In Defense of a Trope-Based Ontology for Conceptual Modeling: An example with the foundations of Attributes, Weak Entities and Datatypes

Giancarlo Guizzardi[1,2], Claudio Masolo[1], Stefano Borgo[1]

[1]Laboratory for Applied Ontology (ISTC-CNR), Trento, Italy
{guizzardi,masolo,borgo}@loa-cnr.it
[2] Computer Science Department,
Federal University of Espirito Santo,
Vitoria-ES, Brazil

Abstract. In recent years, there has been a growing interest in approaches that employ foundational ontologies as theoretical tools for analyzing and improving conceptual modeling languages. However, some of these approaches do not always make explicit their ontological commitments. This leads to situations where criticisms resulting from the specific ontological choices made by a particular approach are generalized to the enterprise of ontology as a whole. In this paper we discuss an example of such a case involving the BWW approach. First, we make explicit the ontological commitments underlying that approach by relating it to other possible philosophical alternatives. Second, we construct an ontological theory which commits to a different philosophical position. Third, we show how the ontology proposed here can be used to provide real-world semantics and sound modeling guidelines for the modeling constructs of Attributes, Weak Entities and Datatypes. Finally, we compare the ontology proposed here with BWW, thus demonstrating its benefits.

1 Introduction

In recent years, there has been a growing interest in the use of foundational ontologies for: (i) evaluating conceptual modeling languages; (ii) developing guidelines for their use; (iii) providing real-world semantics for their modeling constructs (e.g., [5,7,18]).

A well-known example of a foundational ontology in the conceptual modeling/information systems area is the *Bunge-Wand-Weber (BWW)* ontology proposed by Wand and Weber in a series of articles (e.g., [5,20]) on the basis of the original metaphysical theory developed by Bunge in [2]. Recently, this ontology has received a number of criticisms in the literature, mainly due to the contrast between the modeling rules proposed by the BWW ontology, on one side, and what is prescribed by linguistic and cognitive studies as well as empirical sessions with practitioners, on the other (e.g., [8,15,16]). One of the strong points of disagreement between BWW and these approaches is the BWW-rule that states that intrinsic properties (roughly attributes) and associations should never be modeled as entity types in an ontologically correct conceptual model.

In a series of papers, Veres and colleagues (e.g., [8,15,16]) offer a detailed analysis and criticism of the general assumptions of the BWW approach. More specifically, in [16], they provide empirical evidence to support a case against the BWW treatment of associations. The danger in many of these criticisms is that they are formulated as general criticisms to ontology, not as specific criticisms to BWW. In other words, criticisms which are consequent of specific choices made in that ontology are generalized to the whole enterprise of ontological foundations for conceptual modeling. However, in the case of Veres et al., the criticisms cannot be against ontology *per se*, since the authors themselves state that they "describe an ontology of conceptual structure" or "psychologically motivated ontology" for the same purpose.

The purpose of this article is three fold. First, we want make explicit some ontological choices made by the BWW approach, and to show that the specific theory of universals underlying this approach is only one among many other philosophically correct theories. Second, we want to propose an alternative foundational theory, and to show how it can be used to provide an ontological interpretation for some conceptual modeling fundamental constructs. In particular, we want to create an ontology that countenances the existence of property instances, and a derived approach for conceptual modeling that accepts the representation of both attributes and associations as classes. Third, we intend to demonstrate that a trope-based ontology such as the one proposed here leads to better results as a foundational theory for conceptual modeling from philosophical, cognitive and practical points of view.

In section 2, we discuss different theories of universals and make explicit the BWW choices regarding these theories. In section 3, we propose a trope-based ontology, which is used in section 4 to provide ontological semantics for the conceptual modeling constructs of *attribute*, *weak entity* and *datatype*. In section 5, we compare the results of section 4 with the approach proposed in [5] that uses the BWW ontology as a foundation for UML as a conceptual modeling language. Section 6 presents some final considerations.

2 Universals, Tropes and Properties

Properties, their interpretation and nature have been discussed at length in the western philosophical tradition giving rise to subtle distinctions and disparate characterizations. Here, we introduce and discuss two general views, namely *universalism* and *trope theory*, and a third position that merges both universals and tropes. The discussion of these theories requires a terminological clarification so our first goal is to introduce a few concepts.

We use the term *particular* to refer to entities that have no instances, that is, entities that cannot be predicated of others; for instance the Tour Eiffel or the Mars planet. Contrast this notion with the notion of *universal* which, on the contrary, characterizes any entity that can have instances, e.g. the color Red or the car model Ferrari 250 GTO. Roughly, the properties (and the relations) used in a language are generally taken to correspond to universals since they are attributed to other entities. The notion of class is generally taken as a formal counterpart of the notion of universal. However, this may be misleading. By universal we mean a characterizing qualification of entities like “a Ferrari 250 GTO”, i.e., a property that an entity may satisfy. The corresponding class is the collection of entities that satisfy that property. Another important notion we need to include is the notion of *trope*. Intuitively, a trope is an instance of a property (i.e., the instance of an objectified property) of a specific entity: the redness of John’s T-shirt is a trope that *inheres* to John’s T-shirt (the host).

