Understanding Predication in Conceptual Spaces

Claudio MASOLO and Daniele PORELLO
Laboratory for Applied Ontology, ISTC-CNR, Trento, Italy

Abstract. We argue that a cognitive semantics has to take into account the possibly partial information that a cognitive agent has of the world. After discussing Gärdenfors's view of objects in conceptual spaces, we offer a number of viable treatments of partiality of information and we formalize them by means of alternative predicative logics. Our analysis shows that understanding the nature of simple predicative sentences is crucial for a cognitive semantics.

Keywords. Conceptual spaces, cognitive semantics, predication, objects, concepts, rough sets, three-valued logics.

1. Introduction

Conceptual spaces have been proposed by Gärdenfors [1] as a general framework for modelling representations of concepts of cognitive agents. One of the aims of Gärdefors's research program is to build a cognitive semantics of natural language that has to be capable of modelling the relationship between the language and the corresponding mental representations. In philosophical logic, analogous motivations drove the development of a number of non-classical logics, insofar as they are intended to approximate the reasoning capability of a knowing subject. For this purposes, the epistemological dimension—what is known by a subject or the information that is accessible to a subject—plays a fundamental role in the choice of the logical formalism that models reasoning.

A cognitively situated subject has, in general, only partial information about the world. In [2], we approached partial information for the case of propositional reasoning. The present work is a first step towards developing a cognitive semantics—that is grounded on the framework of conceptual spaces—for a fragment of a *first-order* language. The fragment that we discuss includes only individual constants, unary predicates, and the connectives $\neg$, $\wedge$, $\vee$. As we shall see, although the fragment is quite simple, its interpretation in terms of conceptual spaces raises a number of interesting ontological, epistemological, and logical problems. In particular, once the partiality of information is assumed, the very nature of predication—the object a falls under the concept C—requires careful examination. Gärdenfors views objects as points and concepts as regions in a multidimensional space. In this way, predication is simply understood as set-theoretic membership. We argue that Gärdenfors's view of objects as points in conceptual spaces is problematic from both an ontological and cognitive perspective. Moreover, when one considers epistemological underdetermination of objects—i.e., when not all the properties of an object are known at the maximal resolution offered by the conceptual spaces—

it is simply unfeasible to demand that objects are always understandable as fully determined points, i.e. vectors of values. By taking an epistemological perspective, it becomes then relevant to examine also 'underdetermination' of concepts. As objects shall not necessarily be interpreted as points of a conceptual space, concepts shall not necessarily correspond to crystal clear regions, they can be interpreted as *rough* (as opposed to sharp) regions of conceptual spaces. As we shall see, this double epistemological uncertainty challenges our understanding of the notion of predication itself.

The paper is organized as follows. In Section 2, we briefly introduce the framework of conceptual spaces. Section 3 criticizes the reduction of objects to points in conceptual spaces and introduces the notions of *underdetermined object* and *rough concept*. Our formal framework is described in Section 4, where we introduce and analyze a number of logics for modelling different combinations of underdetermination of objects and roughness of concepts. Section 5 concludes the paper.

2. Conceptual spaces

Gärdenfors [1] proposes a cognitive model of representations based on *conceptual spaces*. A conceptual space is a collection of *domains* that, in their turn, are decomposed into (*quality*) *dimensions*. Dimensions—e.g., temperature, weight, pitch, brightness—correspond to "the different ways stimuli are judged to be similar or different" [1, p.6] and they are built on the basis of these judgments of similarity that "reveal the dimensions of our perceptions and their structures" [1, p.5]. A set S of dimensions is *integral* if an object located in one dimension is necessarily located also in all the other dimensions in S, e.g., $\{hue, brightness\}$. A set of dimensions is *separable* if it is not integral, e.g., $\{hue, pitch\}$. Domains are maximal sets of integral dimensions—e.g., the color domain $\{hue, chromaticness, brightness\}$—and they usually have a *metrics* (or a *geometrical* structure). The way in which the distance defined on a domain is computed is actually an empirical question that is usually approached by means of *multidimensional scaling* (see [3] for an introduction), a methodology to represent similarity judgments between objects in terms of the distance between their locations, i.e., points in a multi-dimensional domain. However, usually, the distance defined on a domain is reduced to the distances defined on its quality dimensions. For instance, the distance between colors is a function of the distance between the values of hue, chromaticness, and brightness.

A *natural property* is a convex region (a convex set of points) of a domain. The separability condition assures that individuals can be classified within a domain *independently* of their classification within other domains. For instance, on can ascribe the weight and the color to an object independently. Conceptual spaces are defined as collections of one or more domains and *concepts* are represented as regions in conceptual spaces. In particular, a *natural concept* is represented as a set of convex regions in a number of domains "together with as assignment of salience weights to the domains and information about how the regions in different domains are correlated" [1, p.105]. Note that concepts are *static* theoretical entities "in the sense that they only describe the *structure* of representations" [1, p.31].

3. Objects and Conceptual Spaces

In [1,4], Gärdenfors assumes that an *object* is represented by a point in a conceptual space, i.e., a vector of coordinates, one for each dimension in the space. As acknowledged by Aisbett and Gibson [5], this choice is problematic for identifying the domains that are "irrelevant or inappropriate to a particular object or concept" [5, p.205]: by assuming a unique n-dimensional conceptual space all the objects must have a location in all the n dimensions. Gärdenfors himself recognizes that *abstract* entities, as opposed to *physical* ones, have no location in space and time, "so their underlying domains are different from those of physical objects" [1, p.135]. To formally account for irrelevant domains, Aisbett and Gibson extend the domains with the distinguished point denoted by '$*$', i.e., "the point at infinity", for which the distance is $d(a,*) = \infty$ for all $a \neq *$ in the domain [5, p.196]. Later, Gärdenfors adds an *epistemological* argument: "[i]n general, one will not know all the properties of an object" [1, p.135]. In this case the dimension is relevant for the object but there is no information about the location of the object. Gärdenfors proposes to capture this lack of knowledge by assuming that objects are represented by *partial* vectors, i.e., "points where the arguments for some dimensions are undetermined" [1, p.135]. Here we extend this idea by allowing objects to be located in a region (rather than a point) of a domain, i.e., one just have *imprecise* information about the object.

