

Lexical Inheritance with Meronymic Relationships

Pablo Gamallo

Received: 23 April 2010 / Accepted: 20 September 2010
© Springer Science+Business Media B.V. 2011

Abstract In most computational ontologies, information inheritance is based on the taxonomic relation *is_a*. A given type inherits from other type only if the latter subsumes the former. We assume, however, that inheritance can be related, not only to the taxonomic relation, but also to the meronymic relationship between parts and wholes. The main aim of this paper is to organise upper-level ontologies associated with lexical information by taking into account part-whole subsumption. As we consider that parts may subsume wholes under specific conditions, ontologies can be defined in terms of systems in which wholes inherit information from its parts. In this article, we describe how part-whole subsumption and, then, meronymic inheritance can be used to deal with type mismatch and metonymic interpretation of polysemous nouns. For this purpose, we attempt to merge old assumptions from both formal ontology and lexical semantics into a homogeneous framework.

Keywords Formal ontology · Mereology · Lexical semantics · Inheritance

1 Introduction

Computational thesauri and ontologies are not simply ordinary repositories of definitions about static words. They also integrate structuring devices such as *inheritance*. The main motivation for inheritance lies in the interest of the economy principle whereby information is stated in the most efficient manner (Ooi 1998). Most of work on inheritance solely makes use of the *is_a* relation, i.e., the *hyponym-hyperonym* or subtyping link, for organising type hierarchies. Simple or multiple inheritance systems are based on the subtyping relation. This relation is used as a

P. Gamallo (✉)
Centro de Investigação em Tecnologias da Informação (CITIUS), Dept. de Língua Espanhola,
Universidade de Santiago de Compostela, 15782 Santiago de Compostela, Galiza, Spain
e-mail: pablo.gamallo@usc.es

channel allowing information to be transferred through the ontological levels. More precisely, this relation is used to build up type hierarchies or taxonomies in such a way that information is transferred from the generic parent-type (the hyperonym) to the more specific daughter-type (the hyponym). This way, the daughter-type inherits information from the parent-type.

When ontological taxonomies are designed to model lexical meaning, it is generally claimed that polysemous items should be represented by means of multiple inheritance. This results in baroque and overloading nets of subtype links at the upper levels of ontological organisation. In such baroque systems, inheritance cannot be easily represented (Asher and Pustejovsky 2000; Pustejovsky 1998; Guarino 1998). To overtake subtyping overloading, this paper describes how information can be naturally transferred by means of other sorts of ontological relations. In particular, the major concern of the paper is to define a particular inheritance system based upon meronymic relations, i.e., *part-whole* relationships. This relies on the following two assumptions:

- A part may *subsume* the whole under specific conditions. That is, particular circumstances may lead a part to be perceived as a schematic image abstracted from the whole object. In this case, the part is represented as a generic scheme of the whole.
- Whereas the basic ontological structure of world entities is mainly organised by subtyping taxonomies, we assume that polysemic lexical information could be mostly structured in part-whole hierarchies. Part-whole subsumption is a suitable organisational device to account for lexical issues such as metonymy, violation of selection restrictions, or regular polysemy. Since part-whole subsumption is the main device used to represent lexical information of polysemic expressions, it is not necessary to build up *is_a* taxonomies as complex multiply-typed hierarchies with multiply-labelled inheritance. Rather, we only require singly-typed hierarchies, which reflect basic ontological relations between classes of entities.

To summarise, this paper presents an information principle organising lexical meaning which is based on the following synthetic idea: as parts subsume wholes, information can be transferred from parts to wholes. To expand and develop such an information principle, the remainder of the paper is structured as follows. Section 2 introduces some cases of violation of selection restrictions, also called type mismatch, as well as the lexical devices (type coercion, predicate transfer) that have been proposed in previous work to solve them. We focus the discussion on some linguistic expressions involving both physical objects and their main internal parts. In Sect. 3, some ontological notions are introduced to describe physical objects in more detail. Section 4 proposes a metonymic principle that could be used to explain type mismatch. This principle leads us in Sect. 5 to define an operation of subsumption based on the part-whole relation. Then, in Sect. 6, we describe a type hierarchy partly organised according to the general principles of meronymic inheritance and part-whole subsumption. Section 7 discusses some issues concerning the degree of complexity of wholes, and the different constraints on meronymic inheritance that should be introduced to account for complex and heterogeneous

wholes. Finally, Sect. 8 studies the internal structure of the horizontal dependencies between parts. Horizontal dependencies will be analysed considering their influence on vertical part-whole relationships.

2 Type Mismatch and Sense Shifting

Many works on lexical semantics are concerned with the analysis of semantic restrictions on linguistic composition, as well as the description of semantic well-formedness for combining lexical expressions. For instance, some research attempts to explain how the semantic type of nouns comes into conflict with the selection restrictions imposed by predicative expressions such as verbs or adjectives. This arises when the type expected by the predicative expression is not compatible with the argument's type. For example, in the expression *listen to the piano*, the verb *listen* requires an argument of type *sound*, whereas *piano* denotes a spatio-material entity, i.e. a physical object noted as *phy*. Likewise, in *drink the bottle*, the verb *drink* expects an argument of type *liquid*, while *bottle* rather denotes an entity perceived as a spatio-material container with the type *phy*.

In order to solve type mismatch (i.e., violation of selection restrictions), two different sorts of assumptions were considered in the past. On the one hand, It is argued that a particular sense shifting of the nominal argument is needed for matching the type required by the predicative expression (Kayser 1987; Pustejovsky 1995; Jayez and Godard 1995). Sense shifting is licensed if we find within the lexical structure of the noun a related semantic information of the correct type. For instance, if *drink* requires a noun of type *liquid* and its argument is the noun *bottle* with the incompatible type *phy*, then it is necessary to coerce the incorrect type of *bottle* into the correct type *liquid*. This coercion is not licensed unless the lexical structure of *bottle* carries ontological information on the type of substance (namely *liquid*) that a bottle is likely to contain.

On the other hand, it can also be argued that a sense transfer should be applied, not on the nominal argument, but on the predicative expression itself (Langacker 1987, 1991; Nunberg 1995). Predicate transfer is succeeded if we have a particular pragmatic or cognitive principle relating the initial predicate to a new predicate requiring the correct selection restrictions. For instance, to solve the type mismatch occurring in *drink the bottle*, a pragmatic function should be applied on the source predicate associated with *drink* (e.g., *drink-liquids*) in order to obtain a new predicate holding the selection restrictions required by *bottle* (e.g., *drink-containers*). Both the source and the derived predicates must be semantically, pragmatically, or cognitively related. The existence of the derived predicate *drink-containers* allows the verb *drink* to be combined with the nominal argument *bottle* referring to a specific container.