Both John’s T-shirt and the redness of John’s T-shirt are particulars. However, they are particulars of very different natures. Tropes are particulars which can only exist in other individuals, i.e., they are *existentially dependent* on other individuals in the way, for instance, the color of an apple *a* depends on *a*, and the electric charge of a conductor *c* depends on *c*. In contrast, particulars such as John, the apple *a*, and the conductor *c* do not inhere in other individuals and, hence, are not existentially dependent entities in this sense. In this article, we give the name *Object* to the latter type of particular.

This brief and rough discussion of objects, universals, classes and tropes tells us that these concepts correspond to different categories of entities. However, which of these entities as well as the relations between them which are countenanced in one’s ontology depend very much on one’s philosophical position w.r.t. the so-called *Problem of Universals* [1,7]. This problem can be summarized as follows: We know that proper names (e.g., Noam Chomsky or Spot) refer to individual entities, but what do general terms (or universal properties) refer to (if anything at all)? We classify objects as being of the same type (e.g., person) and use the same predicate or general term (e.g., red) to different objects. What exactly *is the same* in different objects that justify their belonging to the same category?

Figure 1 illustrates three different representations of the fact: “the particulars a and b share the property being red”. *Universalism* claims that a and b both *instantiate* (I) the being red (Red) universal, i.e. the universal being red is a spatiotemporal independent entity which is somehow wholly present in both a and b (fig.1-left). The trope theory denies the existence of universals as repeatable entities, and considers only tropes and classes of tropes. An important feature that characterizes all tropes is that they can only exist in other individuals, named their *bearers*. A formal relation of *inherence* symbolized as *i(x,y)* is defined to hold between a trope x and its bearer y. Inherence is an irreflexive, asymmetric and intransitive type of existential dependence relation. Moreover, it satisfies the *non-migration principle* [7]. This means that it is not possible for a trope *p* to inhere in two different individuals *a* and *b*. In other words, if we have two particulars *a* (a red apple) and *b* (a red car), and two tropes a_{red} (particular redness of *a*) and b_{red} (particular redness of *b)*, we consider a_{red} and b_{red} to be different individuals, although perhaps qualitatively indistinguishable. What does it mean then to say that *a* and *b* have the *same* color? Due to the non-migration principle, sameness here cannot refer to strict (numerical) identity, but only to a qualitative one (i.e., equivalence in a certain respect). In standard Trope theory, a relation of *resemblance* (≈) is defined between tropes. Hence, tropes can resemble each other to a certain degree and, as in the example above, if they are qualitatively indistinguishable, we say that they exactly resemble each other. This way, Trope theory does not have to commit to the existence of universals as a separate category of abstract entities, since equivalence classes of resembling tropes are enough for predication: a and b have the common property of *being red* because there are two red tropes a_{red} and b_{red} both belonging (∈) to the red class ($|red|_{\approx}$) of tropes that inhere in a and b, respectively. If on the one hand by accepting tropes one does not have to accept universals, on the other hand, these two theories are not incompatible and, actually, they can be merged: a and b have the property being red because the a_{red} trope and the b_{red} trope both are instances of the universal Red. In this case, universals exist but they are instantiated only by tropes.

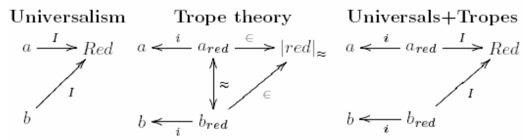

instantiation	I	particular × universal
inherence	i	trope × particular
resemblance	$\approx$	trope × trope
membership	$\in$	trope × class

Fig 1. Different philosophical positions on Universals

2.1 Making Explicit the Ontological Position behind BWW

In BWW, we have a fundamental dichotomy between the notions of *substantial individual* (or *thing*) and *substantial property*. A thing is defined as a substantial individual with all its substantial properties: *"a thing is what is the totality of its substantial properties"* [2, p.111]. Despite of apparently equating a thing which the sum of its properties, Bunge himself does not embraces a type of universalism named the *Bundle of Universals* theory. In fact, he explicitly rejects this theory and, instead, holds another (universalist) position that can be better identified with the *substance-attribute* view [1].

In short, in the former type of theory, particulars are taken as *bundles of universals*, i.e., as aggregates of properties which themselves are repeatable abstract entities. An exemplar theory of this type was proposed by Russel in [12]. For details on this theory as well as for a discussion on the many problems related to it one should refer to [1]. In fact, among the universalist theories, [1] considers the bundle theory of universals to be the weakest one from a philosophical point of view. The *substance-attribute* view makes an explicit distinction between a thing and the properties that the thing has. As a consequence, the theory countenances the existence for every individual of a propertyless *substratum* or *bare particular*. The notion of substratum is strongly associated with the British empiricist philosopher John Locke [1] and due to its mysterious nature it has been the target of strong criticism throughout history. Nonetheless, Bunge claims that as a *"theoretical fiction"* it solves some of the philosophical problems existing in the *bundle of universals* theories [2, p.57]. Hence, for Bunge a thing is a bare particular endowed by all its substantial properties, i.e., he commits to the substance-attribute sort of universalism and, as consequence, denies the existence of *particularized properties*.