Given these problems, and on the basis of some arguments discussed in the remainder of this section, we prefer to take a more general approach where objects reduce neither to complete nor to partial vectors of a space. Gärdenfors does not distinguish (actual or possible) objects from conjunctions of maximally (modulo the resolution provided by the space) specified properties. Vice versa, we assume objects to be different from the concepts under which they are *classified*.[1] We will see that the distinction between the objects and the conceptual classification system allows to represent both the partial information (*underdetermination*) about the objects, and the *roughness* of concepts.

Conceptual Individuation of Objects. From a cognitive perspective, the reduction of objects to (special kinds of) concepts has been recently criticized by Pylyshyn [6]. According to Pylyshyn the individuation of objects cannot be purely conceptual, the "[c]onceptual identification ultimately requires a nonconceptual basis" [6, p.36]. On the basis of several empirical evidences, Pylyshyn supports the idea that this non-conceptual basis is provided by a lower level mechanism built into the visual system, called FINST, that is in charge of the initial individuation and tracking of objects. At this level, an object is just a "bare demonstrative—it [the FINST mechanism] picks out things without doing so by their properties" [6, p.18], i.e., one does not know *what* has been selected (the type or kind of the object) but one knows *which* object it is" [6, p.94]. This move *decouples* (re-)identification from classification, one no longer needs to rely on some necessary conditions, some sortal concepts, to distinguish an object from the other ones or to follow it through change.

The accurateness of the conceptual classification of an object is dynamic, it depends on the collected information about the object. In terms of the theory of *object files* [7], one "can think of an object file as a way for information to be associated with objects that are selected and indexed by the FINST mechanism. When an object first appears in view (...) a file is established for that object. Each object file has a FINST reference to

[1] Aisbett and Gibson claim that a "point in the space is a state in the associated conceptual system" [5, p.192].

the particular individual to which the information refers" [6, p.38] The file allows us to group, maintain, and upgrade the information associated to the same individual.[2]

Common-sense Objects. Among the domains of conceptual spaces, Gärdenfors includes *time* and *space*. If objects are vectors of *points*, then they *extend* neither in time (they are instantaneous) nor in space. Ontologically, this position is a sort of extreme *four-dimensionalism*, see [8]. By contrast, common-sense objects are usually persistent and spatially-extended individuals. Usually, four-dimensionalists do not include common-sense objects in the ontological inventory, but they recognize their cognitive relevance: common-sense objects reduce to mereological sums of *temporal slices*, sort of spatio-temporal worms. This problem has also been acknowledged in Artificial Intelligence, in particular by [9,10], where common-sense objects are in fact not identified with vectors of points. However, the determination of the unity criteria that the temporal slices have to satisfy to form an object is notoriously problematic.[3] Gärdenfors faces here an even harder problem because his objects are not only instantaneous, they are atomic with respect to all the domains of the conceptual space (space, color, etc.). But let us assume to have a criterion to identify the worms that correspond to common-sense objects. Consider now an apple a that is not uniformly colored, e.g., it has a small olive dot but the rest of its surface is crimson. In this case, despite the olive dot, one is still inclined to classify a under scarlet. Common-sense objects can be *holistically* classified on the basis of the spatial-*distribution* of their properties, they can have *emergent* properties. One one hand, to ascribe these holistic and emergent properties to a common-sense object, it seems necessary to look at the properties of the set of vectors that 'compose' the corresponding worm (take also properties like *being polka-dotted* that have an intrinsic distributional nature [11]). On the other hand, it seems plausible that the similarity judgments—at the basis of the construction of the domains of the conceptual spaces—regard common-sense wholes and not spatial-color-. . . -slices. We are facing here a circularity (analogous to the one addressed by the FINST mechanism of Pylyshyn) that Gärdenfors does not address in detail.

Objects vs. Clusters of Features. Ontologically, the identification of an object with a conjunction of properties reminds the *bundle theory*, see [12, §4], while, in terms of theory of perception (vision), it reminds the *feature-placing* approach to objects, see [13].

For what concerns the link with the bundle theory, the combination of four-dimensionalism with the reduction of objects to bundles of properties raises some known problems [12, §4]. It is not clear how these problems can be addressed in the context of the theory of conceptual spaces. For instance, the "*complex of compresent universals but no object problem*" regards the distinction between possible and actual objects: one has to distinguish the case of an actual *compresence* of properties vs. the case of an hypothetical one. Or, the "*duplication problem*" concerns the fact that the identity principle for bundles (same properties, same bundle) "entails that, necessarily, objects that are indiscernible with respect to their non-relational properties are identical, but pairs of non-identical objects sharing all their non-relational universals are certainly possible"

[2]"Notice that it [the object file] needs not (and most likely does not) have information about the properties that caused the index to be assigned or caused the object file to be created, nor does it necessarily contain information about which properties allow the individual object to be tracked" [6, p.38].

[3]For instance, spatio-temporal continuity, often advocated as the only needed criterium, fails in the case of complex artifacts or social-objects (like chess-pieces).

[12, p.108]. These problems are particularly challenging for the framework of conceptual spaces because (i) it is not clear how modal properties can be built from similarity judgements; and (ii) all the domains represent *non-relational* properties.

In vision, the standard feature-placing approach (see [13]) assumes features to have an independent status and material objects to be constructed out of them by clustering features according to their spatial location that is "the ultimate referential index". Vice versa, Matthen claims that "material objects come first [and] features are attributed to them after they are identified" [14, p.324]. Vision attributes *features* to objects, *conscious* visual states have an object-attribute form, and "[t]he 'message' that sensory states convey to the perceiver is assembled from (a) sensory referential components, which identify objects, and (b) descriptive components, which identify sense-features" [14, p.87]. A similar position is embraced by Dretske in [15]. One of the main reasons is that "[p]erceptions of change and motion demand an identity that underlies change. Locations do not provide such an identity" [14, p.282]. This view is supported by empirical evidence like the color-*phi*-phenomenon—where two different colored spots lit for 150 msec, with a 50 msec interval, are perceived as a moving color-changing object, see [16]—or by the fact that, conversely, our visual system can track several spatially superimposed objects, see [17]. Thus, Matthen supports the objects-properties decoupling advocated by Pylyshyn already at the level of vision.