Type mismatch is a widespread linguistic phenomenon. It underlies basic metonymic expressions such as *John is parked out back*, *tie up your shoes*, *the bank announced that...*, *put the poem on the table*, *I can't read this sheet of paper*, etc. Nevertheless, it also appears in cases where type mismatch is not so obvious, since

the two types in conflict cannot be conceived as two separate entities having significant and clear boundaries. Let's consider the expression *touch the bottle*. Assuming that *bottle* is a count noun denoting an entity of type *phy*, i.e., a material object having both a specific shape and particular dimensions in the space, then the verb *touch* should be defined as a predicate selecting for nouns of type *phy*. Yet, such a selection restriction does not seem to be appropriate when the direct object of *touch* is filled by uncountable nouns like *butter* or *sand*, which refer to pieces of matter having unspecified shape and indeterminate dimensions in the space: *touch the butter*, *touch the sand*. It follows that *touch* seems to select for two incompatible types:

- a count noun typed as a dimensional physical object (*touch the bottle*)
- an uncountable noun typed as an unbounded matter (*touch the butter*)

This is not a trivial problem. Let's define the verb *touch* as a predicate that requires material entities, i.e., entities of type *matter*. In this case, the expression *touch the bottle* seems to convey some sort of type mismatch since *bottle* does not denote an entity of type *matter*. To solve type mismatch, we may choose between the two alternating solutions outlined above and taken up again here:

1. If the verb *touch* selects for entities of type *matter*, then the argument *bottle*, which denotes an entity of type *phy*, could be coerced into an entity of type *matter*. This is effectively succeeded since the physical container of the bottle is conceptually related to the material constituent it is made of.
2. Assuming that *touch* selects for entities of type *matter*, then we could use a pragmatic function changing the denotation of *touch* into a derived predicate requiring arguments of type *phy*. This way, both the type required by the derived predicate and the type of its argument, *bottle*, are compatible.

Another possible solution could be to use a new complex type, e.g., *matter_phy*, to characterise the selection restrictions of *touch*. Type *matter_phy* could be defined in terms of the entities that are either unbounded material substances or bounded spatio-material objects. So, it represents the disjunction of *matter* and *phy*. In fact, this solution is not formally different from considering *touch* a polysemous verb selecting either for *matter* or for *phy*. This is actually a simplified version of solution 2.

The two possible solutions outlined above will be analysed in more detail in Sect. 4. By now, in the next section, we will introduce some basic assumptions on the ontological structure of physical objects and pieces of matter.

3 Ontological Assumptions upon Physical Objects

In Borgo et al. (1996), physical objects are defined in terms of spatio-material configurations. The authors claimed that a particular physical object, a glass bottle for instance, has not the same *criteria of identity* (or reference) as its internal substrates, namely the chunk of glass it is made of and the space it occupies. First,

they claimed that physical objects must be ontologically distinguished from chunks of matter. When a bottle ceases *to exist* after falling on the ground, the pieces of glass of which it was made are still there. It can be easily seen that the whole bottle and the glass do not refer to the same entity. Second, the authors distinguished between objects and the spatial location they occupy. Indeed, a bottle keeps its identity even if it can be located in different regions of space at different intervals of time. Obviously, the bottle *is not* the spatial region it occupies since the region can be occupied by other objects.

Although a spatial location and a piece of matter cannot be identified with the whole spatio-material configuration, they are related to it by means of dependence relationships: the existence of a physical object necessarily depends on both the existence of a specific piece of matter, and the concrete existence of a region of space. The existence of the whole spatio-material configuration implies the existence of its basic constituents. This way, physical objects are characterised as complex entities holding meronymic relations with their internal parts. In addition to their meronymic structure, physical objects can also be characterised by means of *functional* properties. In particular, physical objects are likely to change their location, i.e. they can move. Movement is represented by a piece of matter occupying successive spatial locations. In other words, movement is a functional property that depends on the complex meronymic structure of physical objects.

Nevertheless, the ontological characterisation of physical objects described by Borgo et al. (1996) still remains very schematic. It does not account for other functional properties than movement. For instance, physical objects are likely to be *visualised*, *touched*, *measured*, etc. Unlike movement, which is a functionality linked to the whole spatio-material configuration, functions of touching, visualising or measuring are directly linked, not to the whole configuration, but to its specific components: both matter and space. Intuitively, when we touch a bottle, we actually touch the piece of glass it is made of and not the spatial region it occupies. Hence, the action of touching seems to be directly linked to one of the meronymic parts of the whole bottle. Likewise, when we measure the size of the bottle, we actually measure the dimensional space that the bottle is occupying. So, the space becomes the more salient part with regard to the action of measuring.

Even though the properties of being touched and being measured are directly associated with the two constituents, respectively matter and space, the whole bottle is somehow touched and measured as well. We will argue later that the whole spatio-material configuration inherits these functional properties from the two parts: matter and space. In the remainder of the paper, inheritance from parts will be used to solve type mismatch by means of metonymic interpretation.

4 Metonymy and Interpretation of Complex Expressions

In Sect. 2, two proposals were made in order to solve type mismatch: argument coercion and predicate transfer. Kleiber (1994, 1999) rejected both solutions.

On the one hand and according to Nunberg, Kleiber claims that argument coercion makes it difficult to explain the anaphoric relations between the coerced noun and any coreferential pronoun:

(1) John is parked out back. He's waiting for his wife.

In accordance with the assumption on argument coercion, the verb parked should coerce John to be an entity of type *vehicle*. That is, the semantic requirements of the verb would lead the proper name¹ John to denote the car he is driving. As a consequence, further anaphoric pronouns should take the vehicle, and not the individual John, as the antecedent. Nevertheless, in utterance (1), the anaphoric pronoun he is related, not to the car, but to the individual denoted by John. It follows that John keeps its natural referent (i.e., it is still referring to a human individual) and, then, its basic identity criteria, even though the predicate is parked requires the type *vehicle*. Coercion from *individual* to *vehicle* does not seem to be possible since there is no change of referent.

On the other hand, Kleiber also rejected the meaning transfer of the predicate. First, he claimed that the construction of a derived predicate runs counter to intuition and common-sense because this would require an uncontrolled multiplication of the types associated with verbs. Second, he argued that the pragmatic or cognitive relation between the two predicates (source and derived predicates) presupposes (or is not independent from) the metonymic relation between the two referents associated with the noun. In example (1), the denotation of parked should be transferred from the source predicate *park-vehicle*, which requires objects of type *vehicle*, into the derived predicate *park-individual*, which requires objects of type *individual*. Such a transfer is only possible if there is a pragmatic relation between both *park-vehicle* and *park-individual*. Yet, this relation is not independent from the more basic relation between the owner and the car. It follows that the meaning transfer of the predicate relies on the metonymic links organising the semantic content of the noun argument. Thus, predicate transfer cannot be explained without previously analysing the semantic content of the argument.