In principle, it seems that a thing in BWW could be directly associated to the concept of object in a trope-based theory (figs. 1.center and 1.right). However, there are some important differences between the two. Whilst a BWW-thing can be thought as a substratum instantiating a number of properties (as repeatable abstract entities), objects in a trope-based approach are particulars that bear other particularized properties, or to borrow Simons' phrase, *"particulars in particular clothing"*[14]. Thus, in a trope-theoretical approach, one does not have to make any ontological commitment w.r.t. the nature of the *substratum*. In particular, if necessary, one can dispense with a substratum of a mysterious nature. An example of such a view is the one of Simons' *Nuclear Theory* (ibid.). This approach has the benefits of the substance-attribute view, without having to accept its problems, since the nucleus is akin to a substratum, only not a mysterious one. In BWW, the mysterious substratum cannot be eliminated without putting the theory into a *Bundle of Universals* group. We claim that this flexibility is an advantage of an ontology in which tropes are countenanced.

According to Bunge, only things possess properties. As a consequence, a property cannot have properties, i.e., there are no higher-order properties. This dictum leads to the following BWW modeling principle: *entity types in a conceptual model of a domain should only be used to represent substantial universals* [5]. This principle proscribes the representation of types whose instances are particularized properties, including relations. This claim is not only perceived as counterintuitive by conceptual modeling practitioners (as shown by [8,16]), but it is also controversial from a metaphysical point of view. For instance, Armstrong [1], who as much as Bunge embraces scientific realism as a theory of universals, claims that higher-order properties are *necessary* to represent the concept of a *law*. For Armstrong, a law such as Newton's F = MA describes a second-order relation between the three universals involved. Strangely enough, Bunge also defines the concept of a *Law* (quite a central notion is his approach) as a relation between properties, which makes it a second-order relation [2, p.77]. The view that there are, in fact, material higher-order universals is also shared by other approaches (e.g., [4]). Even simple higher-order relations between universals such as "*Redness is more like Orange than it is like yellow*" cannot be dealt with in the current version of the BWW framework. In contrast, in a trope-based approach, if one wants to dispense with higher-order properties of this kind, this relation can be expressed in terms of first-order inexact resemblance relations between tropes. In fact, in such an approach, traditional properties of properties such as the *hue of a certain color* or the *graveness of a certain symptom* can be modeled in terms of first-order inherence relations between tropes (see fig.4).

If one subscribes to Bunge's theory, however, there is a much stronger reason to argue against the representation of non-substantial universals as types: since Bunge denies the existence of *particularized properties*, one could simply state that properties should not be represented as entity types because they should not be allowed to have instances. However, it is important to emphasize that to accept the claims: (c1) *there are instances of properties*, as well as (c2) *properties can have properties* does not amount to an ontologically incorrect position. The claims (c1) and (c2) are only incompatible with the very specific ontological choices made for the BWW framework. As mentioned above, even if one embraces universalism, (c2) can be accepted. Moreover, the denial of (c1) puts BWW in a singular position among the foundational ontologies developed in the realm of computer science (e.g., [4, 7, 9, 13]). As pointed out by [13], there is solid evidence for (c1) in the literature. On the one hand, in the analysis of the content of perception, particularized properties such as colors, sounds, runs, laughter and singings are the immediate objects of everyday perception. On the other hand, the idea of tropes as *truthmakers* underlies a standard event-based approach to natural language semantics, as initiated by [3] and [11].

3 A Trope-Based Ontology

Figure 2 illustrates the main categories that constitute the ontology proposed in this article. The category of particulars comprises both Objects and Tropes. The relation of inherence is defined between tropes and other particulars, which are not necessarily objects. In other words, we admit that tropes can inhere in other tropes. We also consider the categories of *object kind* and *trope kind* as two possible sorts of *kinds*. We use the term kind here in a broader sense than the term universal, without necessarily committing to the existence of universals, i.e., without choosing *a priori* between position (b) or (c) in figure 1. A kind thus can be considered here simply as something (i) which can be predicated of other entities and (ii) that can potentially be represented in language by *predicative terms*. We also use the relation :: of classification between particulars and kinds. Likewise, classification can be interpreted as instantiation or membership depending on the ontological commitment which is made.

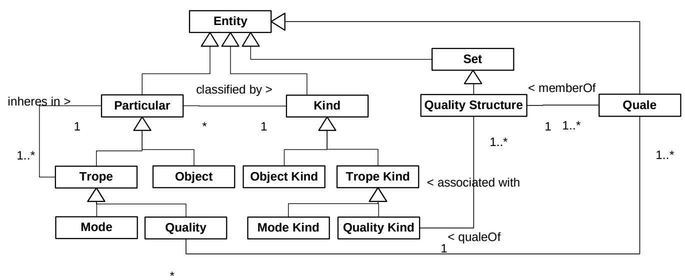

Fig 2. The Categories composing a simple trope-based ontology

Object Kinds classify objects and Trope Kinds classify tropes. Examples of object kinds include Apple, Person and Ferrari 250 GTO. Examples of trope kinds include Color, Electric Charge and Headache. This distinction is also present in Aristotle's original differentiation between what is *said of a subject* (*de subjecto dici*), denoting classification and what is *exemplified in a subject* (*in subjecto est*), denoting inherence. Thus, the linguistic difference between the two meanings of the copula "is" reflects an ontological one. For example, the ontological interpretation of the sentence "Jane is a Woman" is that the object Jane is classified by the object kind Woman. However, when saying that "Jane is tall" or "Jane is laughing" we mean that Jane *exemplifies* the trope kind Tall or Laugh, by virtue of her specific height or laugh.