Coincident objects. The possibility to have different *coincident* objects has been deeply discussed in the philosophical literature on material constitution, see [18]. For instance, during a part of its life, a statue can coincide with a given amount of clay that constitutes it. The statue and the clay, at a given space-time, are not identical because they have different properties, see [19]. For instance, at t, the statue could have some aesthetic properties that the clay does not have (or are not relevant for it). They can also have incompatible properties, for instance, the statue and the amount of clay could have different prices. Different (Gärdenfors) objects can then have the same spatio-temporal coordinates. However, without the reference to objects, it is hard to see how the theory of conceptual spaces can manage the detection of an inconsistency from the one of a (material) constitution. In addition, the distinction between the statue and the clay is usually based on *modal* properties: the amount of clay, but not the statue, can survive a squeeze; vice versa, the statue, but not the clay, can survive a substitution of some of its parts. The modality is not necessarily temporal: in some worlds the amount of clay and the statue could coincide during their whole life. Five-dimensionalism could refer to modal-spatio-temporal worms to solve the problem while, as said, it is not clear how this modal domain can be represented in the context of conceptual spaces.

Epistemological Underdetermination and Roughness. We have seen that Gärdenfors accepts *epistemological* indetermination: some properties of an object could be unknown, i.e., the information about a domain that is relevant to the object is totally lacking. Let us suppose to have the following two partial vectors: $v_1 = \langle p_1, p_2, unknown, unknown\rangle$ and $v_2 = \langle unknown, unknown, p_3, p_4\rangle$. What authorizes Gärdenfors to claim that these vectors represent partial information about two different objects? Why v_1 and v_2 can not be seen as pieces of information that concern the same object? In addition, it is not clear to us why he does not also consider epistemological *underdetermination*, for instance when one knows that an object is *red* but not its exact shade of red (scarlet, crimson, magenta, etc.). Underdetermination is particularly important for generalizing the idea of conceptual spaces to the one of *classification system* used in scientific theories, in particular in measurement theory, see [20]. In this case one needs to take into account the resolution of the instrument used to collect the measurement.[4] Sometimes one can only assign a *determinable* property[5] to an individual and this property would be represented by a region, not a point, in a domain. As limit case one can just know that the object is located inside a domain, e.g., it is colored. By contrast, the distinguished point $*$ represents the fact that the object does not have a location in a given domain.

Finally, note that the metrics in the domains is mainly used, starting from a set of prototypes, to produce a Voronoi tessellation of the domains into convex regions. In this case all the properties and the concepts are *sharp*, it is possible to establish with certitude if an object is or is not classified under them, i.e., concepts are not fuzzy nor rough, a quite implausible situation in cognitive terms. [23] proposes to determinate the regions that corresponds to concepts on the basis of several prototypes of the same concept. Here we do not enter this discussion, we just introduce a form of partial information about *concepts* that we will represent by means of *rough sets* [24].

4. The formal framework

We present our formal framework that relies on previous work on conceptual spaces but departs from the original view of Gärdenfors to account for the ontological and epistemological problems addressed in the previous section and for some technical issues that we shall encounter.

4.1. Definition of Conceptual Space and Concept

Our definition of conceptual spaces is inspired by the formalizations based on vector spaces provided in [5] and [25]. A *domain* Δ is given by a number of n quality dimensions $Q_1,\dots,Q_n$ endowed with a distance d_Δ that usually depends on the distances defined on its quality dimensions. Following [5], we assume that every domain contains the distinguished point $*$.[6] A conceptual space is defined by Gärdenfors as a set of domains $\{\Delta_1,\dots,\Delta_n\}$. We simplify the model by putting the following definition:

Definition 1. *A conceptual space is a subset of the cartesian product of n domains:* $\mathscr{C} \subseteq \Delta_1 \times \cdots \times \Delta_n$.

Our definition is weaker than the one proposed by [5], as we are taking any subset of the cartesian product as a conceptual space. This is motivated just as a simplifying move. Stronger definitions, that express, for instance, separability and integrality of the domains, can be retrieved by putting constraints on $\mathscr{C}$. A point of a conceptual space with n domains is an element $x \in \mathscr{C}$, that is, $x = \langle x_1,\dots,x_n\rangle$ is a vector of values in each domain, i.e., we do not explicitly consider the dimensions of the domains and the reduction of the distances d_{Δ_i} to the ones of the dimensions. These aspects are not relevant to our aims.

---

[4] Interestingly, measurement theories too consider the object under measurement as external to and independent of the measurement instruments, see [21].

[5] Determinable properties are opposed to *fully-determinate* ones, see [22], i.e., maximally resolving properties according to the available sensors that are represented by points in the domain.

[6] Recall that $*$ means that a certain quality is not applicable to a certain object.

Gärdenfors represents *concepts* as regions in conceptual spaces. In addition, he assumes that *natural* concepts correspond to sets of convex regions (that represent natural properties) in a number of domains with a salience assignment. We leave salience for future work and, to provide an interpretation to disjunctions and negations of concepts, we do not concentrate on natural concepts, e.g., the union or the complement of convex regions, in general, is not convex. A *sharp concept* is then represented just as a subset $R \subseteq \mathscr{C}$. Vice versa, we represent *rough concepts* by *rough sets* [24]. Following [26], a rough set of $\mathscr{C}$ is specified by means of a pair of sets $\langle A, B \rangle$ such that $A \subseteq B \subseteq \mathscr{C}$: A represents the interior of the rough set, B is its exterior, and $B \setminus A$ its boundary. The intuition behind a rough concept C is that one does not have a sharp definition of C, i.e., only the properties that belong to its interior are necessary for all the instances of C. The properties that belong to its boundary are, in general, satisfied only by some instances of C. Therefore, for the objects placed at the boundary of C, one can neither conclude that they are C-instances nor that they are not C-instances. In terms of *prototype theory*, one can consider a *graded membership* with two thresholds, one associated with A, the other with B. These thresholds can also be defined on the basis of the distances d_Δ. Finally, note that sharp concepts are just a special case of rough ones, i.e., they can be represented by rough sets with form $\langle A, A \rangle$.