According to Kleiber, neither argument coercion nor predicate transfer is needed to solve type mismatch. He assumes that it is a particular internal constituent of the entity denoted by the noun that becomes salient and representative with regard to the whole entity, under specific conditions. So, the noun argument keeps its identity criteria and its by default type, while the predicate is not derived into a new predicate by pragmatic rules. He calls *metonymic interpretation* the principle asserting that certain parts become salient with regard to the whole under specific conditions. Let's take again the previous example: touch the bottle. Under the conditions concerning the action of touching, the material facet of the bottle becomes the salient part of the entire bottle, and, as such, it may characterise the whole. In other words, the material part of the bottle represents and characterises the complex spatio-material configuration under the conditions imposed by touch.

Considering touch the bottle as a metonymic expression is in accordance with the fact that metonymic interpretations are no longer special cases, but rather standard

¹ To simplify, nouns, proper names and nominal phrases are semantically interpreted in the same way.

and usual ones. Indeed, the interpretation of complex expressions often involves a metonymic inferential device operating on the noun meaning. From this viewpoint, non-metonymic interpretation represents the limit case since the combination of a noun with a predicative expression (verb, adjective, etc.) is commonly interpreted by means of nominal metonymy. Metonymic interpretation is thus at the centre of the the meaning construction process.

Let's take now the expression *small bottle*. If we consider that the adjective *small* merely selects for spatio-material configurations (i.e., physical objects), we hardly account for those cases where the adjective only require spatial or temporal entities such as holes, tunnels, openings, distances, ways, geometrical figures, wells, volumes, etc. These immaterial entities have specific shapes and bounded dimensions regardless of their material facet. They are not spatio-material configurations. Indeed, to account for expressions such as *small hole*, *small line*, *small tunnel*, etc., we need to define *small* as a predicate requiring spatial objects. the adjective *small* seems to directly measure the *space* without presupposing any specific matter, in the same way as the verb *touch* seems to be directly applied to *matter* without presupposing a specific dimensional space. In the expression *small bottle* and under the specific conditions imposed by the adjective, it is, then, the *spatial* facet of the bottle that represents and characterises the overall spatio-material configuration. According to the principle of metonymic interpretation, the restrictions imposed by a dimensional adjective like *small* require the noun *bottle* to give direct access to the type *space*. Under these conditions, the spatial constituent becomes the more salient part of the spatio-material configuration denoted by *bottle*.

Furthermore, it is also easy to find predicates imposing a spatio-material configuration (i.e., type *phy*). For example, in the expression *put the bottle on. . .*, the movement verb *put on* selects for arguments of type *phy*. As has been argued in the previous sections, it is necessary to presuppose both a piece of matter and a set of spatial locations in order to conceptualise physical movement and shift location. Consequently, under the conditions required by this movement verb, it is the overall spatio-material configuration that characterises and represents the whole bottle. It is the limit case: the whole represents the whole.

The metonymic principle lets a nominal expression be compatible with apparently odd selection restrictions, without modifying its identity criteria. In the rest of the paper, our main concern will be to associate this principle with the lexical subsumption based upon the *part-whole* meronymic relation. Meronymic subsumption will be used later to explain type mismatch.

5 Two Sorts of Type Subsumption

5.1 Subtype subsumption

The simplest case of type mismatch is encountered when the predicate requires a more generic type than that directly associated to the argument. Let's assume that *drive* selects for an object with type *vehicle*, while *car* denotes an object having the

specific type *car*. Obviously, since *vehicle* is not equal to *car*, the expression *drive a car* does not satisfy the selection requirements.

The conventional relation between the two mismatched types is one of subtyping (i.e., the taxonomic relation *is_a*). As formal properties accompanying types and subtypes are well known, this case of type mismatch is not considered as being problematic. We only need to apply the subsumption operation.

$$\frac{c \vdash a : t1 \quad \Theta[t1 \leq t2] : t1 \rightarrow t2}{c \vdash \Theta[t1 \leq t2](a) : t2} \tag{1}$$

In Eq. 1, the expression *a* is coerced to a syntactically identical expression with a more general type. We assume a type assignment, $c \vdash a : t1$, with respect to a grammar, such that under assignment *c*, the expression *a* has type *t1* (Pustejovsky 2001). Equation 1 says that, given an item *a* of type *t1*, which is a taxonomic subtype of *t2*, it is possible to assert that *a* is of type *t2*. The subtyping subsumption Θ is a function from *t1* to *t2* since *t1* is a subtype of *t2* ($t1 \leq t2$). So, given that *car* is a subtype of *vehicle*, the item *car*, whose type is *car*, is coerced to be an expression of type *vehicle* (see Eq. 2). This allows *car* to fit the selectional requirements of the verb *drive*. In other words, under the conditions imposing by *drive*, the *car* is coerced to be perceived as a *vehicle*.

$$\frac{c \vdash car : car \quad \Theta[car \leq vehicle] : car \rightarrow vehicle}{c \vdash \Theta[car \leq vehicle](car) : vehicle} \tag{2}$$

Type shifting is here licensed by the basic operation of subsumption. Yet, despite of this type coercion, there is no change of referent (i.e., denotation). Type *vehicle* is able to represent the same entity as *car*, by abstracting the specific information that differentiates cars from other vehicles. Since both types, *vehicle* and *car*, may refer to the same entity, the subsumption operation typically does not modify the identity criteria of the entity denoted by the lexical item. Rather, it modifies the level of specification from which the entity is observed.

The subsumption operation described above is similar, in some respects, to “type coercion”, defined by Pustejovsky. Type coercion is used to solve the selectional mismatching caused by various sorts of metonymic inferences. Type coercion states that, given an item *a* of type *t1*, it is possible to assert that *a* is of type *t2* if there is a specific association between *t1* and *t2* within the qualia structure of *a*. The main difference between the two operations is the following: coercion is licensed by a quale-based association relation (“agentive” or “telic”), while subsumption is licensed by the well known subtyping relation (called “formal” by Pustejovsky). Moreover, whereas subtyping subsumption operates on a singled-typed lattice, qualia-based association requires a multiply-typed lattice, making use of information available through the internal qualia structure associated with the lexical item *a*. It replaces standard inheritance by inheritance *via* qualia labels. Finally, since qualia-based associations involve types characterised by different identity criteria (individuals and events), type coercion also shifts the basic denotation of the lexical item. In the following section, we will define a more generic operation than type coercion, called “meronymic subsumption”, which involves, among others,

constitutive, agentive, and telic relations, but which does not change the identity criteria (or denotational type) of the “coerced” expression.