Here, we capture the intension of a kind by means of an axiomatic specification, i.e., a set of axioms that may involve a number of other kinds representing its essential features. A particular form of such a specification of a kind *U* is called an *elementary specification* (ES). An ES of a kind *U* consists of a number of trope kinds $T_1,\ldots,T_n$ and the inherence relation which attaches instances from the T_i to instances of *U*, expressed by the following schema: $\forall a(a :: K \rightarrow \exists t_1 \ldots t_n \bigwedge_{i \leq n} (t_i :: T_i \wedge i(t_i, a)))$. The relation between a kind U and the trope kinds in its elementary specification is one of *characterization*: A kind U is characterized by a trope kind T iff every instance of K exemplifies T, i.e., iff **∀x (x::U → ∃y y::T ∧ i(y,x))**.

An attempt to model the relation between properties and their representation in human cognitive structures is presented in the theory of *conceptual spaces* introduced in [6]. The theory is based on the notion of *quality dimension*. The idea is that for several perceivable or conceivable trope kinds there is an associated quality dimension in human cognition. For example, height and mass are associated with one-dimensional structures with a zero point isomorphic to the half-line of nonnegative numbers. Other properties such as color and taste are represented by multi-dimensional structures.

Gardenfors [6] distinguishes between *integral* and *separable* quality dimensions: "certain quality dimensions are integral in the sense that one cannot assign an object a value on one dimension without giving it a value on the other. For example, an object cannot be given a hue without giving it a brightness value (…) Dimensions that are not integral are said to be separable, as for example the size and hue dimensions." He then defines a *quality domain* as "a set of integral dimensions that are separable from all other dimensions" and a *conceptual space* as a "collection of one or more domains" (ibid.). Finally, he defends that the notion of conceptual space should be understood literally, i.e., quality domains are endowed with certain geometrical structures (topological or ordering structures) that constrain the relations between its constituting dimensions. In his framework, the perception or conception of a trope can be represented as a point in a quality domain. This point is named here a *quale* [9].

An example of a quality domain is the set of integral dimensions related to color perception. A color quality *c* of an apple *a* takes it value in a three-dimensional color domain constituted of the dimensions hue, saturation and brightness. The geometric structure of this space (the *color splinter* [6]) constrains the relation between some of these dimensions. In particular, saturation and brightness are not totally independent, since the possible variation of saturation decreases as brightness approaches the extreme points of black and white, i.e., for almost black or almost white, there can be very little variation in saturation. A similar constraint could be postulated for the relation between saturation and hue. When saturation is very low, all hues become similarly approximate to grey.

We adopt in this work the term *quality structures* to refer to quality dimensions and quality domains, and we define the formal relation of *association* between quality structure and a trope kind. Additionally, we use the terms *quality kinds* for those trope kinds that are associated with a quality domain, and the term *quality* for a trope classified under a quality kind. We also assume that quality structures are always associated with a unique quality kind, i.e., a quality structure associated with the kind Weight cannot be associated with the kind Color.

Following [9], we take that whenever a quality kind Q is related to a quality domain D, then for every individual quality x::Q there are *indirect qualities* inhering in x for every quality dimension associated with D. For instance, for every particular quality c instance of Color there are quality individuals h, s, b which are instances of quality kinds Hue, Saturation and Brightness, respectively, and that inhere in c. The qualities h, s, b are named *indirect qualities* of c's bearer. Qualities such as h, s, b are named *simple qualities*, i.e., qualities which do not bear other qualities. In contrast, a quality such as c, is named a *complex quality*. Since the qualities of a complex quality x::Q correspond to the quality dimensions of the quality domain associated with Q, then we have that no two distinct qualities inhering a complex quality can be of the same type.

For the same reason, since there are not multidimensional quality dimensions, we have that complex qualities can only bear simple qualities. Moreover, we use predicate *ql(x,y)* to represent the formal relation between a quality individual y and its quale x**.**

Finally, we make a distinction between qualities and another sort of trope named here *modes*. Modes are tropes whose kinds are not directly related to quality structures. Gardenfors [6] makes the following distinction between what he calls *concepts* and *properties*: "Properties…form as special case of concepts. I define this distinction by saying that a *property* is based on *single domain*, while a *concept* may be based on *several domains*". We claim that only trope kinds that are conceptualized w.r.t. a single domain, i.e., quality kinds, correspond to properties in Gardenfors sense. However, there are trope kinds that as much as object kinds can be conceptualized in terms of multiple separable quality dimensions. Examples include beliefs, desires, intentions, perceptions, symptoms, skills, among many others. Like objects, modes can bear other tropes, and each of these tropes can refer to separable quality dimensions. However, since they are tropes, differently from objects, modes are necessarily existentially dependent of some particular.