4.2. Logical Language and Models

We introduce a predicative language $\mathscr{L}$ with a countable set of individual constants $\mathbf{C}$ and a countable set of atomic predicates *Pred*. We only discuss predication, thus we do not introduce variables and quantifiers at this point.

We assume that we can build complex predicates from atomic predicates to model the composition of concepts along the lines of [27]. For instance, we will write sentences $(P \wedge Q)(a)$ to indicate that a satisfies the conjunction of the concepts P and Q. The motivation for this move is the following. In the classical case, when $P(a) \wedge Q(a)$ is true, the individual associated to a belongs to the intersection of the interpretation of P and Q. Since we are here introducing non-standard semantics for predicates and objects, we want to test whether, also in the non-standard cases, if, for instance, the concepts P and Q apply to a—i.e., $P(a) \wedge Q(a)$—then the complex concept $(P \wedge Q)$ also applies to a—i.e., $(P \wedge Q)(a)$. For this reason, we add the concepts that correspond to the conjunction, disjunction, and negation of concepts, by introducing the following set of predicate terms:

$$\mathbf{P} := P \in Pred \mid \neg P \mid P \wedge P \mid P \vee P$$

The set of atomic formula *Atom* is defined as follows: $Q(a) \in Atom$ iff $Q \in \mathbf{P}$ and $a \in \mathbf{C}$. This definition extends to the language as follows.

$$\mathscr{L} := Q(a) \in Atom \mid \neg \phi \mid \phi \wedge \phi \mid \phi \vee \phi$$

We introduce now the structure that we use to define the models of our language.

A *conceptual* structure for $\mathscr{L}$ is a tuple $S = \langle \mathscr{C}, \mathscr{O}, \varepsilon, \iota, \sigma \rangle$ where:

- $\mathscr{C}$ is a conceptual space;
- $\mathscr{O}$ is a non empty set of *objects*;
- ε is a function that maps individual constants into objects, $\varepsilon : \mathbf{C} \to \mathscr{O}$;

- ι is a function that maps predicates into rough sets of $\mathscr{C}$, $\iota : \mathbf{P} \to \mathscr{P}(\mathscr{C}) \times \mathscr{P}(\mathscr{C})$;
- σ is a function that maps objects into regions of $\mathscr{C}$, $\sigma : \mathscr{O} \to \mathscr{P}(\mathscr{C})$.

A *conceptual model* M is then obtained by adding a valuation function $||\cdot||_M$ that maps formulas to a suitable set of truth-values. $||\cdot||_M$ depends on ε, ι, σ but also on the choice of the set of truth-values. In the remainder of this section, we consider several examples of valuations that are designed to meet viable views of concepts and objects.

Contra Gärdenfors, in Section 3, we argued against the reduction of objects to points in a space. In our conceptual structures, the objects in $\mathscr{O}$—that are clearly distinct from the points and the regions of $\mathscr{C}$—may provide a direct representation of common-sense objects without necessarily embracing a realist stance or a strong ontological position. For instance, endurantist, perdurantist, or costructivist views of objects are all compatible with our framework. In a cognitive view, the objects in $\mathscr{O}$ could also be intended as the mental counterparts of the physical objects, e.g., the object files previously discussed.

Formally, the objects in $\mathscr{O}$ provide the denotation of the individual constants $\mathbf{C}$. The function ε, called *extension*, associates individual constants to objects, thus it plays the role of the interpretation function in a standard first-order model. The σ function, called *classification*, locates objects in $\mathscr{C}$, i.e., it characterizes an object in terms of its properties. Different objects may then have the same classification, allowing us to deal with the problem of coincidence. In addition, the underdetermination of objects can be represented by locating them in regions (rather than points) of $\mathscr{C}$, i.e., along some domains, the exact (fully determinate) property of an object is not known. To avoid underdetermination, it is always possible to σ-map all the objects to singletons.

The (rough) concepts of $\mathscr{C}$ provide the semantics of the predicates of $\mathscr{L}$. The function ι, called *intension*, maps a predicate $P \in \mathbf{P}$ into a rough set of $\mathscr{C}$ that we indicate by $\langle \underline{\iota}(P), \overline{\iota}(P) \rangle$. The (rough) extension of a predicate P—the set of objects that (roughly) satisfy P—can be defined on the basis of how an object is positioned in $\mathscr{C}$ (via σ) with respect to the intension of P (see next subsections).[7] In this way, we partially capture the intension of predicates, i.e., their interpretation is not reduced to a set of objects. In addition, to guarantee compositionality, the intension of non-atomic predicates—i.e., $\iota(\neg P)$, $\iota(P \wedge Q)$, and $\iota(P \vee Q)$—must be a function of the intensions of its atomic parts—i.e., $\iota(P)$ and $\iota(Q)$. For instance, it is necessary to build the concept associated (via ι) to $P \wedge Q$ starting from the two concepts associated (via ι) to P and Q. By relying on this function we can then check whether $P(a) \wedge Q(a)$ and $(P \wedge Q)(a)$ have the same truth-value.

Finally, note that conceptual structures do not explicitly consider time. Differently from Gärdenfors, we do not include time as a domain of $\mathscr{C}$. Actually, at this stage, our framework is intended to represent only a snapshot of the world. The full treatment of time would require (i) to add the set $\mathscr{T}$ of times; and (ii) to add a temporal argument to σ, $\sigma : \mathscr{O} \times \mathscr{T} \to \mathscr{P}(\mathscr{C})$, i.e., the location of an object into a conceptual space depends on time. This move requires an extension of the language $\mathscr{L}$ with temporal operators that we leave for future work.

We can distinguish four different kinds of models for our language $\mathscr{L}$:

1. Objects are completely determined and concepts are sharp.
2. Objects are completely determined and concepts are rough.

[7] This is a Fregean perspective on predicates: the intension of a predicate (its *Sinn*) is a mean to obtain the extension, cf. [28].

3. Objects are underdetermined and concepts are sharp
4. Objects underdetermined and concepts are rough.

The first case corresponds to the Gärdenfors's view, with the significant conceptual difference that we do not identify objects with points in the space. Here, objects are associated to points and concepts are represented by subsets of $\mathscr{C}$. The other cases introduce uncertainty concerning an object or a concept and this assumptions is reflected, as we shall see, on the compositional constraints on ι, on the definition of predication, and on the truth values that are needed for reasoning about predicative sentences. The analysis of these four cases is the topic of the following sections.