5.2 Meronymic Subsumption

Both subtyping and meronymy relations have similar linguistic uses and formal properties. As has been said above, type *vehicle* is able to represent the same entity as *car*, by abstracting the specific information that differentiates cars from other vehicles. Likewise, a part of a car, for instance the *coachwork*, can also be used to represent the same entity as *car*, by abstracting the other specific parts. Consider the expression *red car*. The adjective *red* does not require the car to be perceived as a whole constituted by a motor, four wheels, a coachwork, a steering wheel, etc. The car needs only to have a red coachwork to be considered as a red car. For the conditions imposed by the adjective *red*, the only relevant part of the car is the coachwork. The specific colour of the other parts (wheels, motor, seats, ...) is then not relevant. It follows that the coachwork represents the whole car under the specific conditions imposed by the adjective *red* (see Gamallo1998). The way in which the part *coachwork* subsumes the specific whole *car* is similar to the subsumption between the generic type *vehicle* and the specific type *car*.

Moreover, the subtype relation and the meronymic part-whole relation can be defined in terms of the same logical properties. Like subtyping, part-whole meronymy is a partial ordering (i.e., reflexive, antisymmetric, and transitive) relationship (Varzi 1996; Smith 1998). Since meronymy shares the same formal properties than subtyping, we assume that it may also serve to define a special case of subsumption: *meronymic subsumption*. In Eq. 3, the expression *a* is coerced to have the type of one part. It says that, given an item *a* of type *t1*, which has a part with type *t2*, it is possible to assert that *a* is of type *t2* (i.e., under assignment *c*, the expression *a* has type *t2*). Meronymic subsumption Φ is a function from *t1* to *t2* since *t2* is a part of *t1* ($t2 \preceq t1$).

$$\frac{c \vdash a : t1 \quad \Phi[t2 \preceq t1] : t1 \rightarrow t2}{c \vdash \Phi[t2 \preceq t1](a) : t2} \tag{3}$$

Consider again the expression *touch the bottle*. Let’s assume that the type selected by *touch* is *matter*, while *bottle* inherits by subtyping the taxonomic type *phy*. Then, the selectional requirements can be satisfied just in case there could be a whole-part relation between *phy* and *matter*. Given that *matter* is an ontological part of *phy*, the item *bottle*, whose type is *phy*, is coerced to be an expression of type *matter*. Then, *bottle* matches the selectional requirements of the verb *touch*. Type shifting is here licensed by the operation of meronymic subsumption. Equation 4 says that, given the item *bottle* of type *phy*, which has a part of type *matter*, there is a subsumption between *phy* and *matter*, which changes the type of *bottle* from *phy* to *matter*.

$$\frac{c \vdash \text{bottle} : phy \quad \Phi[\text{matter} \preceq phy] : phy \rightarrow matter}{c \vdash \Phi[\text{matter} \preceq phy](\text{bottle}) : matter} \tag{4}$$

Meronymic subsumption allows us to formalise the “metonymic interpretation” principle that we have outlined in Sect. 4. Meronymic subsumption leads *matter* to

be the salient part of the spatio-material object denoted by *bottle*, under the conditions of touching. In other words, when the bottle is touched, the material it is made of becomes the more representative facet of the whole object. This notion of meronymic subsumption is consistent with the observation that metonymic inferences do not modify the identity criteria of the semantic entities.

Notice that meronymic subsumption can be applied, not only to prototypical part-whole relations involving constitutive information, but also to telic, agentive, or further functional relations, since all of them can be perceived as particular cases of part-whole associations.

6 Combining Subtyping and Meronymic Relations in a Type Hierarchy

We argue that subtyping mismatch and metonymic mismatch can be described in terms of an abstract subsumption operation, which is based on both the hyperonym-hyponym and part-whole relationships. Both relations are used here to build type hierarchies. According to the formal properties of these relations, type hierarchies are organised as simply-typed lattices.

6.1 Relation of Information Containment

In order to combine subtyping and meronymic relationships in a type lattice, we assume that lexical types are semantic information structures ordered by the abstract relation of information containment, noted \sqsubseteq . If $t1 \sqsubseteq t2$ we say that $t1$ contains at least the same information than $t2$, that is, every piece of information in $t1$ is contained in $t2$; or conversely, $t2$ is at least as specific as $t1$ (Paiva 1993; Carpenter 1992). This relation holds only between consistent types. So, if $t1 \sqsubseteq t2$ then $t1$ subsumes the information contained in $t2$. Since we do not describe how information associated with types could be represented, a formal description of consistent types cannot be made here. This way, work on different modes of information representation such as typed structure features, postulates of interpretation, conceptual graphs, or qualia structures remains beyond the scope of the paper. We merely assume that each type is associated with a structure of semantic information, regardless of its particular organisation.

A type hierarchy can be defined as an algebra $\langle T, \sqsubseteq \rangle$, where T represents the set of types, and \sqsubseteq the informative containment relation, which is defined as a partial ordering. It is by means of such a relation that the inheritance device enables information to be transferred through the type hierarchy. Let's take the following types:

$$bottle, matter, phy \in T$$

According to ontological requirements, we may organise these three types into a hierarchy where:

$$\begin{aligned} phy &\sqsubseteq bottle \\ matter &\sqsubseteq bottle \end{aligned}$$

Type *bottle* is more specific than both its hyperonym *phy* and its internal part *matter*. So, *phy* and *matter* are somehow generic supertypes subsuming *bottle*. Given that both the hyperonym *phy* and the part *matter* can be considered as two abstract schemes subsuming the more specific type *bottle*, they represent two sources from which *bottle* inherits semantic information.

$$\frac{c \vdash a : t1 \quad \otimes [t2 \sqsubseteq t1] : t1 \rightarrow t2}{c \vdash \otimes [t2 \sqsubseteq t1](a) : t2} \tag{5}$$

The relation of information containment gives rise to a very general notion of coercion: “informative subsumption”. Equation 5 displays type coercion by informative subsumption. Given an item *a* of type *t1*, which is more informative than *t2*, it is possible to assert that *a* is of type *t2*. Informative subsumption \otimes is a function from *t1* to *t2* where *t2* is somehow contained in *t1* (either as a hyperonym or as a part).

6.2 Hyperonyms and Wholes

Intuitively, subtyping and meronymic relationships are two specific instantiations of the relation of information containment \sqsubseteq . Specific subtypes are more informative than their hyperonyms, as well as wholes are more informative than their parts. Wholes and hyperonyms are obtained by applying two different operations on types: unification and generalisation, respectively. Both operations are naturally defined in terms of the relation of information containment.