4 A foundation for Attributes, Weak Entities and Datatypes

Suppose that we have an object kind *Apple* whose elementary specification contains the trope kind *Weight*. Thus, for an instance *a* of Apple there is an instance *w* of the quality kind Weight inhering in *a*, i.e., $\forall$**a (a::Apple** $\rightarrow$ $\exists$**w (w::Weight** $\wedge$ **i(w,a))).**

Associated with the quality kind Weight we have a quality dimension **WeightDim** and, hence, for every instance *w* of Weight there is a quale *c* denoting a particular weight value, i.e., a point in the weight quality dimension such that *ql(c,w)* holds. We take here the weight quality domain to be a one-dimensional structure isomorphic to the half-line of non-negative numbers, which can be represented by a set. The mapping between a substantial *a* and its weight quale can then be represented by the function: **weight: Ext(Apple)** $\rightarrow$ **WeightDim** such that **weight(x) = y |** $\exists$**z z::Weight** $\wedge$ **i(z,x)** $\wedge$ **ql(y,z),** and **Ext(Apple)** represents the set extension of the kind Apple.

In general, let K be a (object or trope) kind and let $Q_{1,}\ldots,Q_n$ be a number of quality kinds. Let E be an elementary specification characterizing the kind U: $\forall$**x (x::U** $\rightarrow$ $\exists q_{1,}\ldots,q_n \bigwedge_{i\leq n}$ **(**q_i**::**Q_i $\wedge$ **i(**q_i**,x)))**. If D_i is a quality domain directly *associated with* Q_i, we can define the function Q_i**: Ext(U)** $\rightarrow$ $\mathbf{D}_i$ (named an *attribute function* for quality universal Q_i) such that for every x::U we have that $\mathbf{Q_i}$**(x)= y | y** $\in$ $\mathbf{D_i}$ $\wedge$ $\exists$**q::**$\mathbf{Q_i}$ $\wedge$ **i(q,x)** $\wedge$ **ql(y,q).**

Let us suppose for now a situation in which every Q_i present in the elementary specification of a kind U is a simple quality kind i.e., Q_i is associated to a one-dimensional quality domain. In this simplest case, the quality kinds appearing in the elementary specification of *U* can be represented in a conceptual model via their corresponding *attribute functions* and associated *quality dimensions* in the following manner: [**<u>Principle 1:</u>** Every attribute function derived from the elementary specification of the kind U may be represented as an attribute of the class C_U (representation of the kind U) in a conceptual model; every *quality dimension* which is the co-domain of one of these functions may be represented as data types of the corresponding attributes in this conceptual model. Finally, relations constraining and informing the geometry of a quality dimension may be represented as constraints in the corresponding data type].

In UML "a data type is a special kind of classifier, similar to a class, whose instances are values (not objects)... A value does not have an identity, so two occurrences of the same value cannot be differentiated" [10, p.95]. A direct representation of Apple's elementary specification in UML according to principle 1 maps the attribute function **weight: Ext(Apple)→WeightDim** to an attribute weight with data type WeightValue in class Apple (figures 3.a-b)**.**

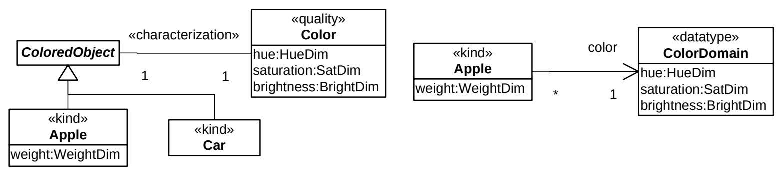

Fig 3.(a, left) - Representing Quality Universals and Indirect Qualities; (b) Representing Qualia in a Multi-Dimensional Quality Domain

Suppose now that we have the following extension of the elementary specification of the kind Apple: **∀a (a::Apple → ∃c∃w (c::Color ∧ i(c,a)) ∧ (w::Weight ∧ i(w,a))).** In order to model the relation between the quality c (color) and its quale, there are other issues to consider. As previously mentioned, the Color quality kind can be associated with a tri-dimensional quality structure composed of quality dimensions hue, saturation and brightness. These dimensions can be considered to be indirect quality kinds exemplified in an apple *a*, i.e., there are quality individuals h, s, b which are instances of quality universals Hue, Saturation and Brightness, respectively, that inhere in the color quality c (which in turn inheres in object *a*). The elementary specification of quality universal Color could then be specified as follows: **∀c (c::Color → ∃h∃s∃b (h::Hue ∧ i(h,c)) ∧ (s::Saturation ∧ i(s,c)) ∧ (b::Brightness ∧ i(b,c))).** In this case, we can derive the following attribute functions from the features in this specification: **hue: Ext(Color) → HueDim**; **saturation: Ext(Color) → SatDim**; **brightness: Ext(Color) → BrightDim**. Together these functions map each quality of a color c to its corresponding quality dimension. One possibility for modeling this situation is a direct application of principle 1 to the Color kind specification. In this alternative, depicted in figure 3.a, the UML class Color directly represents the quality universal color and, its attribute functions hue, saturation and brightness.