4.3. *Objects are determined and concepts are sharp*

In this scenario, objects are always associated to singletons of $\mathscr{C}$. By slightly abusing our previous definitions, we assume here $\sigma(o) \in \mathscr{C}$. In this case, a predicate in the language is associated to a subset of $\mathscr{C}$, that is, $\iota(P) \in \mathscr{P}(\mathscr{C})$. We can then consider the standard set of truth values $\{t, f\}$ and adapt to our framework the usual Tarskian semantic clause of atomic formulas (where $a \in \mathbf{C}$ and $P \in \mathbf{P}$, i.e., P can also be a complex predicate):

$$||P(a)||_M = t \text{ iff } \sigma(\varepsilon(a)) \in \iota(P) \tag{1}$$

i.e., $P(a)$ is true in a model if the vector of values that represents the information concerning the object denoted by a is a member of the set of vectors that represents the information concerning P, the intension of P. The difference with the standard first-order clause is that the predication applies to $\sigma(\varepsilon(a))$ rather than directly to $\varepsilon(a)$.

It is easy to check that, by assuming a classical semantics for the logical connectives $(\neg, \wedge, \vee)$ and by defining the intension of non-atomic predicates as:

$$\iota(\neg P) = \mathscr{C} \setminus \iota(P) \tag{2}$$

$$\iota(P \wedge Q) = \iota(P) \cap \iota(Q) \tag{3}$$

$$\iota(P \vee Q) = \iota(P) \cup \iota(Q) \tag{4}$$

the correspondence between the logical connectives and the corresponding complex predicates is guaranteed, i.e., the mechanism of conceptual composition is aligned with the semantics of connectives.

From (1), the extension of a predicate can be defined by means of its intension:

$$\varepsilon(P) = \{o \in \mathscr{O} \mid \sigma(o) \in \iota(P)\} \tag{5}$$

i.e., the extension is given by the set of objects that are classified into points of $\iota(P)$.

4.4. *Objects are determined and concepts are rough*

In this scenario, objects are still associated to singletons (we assume again $\sigma(o) \in \mathscr{C}$) but predicates are now associated to rough sets: $\iota(P) = \langle \underline{\iota}(P), \overline{\iota}(P) \rangle$. We can define three regions (that depend on P): $POS_P = \underline{\iota}(P)$, $NEG_P = \mathscr{C} \setminus \overline{\iota}(P)$, and $BN_P = \overline{\iota}(P) \setminus \underline{\iota}(P)$. Given the fact that $\sigma(o) \in \mathscr{C}$, a three-valued logic—a logic whit the truth-values $\{t, u, f\}$—can then capture the three situations where $\sigma(o)$ belongs to POS_P, NEG_P, or BN_P; i.e., we can consider the following semantic clauses for predication:

$$||P(a)||_M = t \text{ iff } \sigma(\varepsilon(a)) \in POS_P \tag{6}$$

$$||P(a)||_M = u \text{ iff } \sigma(\varepsilon(a)) \in BN_P \tag{7}$$

$$||P(a)||_M = f \text{ iff } \sigma(\varepsilon(a)) \in NEG_P \tag{8}$$

i.e., the classification of an object under a concept may be partial in the sense that we may be able to place an object only within the boundary of a concept.

It is then possible to check that the following constraints compositionally define the intension of non-atomic predicates:[8]

$$\iota(\neg P) = \langle \mathscr{C} \setminus \bar{\iota}(P), \mathscr{C} \setminus \underline{\iota}(P) \rangle \tag{9}$$

$$\iota(P \wedge Q) = \langle \underline{\iota}(P) \cap \underline{\iota}(Q), \bar{\iota}(P) \cap \bar{\iota}(Q) \rangle \tag{10}$$

$$\iota(P \vee Q) = \langle \underline{\iota}(P) \cup \underline{\iota}(Q), \bar{\iota}(P) \cup \bar{\iota}(Q) \rangle \tag{11}$$

This treatment of complex predicates is aligned with the Kleene three valued logic, where the semantics of the connectives is provided by the following truth tables:

$\neg$	t	u	f
	f	u	t

$\wedge$	t	u	f
t	t	u	t
u	u	u	u
f	t	u	f

$\vee$	t	u	f
t	t	u	t
u	u	u	u
f	t	u	f

For example, take the case $||P(a)||_M = ||Q(a)||_M = u$, we check that $||(P \wedge Q)(a)||_M = u$. By definition, $\sigma(\varepsilon(a)) \in BN_P$ and $\sigma(\varepsilon(a)) \in BN_Q$, thus $\sigma(\varepsilon(a)) \in BN_P \cap BN_Q$. That is $\sigma(\varepsilon(a)) \in (\bar{\iota}(P) \setminus \underline{\iota}(P)) \cap (\bar{\iota}(Q) \setminus \underline{\iota}(Q)) = (\bar{\iota}(P) \cap \bar{\iota}(Q)) \setminus (\underline{\iota}(P) \cup \underline{\iota}(Q))$, which is included in $(\bar{\iota}(P) \cap \bar{\iota}(Q)) \setminus (\underline{\iota}(P) \cap \underline{\iota}(Q)) = BN_{P \wedge Q}$. Therefore $||(P \wedge Q)(a)||_M = u$. The proof of the other cases is similar to the one presented in [26], so we omit it.

The extension of a predicate is then a rough set on $\mathscr{O}$. The intension ι of a predicate P allows to compute its extension which partitions the objects of $\mathscr{O}$ into three sets: the set of those objects that are in the positive part of P, the set of those that are in its boundary, and the set of those that are in its negative part. That is, the extension of P is (the rough set of $\mathscr{O}$) $\varepsilon(P) = \langle \underline{\varepsilon}(P), \bar{\varepsilon}(P) \rangle$ such that:[9]

$$\underline{\varepsilon}(P) = \{o \in \mathscr{O} \mid \sigma(o) \in \underline{\iota}(P)\} \tag{12}$$

$$\bar{\varepsilon}(P) = \{o \in \mathscr{O} \mid \sigma(o) \in \bar{\iota}(P)\} \tag{13}$$

4.5. Objects are underdetermined and concepts are sharp

In this scenario, both objects and concepts are associated to subsets of $\mathscr{C}$, i.e., $\sigma(o) \in \mathscr{P}(\mathscr{C})$ and $\iota(P) \in \mathscr{P}(\mathscr{C})$. In this case, different interpretations of the underdetermination of the objects are possible. Firstly, we may assume that $P(a)$ is verified only in case

[8] This definitions are based on the approach in [26].

