operators *erosion* and *dilation* [20] for implementing approximate methods for contrast and merge, motivated by the observation that operators embedding non-linear functions as max and min carry computational advantages, and are cognitively more plausible than aggregations by average as those used in § 3.1.

References

1. Gärdenfors, P.: Conceptual Spaces: The Geometry of Thought. MIT Press (2000)
2. Dessalles, J.L.: From Conceptual Spaces to Predicates. Applications of conceptual spaces: The case for geometric knowledge representation, 359 Synthese Library, pp. 17–31 (2015)
3. Sileno, G., Bloch, I., Atif, J., Dessalles, J.L.: Similarity and Contrast on Conceptual Spaces for Pertinent Description Generation, Proceedings of KI 2017: Advances in Artificial Intelligence, vol. 10505 LNAI, pp. 262-275 (2017)
4. Tversky, A.: Features of similarity. Psychological Review, 84(4), 327-352. (1977).
5. Gärdenfors, P.: The Geometry of Meaning: Semantics Based on Conceptual Spaces. MIT Press (2014)
6. Gärdenfors, P., Williams, M.A.: Reasoning about categories in conceptual spaces. Proceedings of International Joint Conference on Artificial Intelligence (IJCAI) pp. 385–392 (2001)
7. Raubal, M.: Formalizing Conceptual Spaces. Formal Ontology in Information Systems, Proceedings of the Third International Conference (FOIS) 114, pp. 153–164 (2004).
8. Adams, B., Raubal, M.: A Metric Conceptual Space Algebra. Proceedings 9th International Conference on Spatial Information Theory (COSIT). pp. 51–68.
9. Lieto, A., Radicioni, D.P., Rho, V.: A common-sense conceptual categorization system integrating heterogeneous proxytypes and the dual process of reasoning. Proceedings of International Joint Conference on Artificial Intelligence (IJCAI) pp. 875–881 (2015)
10. Lewis, M., Lawry, J.: Hierarchical conceptual spaces for concept combination. Artificial Intelligence 237, pp. 204–227 (2016).
11. Bechberger, L., Kühnberger, K.U.: A Thorough Formalization of Conceptual Spaces. Proceedings of KI 2017: Advances in Artificial Intelligence, LNAI 10505, pp. 58–71 (2017)
12. Hernández-Conde, J.V.: A case against convexity in conceptual spaces. Synthese, pp. 1–27 (2016).
13. Coradeschi, S., Saffiotti, A.: Anchoring symbols to vision data by fuzzy logic. Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty. LNCS 1638, pp. 104–115
14. Coradeschi, S., Saffiotti, A.: An introduction to the anchoring problem. Robotics and Autonomous Systems 43(2-3), 85–96 (2003).
15. Cohn, A.G., Gotts, N.M.: The 'Egg-Yolk' Representation of Regions with Indeterminate Boundaries. Proceedings of the GISDATA Specialist Meeting on Geographical Objects with Undetermined Boundaries 2, pp. 171–187 (1996)
16. Hampton, J.: Conceptual combination. In: Knowledge Concepts and Categories, pp. 133–159. (1997).
17. Bloch, I.: Fuzzy Relative Position Between Objects in Image Processing: A Morphological Approach. IEEE Transactions on Pattern Analysis and Machine Intelligence 21(7), 657–664 (1999)
18. Bechberger, L., Kühnberger, K.U.: Measuring relations between concepts in conceptual spaces. Proceedings of SGAI 2017: Artificial Intelligence XXXIV, LNAI 10630, pp. 87–100 (2017)
19. Ben-Yosef, G., Assif, L., Ullman, S.: Full interpretation of minimal images. Cognition 171, pp. 65–84 (2018).
20. Bloch, I., Heijmans, H., Ronse, C.: Mathematical morphology. Handbook of Spatial Logics pp. 857–944 (2007).