Another modeling alternative is to use directly the construct of a data type to represent a quality domain and its constituent quality dimensions (figure 3.b). That is, we can define the quality domain associated with the universal Color as the set **ColorDomain ⊂ HueDim × SatDim × BrightDim**. Then, we can define the following *attribute function* for the object kind Apple: **color: Ext(Apple) → ColorDomain** such that **color(x) = {⟨h,s,b⟩ ∈ ColorDomain | ∃c::Color i(c,x) ∧ (h = hue(c)) ∧ (s = saturation(c)) ∧ (b = brightness(c))}** where hue, saturation and brightness are the attribute functions previously defined. In figure 3.b, we use the UML construct of a *structured datatype* to model the **ColorDomain**. In this representation, the *datatype fields* hue, saturation, brightness are placeholders for the coordinates of each of the (integral) quality dimensions forming the color domain. In this way the "instances" (members) of ColorDomain are quale vectors ⟨x,y,z⟩ where $x \in$ HueDim, $y \in$ SatDim and $z \in$ BrightDim. The *navigable end name* ***color*** in the association between Apple and ColorDomain represents the attribute function ***color*** described above**.**

The two forms of representation exemplified in figures 3.a and 3.b do not convey the same information, which we highlight by the use of different stereotypes. In figure 3.a, color instances are one-sidedly existentially dependent on the particulars they are related to via an *inherence* relation. These instances are genuine individuals with a definite numerical identity. In contrast, the members of the ColorDomain are *pure values* that represent points in a quality domain. These values can qualify a number of different objects but they exist independently of them in the sense that a color tuple is a part of quality domain even if no object "has that color".

Both representations are warranted in the sense that ontologically consistent interpretations can be found in both cases. Notwithstanding, we believe that some guidelines could be anticipated regarding which alternative should be pragmatically more suitable in different cases. In situations in which the tropes of a trope all take their values (qualia) in a single quality domain, the latter alternative (shown in figure 3.b) should be preferred due to its compatibility with the modeling tradition in conceptual modeling and knowledge representation. This is the case with quality kinds. Additionally, since the conceptualization of these tropes depends on the combined appreciation of all their quality dimensions, we claim that they should be mapped in an integral way to a quale vector in the corresponding n-dimensional quality domain.

In the sequel, we observe the following principle between quality domains and their representation in terms of data types: [**Principle 2:** Every quality dimension D directly associated to a quality kind Q may be represented as a datatype DT in a conceptual model; Relations constraining and informing the geometry of a quality dimension D may be represented as operators in the corresponding datatype DT. A collection of integral dimension $D_1 \ldots D_n$ (represented by data types $DT_1 \ldots DT_n$) constituting a quality domain QD can be grouped in structured datatype W representing quality domain QD. In this case, every quality dimension D_i of QD may be represented by a field of W of type DT_i. Moreover, the relations between the dimensions D_i of QD may be represented by constraints relating the fields of data type W].

Principle 2 is a generalization of principle 1 in order to account for quality domains. In summary, every quality kind Q that is associated to a quality domain in an elementary specification of kind U can be represented in a conceptual model via attribute functions mapping instances of U to quale vectors in the n-dimensional domain associated with Q. The n-dimensional domains should be represented in a conceptual model as an n-valued structured data type.

Now, let us consider a case where one of the trope kinds M that characterizes a kind U in its elementary specification is a mode kind. We defend here that these are the cases in which we want to explicitly represent a trope kind in a conceptual model. An example of such a situation is depicted in Figure 4, which models the relation between a Hospital, its Patients, and a number of symptoms reported by these patients. Suppose that an individual patient John is suffering from headache and influenza.

John’s headache and influenza are modes inhering in John. Even if another patient, for example Paul, has a headache that is qualitatively indistinguishable from that of John’s, John’s headache and Paul’s headache are two different particulars. Instances of Symptoms can bear tropes themselves (such as duration and graveness) and can participate in relations of, for example, causation or precedence.

In figure 4, the mode kind Symptom is represented by a class construct decorated with the «mode» stereotype. The formal relation of «characterization» between Symptom and Patient is mapped to the inherence relation in the instance level, representing the existential dependence of a Symptom on a Patient. In other words, for an instance s of Symptom there must be a specific instance p of Patient associated with s, and in every situation that s exists p must exist and the inherence relation between the two must hold. A mode kind such as Symptom in figure 4 can be seen as the ontological counterpart of the concept of *Weak entities types* in EER diagrams, which has been lost in the UML unification process [17].

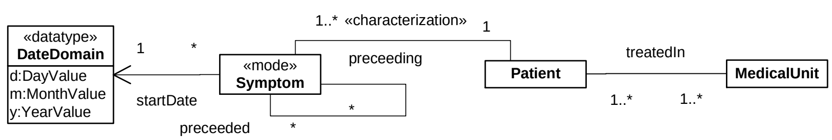

Fig. 4. Representing Object and Mode Kinds and Quality Structures

To summarize this section we can provide the following procedure to represent in conceptual modeling the elementary specification of kinds and their associated trope kinds and quality structures: Take an object kind U with its associated elementary specification E. For every trope kind Q characterizing U do: (1) If Q is a simple quality kind then principle 1 can be applied; (2) If Q is a complex quality kind then principle 2 can be applied; (3) If Q is a mode kind then it should be explicitly represented and should be related to U in a model via a *characterization* relation. Moreover, this procedure can be re-applied to the elementary specification of each trope kind Q in E.