[9] Trivially, $\underline{\varepsilon}(P) \subseteq \bar{\varepsilon}(P)$.

every point in $\sigma(\varepsilon(a))$ is in $\iota(P)$ while having even a bit of information concerning a that contrasts with information concerning P is sufficient to falsify $P(a)$. We call this view *certainty-of-truth* reading: in order to claim that a sentence is true, every bit of information that we have must support the truth of $P(a)$. The semantic conditions are in this case expressible by means of the standard truth values $\{t, f\}$:

$$||P(a)||_M = t \text{ iff } \sigma(\varepsilon(a)) \subseteq \iota(P) \tag{14}$$

$$||P(a)||_M = f \text{ iff } \sigma(\varepsilon(a)) \nsubseteq \iota(P) \tag{15}$$

Secondly, we may view $P(a)$ as falsified only in case every bit of information concerning a contrasts with P. This second option is a *certainty-of-falsity* reading: we can claim a sentence is false only in case every evidence that we have about a contrasts with P. In this reading, the verification of $P(a)$ is the case in which some information concerning a matches those concerning P. Hence, we would have:

$$||P(a)||_M = t \text{ iff } \sigma(\varepsilon(a)) \cap \iota(P) \neq \emptyset \tag{16}$$

$$||P(a)||_M = f \text{ iff } \sigma(\varepsilon(a)) \cap \iota(P) = \emptyset \tag{17}$$

A third option that we list here is to demand both certainty with respect to truth and certainty with respect to falsity. This perspective views verification as the inclusion of every bit of information, falsification as the mismatch of every bit of information, and introduces a third truth value (labelled u) for possibly conflicting information.

$$||P(a)||_M = t \text{ iff } \sigma(\varepsilon(a)) \subseteq \iota(P) \tag{18}$$

$$||P(a)||_M = u \text{ iff } \sigma(\varepsilon(a)) \cap \iota(P) \neq \emptyset \text{ and } \sigma(\varepsilon(a)) \nsubseteq \iota(P) \tag{19}$$

$$||P(a)||_M = f \text{ iff } \sigma(\varepsilon(a)) \cap \iota(P) = \emptyset \tag{20}$$

Since $\iota(P)$ is a set, it is meaningful to define the intension of non-atomic predicates as in Section 4.3 (see (2)-(4)). It is easy to check that this defintion is aligned with the standard semantics for connectives in the case of the first two readings (certainty-of-truth and certainty-of-falsity).

In the case of the third reading, we can show that the same definition of complex intensions is aligned with the Kleene truth table for the connective $\neg$ (see Section 4.4). Suppose $||P(a)||_M = u$, then $\sigma(\varepsilon(a)) \cap \iota(P) \neq \emptyset$ and $\sigma(\varepsilon(a)) \nsubseteq \iota(P)$, which entails that also $\sigma(\varepsilon(a)) \cap (\mathscr{C} \setminus \iota(P)) \neq \emptyset$ and $\sigma(\varepsilon(a)) \nsubseteq (\mathscr{C} \setminus \iota(P))$, that is $||(\neg P)(a)||_M = u$.

However, the situation for the conjunction is more delicate. As we did in Section 4.4, suppose that $||P(a)||_M = ||Q(a)||_M = u$. We check the truth-value of $||(P \wedge Q)(a)||_M$.

From the assumptions, we have that $\sigma(\varepsilon(a)) \cap \iota(P) \neq \emptyset$ and $\sigma(\varepsilon(a)) \nsubseteq \iota(P)$ and $\sigma(\varepsilon(a)) \cap \iota(Q) \neq \emptyset$ and $\sigma(\varepsilon(a)) \nsubseteq \iota(Q)$. By (3), $\iota(P \wedge Q) = \iota(P) \cap \iota(Q)$. Therefore, there are two possible cases that are compatible with the assumptions:[10]

1. $\sigma(\varepsilon(a)) \cap \iota(P \wedge Q) = \emptyset$, therefore $||(P \wedge Q)(a)||_M = f$;
2. $\sigma(\varepsilon(a)) \cap \iota(P \wedge Q) \neq \emptyset$, then, since $\sigma(\varepsilon(a)) \nsubseteq \iota(P)$ and $\sigma(\varepsilon(a)) \nsubseteq \iota(Q)$, $\sigma(\varepsilon(a)) \nsubseteq \iota(P \wedge Q)$, therefore $||(P \wedge Q)(a)||_M = u$.

[10]The first case is possible because a set A may overlap with B and with C, but that does not entail that A overlaps with the intersection of B and C.

This means that the value of the conjunction is not determined in case the values of the conjuncts are both u. This fact has been already noticed in [27].[11] To cope with that, in [27], the authors introduced the following *non-deterministic* truth tables where the value of the conjunction and of the disjunction on the pair of values $\langle u,u\rangle$ is a *set* of values, instead of being a single value as usual.

$\neg$	t	u	f
	f	u	t

$\wedge$	t	u	f
t	t	u	t
u	u	$\{u,f\}$	u
f	t	u	f

$\vee$	t	u	f
t	t	u	t
u	u	$\{t,u\}$	u
f	t	u	f

It is interesting to notice that, in our framework, it is the combination of underdetermination of objects and sharpness of property that leads to a more sophisticated form of uncertainty, that is, the non-determinism of the logical connectives.