The unification of two types is defined to be the most general type which contains all the information provided by both types. Each constituent type subsumes the type obtained by unification. For instance, the unification *matter* \sqcup *space* is well-defined if there is a type *phy* such that *matter* \sqsubseteq *phy* and *space* \sqsubseteq *phy*. So, the whole type *phy* results from unifying its parts, i.e., *matter* \sqcup *space* = *phy*. Unification can be applied to a not restrictive number of parts.

The generalisation of two types is defined in terms of the most specific type that contains only information found in both types. The type obtained by generalisation subsumes the two combined types. For instance, the generalisation *natural_kind* \sqcap *artefact* is well-defined if there is a type *phy* such that *phy* \sqsubseteq *natural_kind* and *phy* \sqsubseteq *artefact*. Thus, type *phy* is the generalisation of its immediate hyponyms: *natural_kind* and *artefact*. However, there is not a restrictive number of hyponyms that can be merged by generalisation. In fact, type *phy* also results from applying the generalisation operation to specific physical objects such that *bottle*, *car*, *plant*, etc.

A given type may be obtained from both unification and generalisation. That is, a type may represent both a whole and a hyperonym. For example, type *phy* consists in combining its parts and its hyponyms (i.e., its subtypes):

- *phy* is the result of integrating by unification its mereological components:
phy = *matter* \sqcup *space*

- *phy* is the result of extracting the shared information from its specific hyponyms (*bottle*, *car*, *plant*, etc), by means of the generalisation operation:
 $phy = bottle \sqcap car \sqcap plant \sqcap \dots$

It follows that $t1 \sqsubseteq (t1 \sqcup t2)$ represents a meronymic relation, whereas $(t1 \sqcap t2) \sqsubseteq t1$ represents a taxonomic relation. Hierarchies joining types by means of unification are perceived as “meronymies”. Meronymies are ordered by the parthood relation, noted \sqsubseteq_{mer} . Hierarchies organising types by means of generalisation are perceived as “taxonomies”. Taxonomies are ordered by the subtyping relation, noted \sqsubseteq_{tax} . Meronymies and taxonomies are merely two complementary modes of combining and relating types in $\langle T, \sqsubseteq \rangle$.

6.3 Taxonomic and Meronymic Inheritance

Typically, lexical thesauri and ontologies are mainly organised on the basis of taxonomic hierarchies. In taxonomies, information inheritance solely depends on the subtype relationship *is_a*. Further ontological relations, such as whole-part relationships, do not play any role in inheritance because they are not associated with the formal properties enabling the transfer of information throughout the hierarchy. We assume, however, that the whole-part relation must be conceived as an internal part of the organising skeleton of upper-level ontologies. Given that parts can be defined in terms of schematic concepts subsuming wholes under specific conditions, we argue that wholes inherit information from their parts. Such an assumption is in accordance with the metonymic principle described in the previous section.

Let’s consider again the description of spatio-material configurations (i.e., physical objects). Physical objects may be ontologically characterised by means of functional properties like *being moved*, *being used in a certain way*, *being created*, etc. These properties represent the informative content associated with type *phy*, regardless of the formal language (feature structure, conceptual graphs, etc.) used to represent it. Further elaborate descriptions concerning these properties could lead us to characterise subtypes of *phy* at various levels of schematicity. Indeed, the *being used* property could allow us to perceive physical objects as tools, pieces of decorations, etc. Likewise, specific information concerning the *being created* functionality would enable us to distinguish artefacts from natural kinds. The two basic constituents of physical objects, *matter* and *space*, also possess characteristic functions. Pieces of matter are characterised by means of functional properties such as *being touched*, *being stained*, *getting wet*, *having a specific colour*, etc. Dimensional spaces are constituted by properties and functions like *being measured*, *having a specific length*, *being designed*, *being passed through*, etc. As *phy* is subsumed by its parts, it inherits the functional information characterising pieces of matter and bounded spaces. So, physical objects, not only can be moved, but also can be touched, measured, etc. Even though *phy* inherits all the information accessible from its parts, it shouldn’t inherit their identity criteria. Indeed, a physical object is not identified as the matter it is made of, nor the bounded space it occupies. That is, physical objects do not denote unbounded chunks of matter nor immaterial

spaces. Rather, the clue for the identification of a physical object could be the existence of a particular relation between a piece of matter and the spatial dimension that it occupies.

Type *phy* is also related to its more specific subtypes: *bottle*, *car*, *plant*, etc. Given that subtypes are subsumed by the hyperonym *phy*, they inherit all the information characterising physical objects (as well as pieces of matters and bounded spaces). Figure 1 illustrates the hierarchical path from specific physical objects at the bottom of the hierarchy, to their schematic constituents at the top. Figure 1 can be described as follows:

- From the taxonomic viewpoint, *phy* is perceived as the generalisation of the elaborate types *bottle*, *car*, *plant*...:

$$phy = bottle \sqcup car \sqcup plant \sqcup \dots$$


That is to say, *phy* represents the type containing the information shared by its subtypes. In Fig. 1, thin lines stand for taxonomic relationships.

- From the meronymic viewpoint, *phy* is a spatio-material entity, i.e., the type obtained by unifying both *space* and *matter*: $phy = space \sqcup matter$. Thick down-up arrows represent meronymic relationships.

The level of elaboration concerning the internal components of physical objects depends on the particular applications, designs, and descriptive goals of the ontology. If the application would require deeper knowledge about physical objects, we should take into account finer-grained elaborations of the *phy* components: colour, texture, shape, etc.

To summarise, Fig. 1 illustrates a type hierarchy whose inheritance device is activated not only by subtyping subsumption but also by whole-part relations. Indeed, the *bottle* type inherits information not only from its hyperonym *phy*, but also from its internal constituents *matter* and *space*. Since subtype relations only connect types sharing similar identity criteria, taxonomies are not any longer multiply-typed and multiply-labelled hierarchies resulting in an overloading of subtype links (Guarino 1998). In our model, taxonomies are simplified since we only use meronymies to account for lexical phenomena such as selection restrictions, metonymic interpretation and polysemy. While part-whole relations are mainly used to organise lexical information, taxonomies are merely used to structure the ontological relations between entities of the world.