5 A Comparison with the BWW Approach

One of the most defended principles of the BWW approach is the one that states that *“properties cannot have properties”*. So, a question that comes to the mind is: how would one model in the BWW approach situations such as the ones depicted in figures 3 and 4? Take for example, the model in figure 3. In [5], Color is one of the examples used for a property. However, if both Color and Hue (Saturation, Brightness) are properties, how can this conceptualization be modeled in an approach that proscribes the representation of properties of properties?

According to [5], in BWW, the intrinsic (as opposed to relational) properties of a thing *must* be modeled as attributes of the type instantiated by that thing. Since, only substantial types can have attributes, we have that intrinsic properties must be modeled as attributes of substantial types. Thus, a solution to the problem mentioned above is to consider Hue, Saturation and Brightness to be direct properties of Apple, not of Color. The latter, in turn, is then considered to be a conjunction of these three properties, i.e., to instantiate a specific super-determinate shade of red is to instantiate the specific values of Hue, Saturation and Brightness that compose this color. However, in order to be complete, such a solution must also account for the constraints that restrict the possible values that these three dimensions together can assume.

In BWW, a type is represented by a model named a *functional schema*. A functional schema comprises a finite sequence of functions $F = \langle F_1..F_n \rangle$, such that each function F_i (named an *attribute*) represents a property shared by the members of the type described by the functional schema. For every attribute F_i there is a co-domain V_i of values. Bunge defines a function $F(t)$ as the *state function* of the thing, such that $F(t')= \langle F_1(t')..F_n(t') \rangle$ is said to represent the *state of a thing* at time t'. The set $V_1 \times \ldots \times V_n$ is termed the *state space of a thing*. Now, there are certain sorts of types named *Natural Kinds* whose instances have properties which are lawfully related. For these types, it is not the case that the coordinates of the state vectors representing their properties can vary freely. The subset of $V_1 \times \ldots \times V_n$ constrained by the laws of that type being described is named by Bunge the *lawful state space* of a thing. In other words, the lawful state space associated with a natural kind defines all possible states that instances of that kind can assume.

Compared to the approach advocated in this article, we claim that the solution just discussed has two drawbacks. First, as exemplified in figure 3.a, the constraints relating the properties of Hue, Saturation and Brightness are not intrinsic to the type Apple but to the geometry of the Color quality structure and, thus, are reflected in all colored objects. Moreover, these properties form a closure set w.r.t. to mutual dependence and, thus, define a quality domain. In other words, these properties are *integral* and one *"cannot assign an object a value on one dimension without giving it a value on the other"*. For these reasons, we claim that the proposal advanced here of explicit representation of quality domains as datatypes provides the following modeling benefits: (i) a further degree of structuring on lawful state spaces by acknowledging that the co-domains V of attribute functions can also be multidimensional. In fact, this allows for the representation of richer conceptual structures such as the one modeled in figure 5 in which the same Color trope can be measured (take its value) in alternative quality domains; (ii) a structured datatype representing a quality domain can reinforce (via its constructor method) that its tuples will always have values for all its integral dimensions, and only values which obey the constraints imposed by the geometry of that domain; (iii) it also allows for a potential reuse of specifications of multidimensional value co-domains. In this example, once the constraints representing the geometry of the color domain are captured in the specification of the ColorDomain datatype, this specification can be consistently re-used for all colored objects.

The second problem with the solution previously discussed can be defined as follows. If Hue, Saturation and Brightness and its relating constraints are represented in the specification of all types whose instances are colored objects, then by the BWW definition of a natural kind, we can define a natural kind whose instances are all particulars that exemplify the lawfully related properties of Hue, Saturation and Brightness. This allegedly natural kind would be analogous to the type *ColoredObject* depicted in figure 3.a. However, the typical notion of natural kinds in philosophy implies that [7]: (i) they are rigid designators, i.e. that they classify necessarily (in the modal sense) their instances; (ii) that they afford the best inductive generalizations, i.e., that knowing that a particular x is of a kind A also imply knowing that x has all essential properties which are common to instances of A; (iii) that they are associated with a criterion of individuation. These characteristics (i-iii) can all be found in *Apple* but none of them in *ColoredObject*.