Although the intension of a predicate P is in this case sharp, i.e. $\iota(P)$ is a set, we may still wonder whether its extension is a rough set. Denote by $\underline{\varepsilon}(P)=\{o\in\mathscr{O}\mid\sigma(o)\subseteq\iota(P)\}$ and by $\bar{\varepsilon}(P)=\{o\in\mathscr{O}\mid\sigma(o)\cap\iota(P)\neq\emptyset\}$. The rough set that corresponds to the extension of P is defined as $\varepsilon(P)=\langle\underline{\varepsilon}(P),\bar{\varepsilon}(P)\rangle$ where, clearly, $\underline{\varepsilon}(P)\subseteq\bar{\varepsilon}(P)$. In order to respect conditions (18), (19) and (20), the extension of a predicate P can be computed via the following function ε_P:

$$\varepsilon_P(o)=\begin{cases} t \text{ if } o\in\underline{\varepsilon}(P) \text{ iff } \sigma(o)\subseteq\iota(P)\\ u \text{ if } o\in\bar{\varepsilon}(P)\setminus\underline{\varepsilon}(P) \text{ iff } \sigma(o)\cap\iota(P)\neq\emptyset \text{ and } \sigma(o)\not\subseteq\iota(P)\\ f \text{ if } o\in\mathscr{C}\setminus\bar{\varepsilon}(P) \text{ iff } \sigma(o)\cap\iota(P)=\emptyset\end{cases} \tag{21}$$

For instance, the extension of $\neg P$ can be defined by $\varepsilon(\neg P)=\langle\underline{\varepsilon}(\neg P),\bar{\varepsilon}(\neg P)\rangle$, that functionally depends on $\varepsilon(P)$, since $\underline{\varepsilon}(\neg P)=\mathscr{C}\setminus\bar{\varepsilon}(P)$ and $\bar{\varepsilon}(\neg P)=\mathscr{C}\setminus\underline{\varepsilon}(P)$ (cf. [27]). Therefore $\varepsilon(\neg P)=\langle\underline{\varepsilon}(\neg P),\bar{\varepsilon}(\neg P)\rangle=\langle\mathscr{C}\setminus\bar{\varepsilon}(P),\mathscr{C}\setminus\underline{\varepsilon}(P)\rangle$. But, for conjunctions and disjunctions, the argument above fails. The non-determinism of truth values entails that the value of the function $\varepsilon_{P\wedge Q}$ that computes the extension of $P\wedge Q$ is not functionally dependent on the values of ε_P and ε_Q. We can still associate a rough set to $P\wedge Q$, defined by $\varepsilon(P\wedge Q)=\langle\underline{\varepsilon}(P\wedge Q),\bar{\varepsilon}(P\wedge Q)\rangle$. However, this set is not dependent on $\varepsilon(P)=\langle\underline{\varepsilon}(P),\bar{\varepsilon}(P)\rangle$ and $\varepsilon(Q)=\langle\underline{\varepsilon}(Q),\bar{\varepsilon}(Q)\rangle$. Thus, in this case the extensions are not compositional, whereas intensions are.

Beside the three readings that have we explored, more sophisticated approaches to predication in this case require defining *aggregation procedures* that measure how many points of $\sigma(\varepsilon(a))$ are demanded to be in $\iota(P)$, cf. [29]. E.g. one can define a majoritarian aggregation where $P(a)$ is true (false) only if a majority of points is in $\iota(P)$ ($\mathscr{C}\setminus\iota(P)$). We leave this approach for future work.

4.6. *Objects are underdetermined and concepts are rough*

In this scenario, concepts are associated to rough sets and objects to sets of points. Note that this case is the most general of the four. As we have done is Section 4.5, we may

[11] The cause of non-determinism is explained in [27] by noticing that the intersection of the positive parts of two sets is only included in and not equal to the positive part of the intersection of two sets. Therefore, the boundary of the intersection is not the intersection of the boundaries. Analogous situation for the union. Note that the proof of the other combinations of truth values is similar to the one presented in [27].

choose how to read the underdetermination of objects. In the certainty-of-truth reading, it makes sense to assume just two truth values and put the following definition:

$$||P(a)||_M = t \text{ iff } \sigma(\varepsilon(a)) \subseteq POS_P \tag{22}$$

$$||P(a)||_M = f \text{ iff } \sigma(\varepsilon(a)) \nsubseteq POS_P \tag{23}$$

while, for the certainty-of-falsity reading, we have the following definition:

$$||P(a)||_M = t \text{ iff } \sigma(\varepsilon(a)) \nsubseteq NEG_P \tag{24}$$

$$||P(a)||_M = f \text{ iff } \sigma(\varepsilon(a)) \subseteq NEG_P \tag{25}$$

However, these two readings basically amount to forget the structure of rough sets. The combined reading is more interesting, we phrase it for this scenario as follows.

$$||P(a)||_M = t \text{ iff } \sigma(\varepsilon(a)) \subseteq POS_P \tag{26}$$

$$||P(a)||_M = u \text{ iff } \sigma(\varepsilon(a)) \nsubseteq POS_P \text{ and } \sigma(\varepsilon(a)) \nsubseteq NEG_P \tag{27}$$

$$||P(a)||_M = f \text{ iff } \sigma(\varepsilon(a)) \subseteq NEG_P \tag{28}$$

Certainty of truth means that all the points in $\sigma(\varepsilon(a))$ are certainly P, that is, they are in the positive part of P. The case of falsity means that all the points of $\sigma(\varepsilon(a))$ are certainly not P, that is, they are in the negative part of P. Note that the case of $||P(a)||_M = u$ does not entail that $\sigma(\varepsilon(a))$ is included in BN_P, $\sigma(\varepsilon(a))$ may spread across all the three regions POS_P, NEG_P, and BN_P. We can therefore adapt the treatment in Section 4.5 to this case, noticing that non-determinism of logical connectives is back.