Fig. 1 Type hierarchy of *phy* (physical object)



6.4 Type Mismatch and Metonymic Interpretation Principle

The interpretation principle defined in Sect. 4 asserts that a part or a partial aspect of an entity may represent the whole entity under specific conditions. Whole-part subsumption and meronymic inheritance combine with subtyping subsumption and taxonomic inheritance so as to make the metonymic principle operational. Let’s analyse again the expression *small hole*. Whereas the adjective *small* selects for dimensional and bounded spaces, i.e., objects with type *space*, *hole* is a noun of type *hole*. Then, the interpretation of that expression requires the subtyping subsumption operation illustrated in Eq. 6. The item *hole* is coerced to be an expression with the more generic type *space*. More precisely, Eq. 6 states the following: given the noun *hole* of type *hole*, which is a subtype of *space*, it is possible to apply on it the informative subsumption operation \otimes , which changes the type of *hole* from *hole* to *space*.²

$$\frac{c \vdash \text{hole} : \text{hole} \quad \otimes [\text{hole} \sqsubseteq_{\text{space}}] : \text{hole} \rightarrow \text{space}}{c \vdash \otimes [\text{hole} \sqsubseteq_{\text{space}}](\text{hole}) : \text{space}} \quad (6)$$

Let’s take now the expression *small bottle*. Assuming that *bottle* has the specific type *bottle*, which is a subtype of *phy*, then the interpretation of that expression needs for both subtyping and meronymic subsumptions. In Eq. 7, the noun *bottle* is coerced to a noun with the more generic type *phy*, which is coerced in turn to a noun with the type of one of its internal constituents: *space*. Given that *bottle* is a subtype of *phy*, and *space* a part of *phy*, then the noun *bottle* is coerced twice by the informative subsumption operation \otimes . First, it is coerced by subtype subsumption to be an item of type *phy*. Second, the item of type *phy* obtained in the previous operation is coerced by meronymic subsumption to be a noun of type *space*.

$$\frac{c \vdash \text{bottle} : \text{bottle} \quad \otimes [\text{bottle} \sqsubseteq_{\text{tax}} \text{phy}] : \text{bottle} \rightarrow \text{phy} \quad \otimes [\text{phy} \sqsubseteq_{\text{mer}} \text{space}] : \text{phy} \rightarrow \text{space}}{c \vdash \otimes [\text{phy} \sqsubseteq_{\text{mer}} \text{space}](\otimes [\text{bottle} \sqsubseteq_{\text{tax}} \text{phy}](\text{bottle})) : \text{space}} \quad (7)$$

The recurrent application of the two subsumption operations allows us to account for the very subtle sense shift that the noun *bottle* undergoes in the context of the expression *small bottle*. The shift in question is not an entire denotational transformation where the noun *bottle* would change its criteria of identity. It is not a referential change from a spatio-material object into an immaterial space. The referential entity denoted by *bottle* keeps its identity criteria: it is still identified as a spatio-material configuration from the taxonomic viewpoint (i.e., $\text{bottle} \sqsubseteq_{\text{tax}} \text{phy}$). Yet, the whole configuration is perceived from one of its internal facets—*space*—under the conditions required by the dimensional adjective *small*. So the dimensional conditions make the spatial part salient with regard to the whole *bottle*.

² In research on formal ontology, holes are not considered as mere spatial locations but as physical objects which never have matter associated in any location; they are not spaces but rather *immaterial* objects. For the general purpose of this article, the ontological difference between bounded spaces and immaterial objects will not be taken into account.

Likewise, the verb *touch* requires the direct object to be a material entity, i.e., an entity of type *matter*. In our ontology, type *bottle* fill such conditions because it can inherit material properties from its constituent *matter*, through its hyperonym *phy*:

$$bottle \sqsubseteq_{tax} phy \sqsubseteq_{mer} matter$$

There is no need for constraining physical objects like bottles to change into entities identified as unbounded material substances. Type *bottle* is not a subtype of *matter*; rather it is identified as a spatio-material object likely to be described in terms of its material constituent under specific conditions. More precisely, the action of touching makes the part *matter* salient with regard to the whole bottle.

In the following sections, we describe how meronymic inheritance could be used to deal with other cases of metonymy, in particular those cases related to more complex and heterogeneous types.

7 Complex Types and Constraints on Meronymic Inheritance

Types are unified at various levels of organisation giving rise to “complex types”. Complex types represent polymorphic and heterogeneous objects, such as those obtained by unifying a physical support with a symbolic content, a container with the substance it contains, or a human organisation with the product it makes. Nevertheless, the parts building complex types do not always have the same degree of salience and representativeness with regard to the whole (Miéville 1998). We argue that meronymic inheritance must be modified and adjusted to a certain extent by taking into account prominence and salience of parts. To be more precise, not salient parts constrain the meronymic inheritance device in such a way that they do not transfer to the complex type all the information they contain. A complex type inherits from its not salient parts only the information that can be pertinent to salient parts (Gamallo 2000a, b). Let’s us analyse the expressions illustrated in Table 1.

Take the word *poem*. It usually matches the semantic conditions requiring symbolic entities, such as in *interesting poem*, or *learn the poem*. Nevertheless, according to the expressions 1a–c, *poem* does not refer only to a symbolic content, but to the physical object holding the symbolic information. Indeed, the verbs *take*, *burn*, and *put* require nominal phrases denoting physical entities. Since the symbolic content and the physical holder are entities ontologically related by an underspecified *container-content* link, we assume that both entities can be unified into the whole type $phy \sqcup symb$. However, *poem* does not give access to the two parts of this polymorphic type in the same way. Oddness of expressions 1d–f shows that the physical facet (i.e., *space* and *matter*) is not entirely accessible from *poem*. Given that this noun hardly matches the specific physical conditions imposed by *touch*, *blue*, or *rectangular*, the physical holder appears to be less salient than the symbolic part. So, noun *poem* makes *symb* more salient than *phy*. This construal asymmetry constrains the meronymic inheritance device in such a way that there is some information directly associated with the physical part that cannot be inherited by the complex type. This way, the physico-symbolic type does not inherit those physical aspects that are ontologically independent of the symbolic facet: for instance, the

Table 1 Metonymic expressions

Correct metonymies	Odd metonymies
(1)	(1)
a. Take the <i>poem</i>	d. Touch the <i>poem</i>
b. Burn the <i>poem</i>	e. Blue <i>poem</i>
c. Put the <i>poem</i> on the table	f. Rectangular <i>poem</i>
(2)	(2)
a. Take the <i>wine</i>	c. Glass <i>wine</i>
b. Put the <i>wine</i> on the table	d. Small <i>wine</i>
(3)	(3)
a. <i>Newspaper</i> on strike	c. The <i>newspaper</i> is drinking beer
b. The <i>newspaper</i> announced that. . .	d. Tall <i>newspaper</i>
(4)	(4)
a. The <i>lawyer</i> is parked out back	c. The <i>lawyer</i> is blue
b. The <i>lawyer</i> is at the top of the list	d. The <i>lawyer</i> has 5 letters

colour or the shape of the holder are not related to the symbolic content of a poem. The whole type only inherits those physical properties or functions, e.g., movement or destruction, that are somehow related to the symbolic part. If someone changes the location of the physical holder, then the symbolic content is also being moved (take the poem). Likewise, in someone destroys the holder, then its symbolic content is no more available (burn the poem).