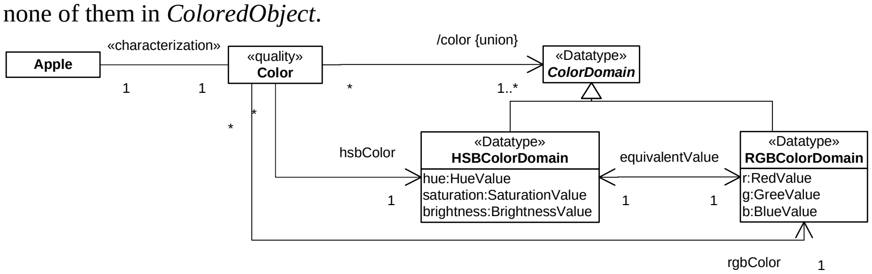

Fig. 5. Explicitly representing quality universals and quality spaces

Let us now consider the case depicted in figure 4. Here, once more, the trope kind symptom can be modeled by having its properties ascribed directly to the type Patient. However, take the property of graveness. "Being grave" is not a property of a particular Patient but a property of a symptom of that patient. Suppose that graveness can be valued in a range 0-5. It is still possible to represent the values in this range as different sets of values of other attributes of symptom, but the introduction of a graveness-space is conceptually clearer. The latter is only possible with a reification of symptom, as illustrated in figure 4 using tropes.

A similar case regards the expression of relations between tropes as the relation of precedence (but also causality) between symptoms depicted in figure 4. According to this model, a symptom such as headache or fever can be caused by another one, for example, influenza. However, differently from the cases mentioned above, these relations cannot be described in general terms; they are indeed relations between instances of these properties. To put it differently, it is John's fever which has been caused by his influenza of a certain graveness. Paul's fever in turn has been caused by his pneumonia.

6 Final Considerations

Despite the perceived usefulness of ontologically well-founded principles and tools for the practice of conceptual modeling, a number of recent results have pointed out the incongruence between what is prescribed by the BWW ontology, on one side, and what is indicated by cognitive and linguistically motivated theories, as well as empirical results of experiments with conceptual modeling practitioners, on the other. The position defended in this paper is in line with some of these criticisms to the BWW ontology. In particular, we reject the BWW-rule that in conceptual modeling only substantial universals should be represented as classes. However, as we have pointed out, it is a mistake to generalize these criticisms to the enterprise of ontology-based conceptual modeling as a whole. As we have shown in the paper, the modeling principles advocated by the BWW framework are a consequence of the very particular type of ontological theory sponsored by its proponents, and their ontological view is only one among many other alternatives.

Furthermore, in this paper we have proposed an alternative ontology which has been used as a foundation for the conceptual modeling primitives of attribute, datatype and weak entities. The ontology presented here is only a fragment of a larger theory which has been extended elsewhere to account for other modeling constructs, such as, classifiers (kinds. roles, phases, mixins), association, part-whole relations, among others [7]. In particular, as demonstrated there, when relational properties are considered, a trope-based approach such as this one not only escapes the criticisms pointed out in [8,15,16], but it also brings a number of additional benefits from a modeling point of view.

References

1. Armstrong, D.M. ‘Universals: An Opinionated Introduction’, Westview Press, 1989.
2. Bunge M.: Treatise on Basic Philosophy. Vol. 3. Ontology I. The Furniture of the World. D. Reidel Publishing, New York, 1977.
3. Davidson, D. ‘The Logical Form of Action Sentences’, Essays on Actions and Events, Oxford University Press, 1980.
4. Degen, W., Heller B., Herre H., Smith, B. GOL: ‘Towards an axiomatized upper level ontology’. Proc. of FOIS’01, Maine, USA, ACM Press, 2001.
5. Evermann J., Wand Y.: Towards ontologically based semantics for UML constructs, Proceedings of ER 2001, pages 354–367. Springer-Verlag, 2001.
6. Gärdenfors, P. ‘Conceptual Spaces: the Geometry of Thought’, MIT Press, USA, 2000.
7. Guizzardi, G. ‘Ontological Foundations for Structural Conceptual Models’, PhD Thesis, University of Twente, The Netherlands, 2005.
8. Hitchman, S. “An interpretive study of how practitioners use entity-relationship modeling in a ternary relationship situation”,Communications of the Association for Information Systems, 2003, 11, 451-485.
9. Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A. ‘Ontology Library’, WonderWeb Deliverable D18, 2003.
10. Object Management Group, ‘UML 2.0 Superstructure Specification’, ptc/03-08-02, 2003.
11. Parsons, T. ‘Events in the Semantics of English: A Study in Subatomic Semantics’. Cambridge/MA: MIT Press, 1990.
12. Russel, B. ‘Human Knowledge, its Scopes and Limits’, Allen and Unwin, 1948.
13. Schneider, L. ‘Formalised Elementary Formal Ontology’, ISIB-CNR Technical Report 03/2002, [online: http://www.loa-cnr.it/Publications.html], 2002.
14. Simons, P. Particular in Particular Clothing: Three Trope theories of Substance’, Philosophy and Phenomenological Research, 54, 553-576, 1994.
15. Veres, C.; Hitchman, S, ‘Using Psychology to Understand Conceptual Modeling’, 10th European Conference on Information Systems (ECIS 2002), Poland.
16. Veres, C.; Mansson, G., ‘Cognition and Modeling: Foundations for Research and Practice’, Journal of Information Technology Theory and Application, v.7, n.1, 93-10, 2005.
17. Vigna, S. ‘ERW: The Manual’, [online: http://erw.dsi.unimi.it/ERW/index.html], 2004.
18. Wand Y.,Storey V.C., Weber R.: An ontological analysis of the relationship construct in conceptual modeling. ACM Transactions on Database Systems, 24(4):494–528, December 1999.