By (10), $\iota(P \wedge Q) = \langle \underline{\iota}(P) \cap \underline{\iota}(Q), \overline{\iota}(P) \cap \overline{\iota}(Q) \rangle$. Suppose that $||P(a)||_M = u$ and $||Q(a)||_M = u$. That means that $\sigma(\varepsilon(a)) \nsubseteq POS_P$, $\sigma(\varepsilon(a)) \nsubseteq NEG_P$, $\sigma(\varepsilon(a)) \nsubseteq POS_Q$, and $\sigma(\varepsilon(a)) \nsubseteq NEG_Q$. From $\sigma(\varepsilon(a)) \nsubseteq POS_P$ and $\sigma(\varepsilon(a)) \nsubseteq POS_Q$, we can infer that $\sigma(\varepsilon(a)) \nsubseteq POS_{P \wedge Q}$, thus $||(P \wedge Q)(a)||_M$ is not t. However, $\sigma(\varepsilon(a)) \nsubseteq NEG_P$ and $\sigma(\varepsilon(a)) \nsubseteq NEG_Q$ does not entail that $\sigma(\varepsilon(a)) \nsubseteq NEG_{P \wedge Q} = \mathscr{C} \setminus (\overline{\iota}(P) \cap \overline{\iota}(Q))$. We have again two cases, $\sigma(\varepsilon(a)) \subseteq \mathscr{C} \setminus (\overline{\iota}(P) \cap \overline{\iota}(Q))$, and in this case $||(P \wedge Q)(a)||_M = f$, or $\sigma(\varepsilon(a)) \nsubseteq \mathscr{C} \setminus (\overline{\iota}(P) \cap \overline{\iota}(Q))$, and in this case $||(P \wedge Q)(a)||_M = u$.

We conclude by noticing that this scenario seems to require more than three truth-values, to account for all the possible combinations of spreading an object through the three parts of a concept. We leave a proper treatment of this case for future work.

5. Conclusion and future work

We have seen that the nature of predication in conceptual spaces depends on our understanding of uncertainty towards objects or concepts. Moreover, we have argued that managing uncertainty is unescapable when doing semantics with a cognitive perspective. We have approached here the semantics of predicative sentences and of propositional connectives; future work shall extend this treatment to a full first-order language, in particular, we shall approach the delicate issue of understanding n-ary relations in conceptual spaces [30] and propose a cognitively justified modelling of quantification.

References

[1] P. Gärdenfors. *Conceptual spaces - the geometry of thought*. MIT Press, 2000.

[2] C. Masolo and D. Porello. A cognitive view of relevant implication. In A. Lieto, C. Battaglino, D. P. Radicioni, and M. Sanguinetti, editors, *Proceedings of the 3rd International Workshop on Artificial Intelligence and Cognition*, pages 40–53, Turin, Italy, 2015. CEUR.

[3] S. S. Jones and L. M. Koehly. Muldimensional scaling. In G. Keren and C. Lewis, editors, *A Handbook for Data Analysis in the Behavioral Sciences*, pages 95–163. Lawrence Erlbaum Associates, 1993.

[4] P. Gärdenfors. *The geometry of meaning: Semantics based on conceptual spaces*. MIT Press, 2014.

[5] J. Aisbett and G. Gibbon. A general formulation of conceptual spaces as a meso level representation. *Artificial Intelligence*, 133(1):189–232, 2001.

[6] Z. W. Pylyshyn. *Things and Places. How the Mind Connects with the World*. MIT Press, 2007.

[7] D. Kahneman, A. Treisman, and B. J. Gibbs. The reviewing of object files: Object-specific integration of information. *Cognitive psychology*, 24(2):175–219, 1992.

[8] T. Sider. *Four-Dimensionalism. An Ontology of Persistence and Time*. Clarendon Press, Oxford, 2001.

[9] A. Chella, M. Frixione, and S. Gaglio. A cognitive architecture for artificial vision. *Artificial Intelligence*, 89(1):73–111, 1997.

[10] A. Chella, M. Frixione, and S. Gaglio. Understanding dynamic scenes. *Artificial intelligence*, 123(1):89–132, 2000.

[11] J. Parsons. Distributional properties. *Lewisian themes*, pages 173–180, 2004.

[12] D. Ehring. *Tropes: Properties, Objects, and Mental Causation*. Oxford Univeristy Press, 2011.

[13] A. Clark. *A Theory of Sentience*. Oxford University Press, 2000.

[14] M. Matthen. *Seeing, Doing, and Knowing. A Philosophical Theory of Sense Percetion.* Oxford Univeristy Press, 2005.

[15] F. Dretske. *Naturalizing the Mind*. MIT Press, 1995.

[16] D. C. Dennett. *Consciousness Explained*. Touchstone Books, Simon & Shuster, 1991.

[17] E. Blaser, Z. W. Pylyshyn, and A. O. Holcombe. Tracking an object through feature space. *Nature*, 408:196–99, 2000.

[18] M. Rea, editor. *Material Constitution*. Rowman and Littlefield Publishers, 1996.

[19] L. R. Baker. *The Metaphysics of Everyday Life*. Cambridege Univerity Press, 2007.

[20] L. Mari and A. Giordani. Quantity and quantity value. *Metrologia*, 49:756–764, 2012.

[21] A. Frigerio, A. Giordani, and L. Mari. Outline of a general model of measurement. *Synthese*, 175(2):123–149, 2010.

[22] D. H. Sanford. Determinates vs. determinables. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Spring 2013 Edition.

[23] L. Decock and I. Douven. Similarity after goodman. *Review of philosophy and psychology*, 2(1):61–75, 2011.

[24] Z. Pawlak. *Rough sets: Theoretical aspects of reasoning about data*, volume 9. Springer, 1991.

[25] M. Raubal. Formalizing conceptual spaces. In *Formal ontology in information systems, proceedings of the third international conference (FOIS 2004)*, volume 114, pages 153–164, 2004.

[26] S. Akama and T. Murai. Rough set semantics for three-valued logics. In K. Nakamatsu and J. Abe, editors, *Advances in Logic Based Intelligent Systems*, volume 132 of *Frontiers in Artificial Intelligence and Applications*, pages 242–247. IOS press, 2005.

[27] A. Avron and B. Konikowska. Rough sets and 3-valued logics. *Studia Logica*, 90(1):69–92, 2008.

[28] G. Frege. On function and concept. In *Translations from the Philosophical Writings of Gottlob Frege*, pages 1–128. Oxford: Blackwell, third edition, 1980.

[29] D. Porello and U. Endriss. Ontology merging as social choice: judgment aggregation under the open world assumption. *J. Log. Comput.*, 24(6):1229–1249, 2014.

[30] C. Masolo and D. Porello. A note on the representation of relations in conceptual spaces. In *Atti della XII Conferenza Annuale della Associazione Italiana di Scienze Cognitive*, volume 9, pages 117–121. NeaScience, 2015.