Let's analyse now a noun typically referring to a particular substance: wine. According to 2a–b, this noun is not only able to denote a specific unbounded substance, but also the physical container. As in the physico-symbolic configuration defined above, the substance and its physical container are ontologically related and then unified into a complex type. Yet, both constitutive entities are not accessible from wine in the same way. Whereas the properties of the substance seem to be entirely accessible, specific aspects of the physical container remain out of the semantic scope of the noun: e.g., material and size of the container (2c–d). This information is not inherited by the whole type (and then is not accessible from the noun) because it is not directly dependent of the properties of the more salient part, the substance.

Take the more complex noun newspaper. Apart from denoting an entity of type $phy \sqcup symb$ (i.e., a physico-symbolic entity), this noun is also able to refer to a human staff under certain conditions (see 3a–b). Given that the physico-symbolic entity and the human staff are ontological related by a *product-producer* link, it can be assumed that both types are unified into a more complex type where the product is instantiated by the physico-symbolic object and the producer by the staff. Yet, the noun newspaper does not give an entirely access to all properties of the producer, which is the not salient part. Indeed, specific information on the staff remains inaccessible from the noun, e.g., physical properties of the staff's members, or any type of activity that does not relies on the production process:

drinking beer, smoking, etc. (see 3c–d). These properties remain inaccessible because they are independent of the more salient part, the physico-symbolic product. Thus, they are not inherited by the whole type. This type only inherits from the human producer those properties that are relevant for the product and the production activity.

Finally, let's analyse the complex use of nouns and pronouns denoting human beings, e.g., father, friend, you, I, president, lawyer, etc. These items do not only refer to individuals perceived as biological entities (the *lawyer* is eating eggs with ham), but also to inanimate objects somehow related to them: cars, administrative names, etc. (see 4a–b). Indeed, human beings and specific objects such as functional tools or cultural artefacts can be unified in order to build complex social beings. From the meronymic viewpoint, the biological facet of a human being is only a part of the complex social configurations used to characterise the polymorphic meaning of human nouns like *lawyer*. However, whereas the properties of the biological facet are all accessible from the noun (the *lawyer* is eating, the *lawyer* is sick, etc.), there are significant restrictions to access to the properties of the related social artefacts: cars, names, etc. For instance neither the colour of the car, nor the size of the name are accessible from *lawyer* (4c–d). Except under very particular conditions, these specific properties do not seem to modify the main properties of the salient part: the biological human being. So, they are hardly inherited by the whole type denoting the social being. This type only inherits from the social artefacts (the car and the name, for instance) those features that somehow modify the biological individual: the localisation of his car determines his own localisation (4a), and the place of his name on the list may be related to his professional prestige (4b).

The more complex the whole, the more asymmetric the dependence between the immediate parts. Asymmetry constrains the whole to filter only some of the properties of not salient parts. Access to (or inheritance from) not salient parts seems to be constrained by the following ontological principle (Gamallo 2000b):

Complex wholes only inherit from their not salient parts those properties that can modify the properties of their salient parts.

This principle will be analysed in the following section. For this purpose, the fuzzy notion of *salient part* will be defined considering some concepts and ideas of Formal Ontology.

8 Theory of Dependence and Complex Types

The notion of *salience* is concerned with the horizontal dependencies between co-existing parts. The specific properties of these dependencies determine the degree of integrity and unity of the whole: the more interdependent the parts are, the more homogeneous and compact the whole is (Smith 1998). We assume that issues concerning dependencies between parts and degree of integrity of wholes serve to define constraints on meronymic inheritance. Let's analyse the whole integrity of types such as *space* \sqsubset *matter* and *phy* \sqsubset *symb*.

8.1 Integrity of Spatio-Material Configurations

It seems to be obvious that the type *space* \sqcup *matter* (i.e., *phy*) is constituted by strongly interdependent parts. Indeed, *space* and *matter* cannot be conceived separately if they are integrated into this more complex type. They are related by a *symmetric dependence*. The symmetric dependence between two parts leads us to infer the two following properties: first both parts are necessary, second both parts are salient.

Symmetric dependence causes parts to be necessary. They must co-exist regardless of the conditions imposed by the predicate. Even if the specific conditions imposed by the predicate require information on only one of the parts, the existence of the other part has to be inferred. Let's consider again touch the bottle. As has been said, the conditions imposed by the verb only concern an entity of type *matter*. Nevertheless, since the parts integrating the object denoted by bottle (i.e., *space* \sqcup *matter*) are necessarily dependent, the spatial constituent (i.e., *space*) must also be integrated into the conditions of touching. The two parts cannot be separated without breaking the ontological integrity of the whole. So, the noun bottle denotes the spatio-material configuration even if only the material constituent (i.e., *matter*) is actually accessible under the conditions imposed by touch. In other words, although only one part (*matter*) is activated, the referential entity denoted by bottle is taxonomically identified as a spatio-material configuration.

Symmetric dependence also causes the parts to be salient. There are no specific contexts that could cut off the access to the properties of one of the parts. As parts and their properties are mutually dependent regardless of the specific conditions imposed by the predicate, there are no constraints on meronymic inheritance. The properties of the parts are all potentially accessible. So, all properties of *space* and *matter* are inherited by *phy* and, then, are entirely accessible to those nouns denoting physical objects. Given that there are no restrictions on inheritance, we consider that two symmetric parts are both equally salient with regard to the whole.

Note that when the symmetric dependence is broken, for instance when a glass bottle falls on the ground, the object resulting of this event is no more identified as a bottle. The small pieces of glass on the ground are no more a particular spatio-material configuration representing a receptacle with a particular shape and size. These pieces of glass are identified as an unbounded object which is no more a bottle.³

8.2 Integrity of Physico-Symbolic Entities

Let's analyse now the kind of dependence between a symbolic entity and its physical support when they are integrated into the physico-symbolic type: *phy* \sqcup *symp*. It can be considered that, in some cases, the relation between them is less strong than for symmetric dependencies. In those cases, they are related by an

³ By contrast, grinding metonymies (lamb-animal versus lamb-food, see Copestake and Briscoe 1995) represent those cases where the object is still identified by means of the same noun even when the spatial properties (shape and size) of the whole spatio-material configuration have been lost.

asymmetric dependence. The asymmetric dependence between two parts leads us to infer the two following properties: first the parts are in a dependent-autonomous relationship, second the autonomous part is more salient than the dependent part.

Asymmetry leads one part to be conceived as ontologically dependent of the other, but the inverse is not true. So, we have a *dependent part*, which is not salient, and an *autonomous part*, which is salient. The dependent part does not always co-exist with the autonomous one. Consider for instance the noun poem. It refers to an asymmetric physico-symbolic entity, where the physical facet is perceived as dependent to the autonomous symbolic part. Even though the specific physical support would be destroyed or removed, the symbolic content could still remain either in other physical supports or in the memory of somebody. It follows that even if the specific conditions imposed by a predicate requires only symbolic contents, for instance “to learn”, it is not necessary to infer the existence of a physical support. One could learn the content of a poem without reading it on a paper. On the contrary, the predicative conditions that require only physical objects are not filled by poem if the symbolic content is not somehow presupposed. For instance, poem does not fill the physical conditions imposed by touch because they do not lead us to infer any modification on the symbolic part. Note that words such as *sheet (of paper)* or *file* reverse the order of the dependent-autonomous relation. They consider the physical support as the autonomous part and the symbolic content as the dependent one. The symbolic content of a sheet of paper can be destroyed or removed without destroying the physical support. So, the symbolic content is only inferred if it is required by the conditions imposed by the predicate: for instance, since the verb *read* requires the symbolic content to be impressed on a physical support, the expression *to read the sheet of paper* is meaningful. Yet, since one need not a physical support to learn any symbolic content, the expression *to learn the sheet* seems to be at least odd. In asymmetric dependencies, not only the whole but also the autonomous part may be considered as the direct referent of the noun. So, poem is not only taxonomically identified as a physico-symbolic entity, but also as a symbolic content. Likewise, the noun *sheet* can be taxonomically identified as both a physico-symbolic entity and a physical object.

Asymmetry also causes the autonomous part to be more salient than the dependent part. There are some specific contexts that could cut off the access to the properties of the dependent part. To be more precise, only the properties of the dependent part that somehow presuppose the autonomous part are inherited by the whole configuration. In other words, the access to the properties of the dependent part (i.e., not salient constituent) is only possible if the conditions imposed by the verb lead us to infer some modification on the autonomous part (i.e., the salient constituent).

However, $phy \sqcup symb$ can also be perceived as a symmetric relation. Unlike poem and sheet (of paper), the noun *book* organises the physico-symbolic relation as a symmetric dependence, where the two constituents are entirely accessible. This way, the object denoted by *book* may inherit all properties from both the symbolic content (*learn the book, interesting book*) and the physical support (*touch the book, blue book*).

In this section, we attempted to introduce some informal ideas to build up a theory of dependence on the basis of assumptions concerning constraints on meronymic inheritance. The specific properties of the horizontal relationship (symmetric or asymmetric) determine the mode of access to the information of the parts from the whole.

9 Conclusion and Final Remarks

This article is a contribution to the analysis of lexical inheritance and type subsumption in lexical ontologies. Unlike many works on this research domain, our analysis focused, not on taxonomic subtyping, but on the whole-part relation. We considered parts to be schematic supertypes of wholes relative to specific properties and functional conditions. This way, we described how some meronymic parts are able to characterise and represent the overall whole under the conditions imposed by a specific predicative expression. When a part is conceived as salient relative to the selective conditions of a predicate, it transfers its informative content to the whole by means of a meronymic inheritance device based upon the whole-part link. Meronymic inheritance (information transferred from parts to wholes) was described as an operation of subsumption, where parts subsume wholes. The main objective of the paper was to attempt to show that meronymic subsumption could be adapted in a suitable way to deal with type mismatch and transfer of meaning in the interpretation of complex expressions. For this purpose, we introduced some metonymic principles to identify those parts that subsume a given whole.

A complementary objective of this paper was to use meronymic relations to account for complex and polymorphic types and, then, to describe polysemic information. By contrast, subtype relations were used to organise types only in terms of identity criteria. It follows that taxonomies based on the *is_a* relationship are not any longer multiply-typed hierarchies resulting in an overloading of subtyping links. Rather, they should be constructed as simply-typed lattices organising the ontological relations between entities.

Nevertheless, our current work cannot progress without inquiring the basic assumptions of *formal ontology* and *mereology* (Guarino 1998; Smith 1998; Varzi 1996, 1998). Pertinent questions that could be addressed are the following: What count as a part of a given whole? Under which conditions a part becomes salient and is able to subsume the whole? Is it possible to find specific functional conditions under which the cork of a bottle is able to subsume the whole bottle? What is the ontological role of the more salient parts of a given whole? Does a change of salient parts affect identity?

In order to answer these questions in a suitable way, we claim that theoretical assumptions from both formal ontology and lexical semantics should merge into a more consistent and cohesive theory of word meaning.

Acknowledgments This work has been supported by the Spanish Government (MICINN), within the project FFI2010-14986.

References

- Asher N, Pustejovsky J (2000) The metaphysics of words in context. Ms. Dept. of Computational Linguistics, University of Stuttgart, Germany
- Borgo S, Guarino N, Masolo C (1996) Stratified ontologies: the case of physical objects. In workshop on ontological engineering (ECAI-96)
- Carpenter R (1992) The logic of typed feature structures. Theoretical computer science. Cambridge University Press, Cambridge
- Copestake A, Briscoe T (1995) Semi-productive polysemy and sense extension. *J Semantics* 12:15–67
- Gamallo P (1998) Construction conceptuelle d'expressions complexes: traitement de la combinaison nom-adjectif. PhD thesis, Université Blaise Pascal, Clermont-Ferrand, France
- Gamallo P (2000a) Lexical inheritance in upper-level ontologies. In workshop on ontologies and Lexical knowledge bases (OntoLex2000), Sozopol, Bulgaria. Springer, Berlin
- Gamallo P (2000b) L'interprétation d'expressions complexes nominales analysée comme un processus metonymique. *Revue de Sémantique et Pragmatique* 7:29–58
- Guarino N (1998) Some ontological principles for designing upper level lexical resources. In first conference on language resources and evaluations. Granada, Spain
- Jayez J, Godard D (1995) Principles as lexical methods. In AAAI workshop on representation and acquisition of lexical knowledge. Stanford University, Stanford, pp 27–29
- Kayser D (1987) Une sémantique qui n'a pas de sens. *Langage* 87:33–45
- Kleiber G (1994) Nominale. Essais de sémantique référentielle. Armand Colin, Paris
- Kleiber G (1999) Problèmes de Sémantique. Presses Universitaires Septentrion, Lille
- Langacker RW (1987) Foundations of cognitive grammar: theoretical prerequisites, vol. 1. Stanford University Press, Stanford
- Langacker RW (1991) Foundations of cognitive grammar: descriptive applications, vol. 2. Stanford University Press, Stanford
- Miéville D (1998) Associative anaphora: an attempt at formalisation. *J Pragmat* 31:327–337
- Nunberg N (1995) Transfers of meaning. *J Semantics* 12:109–132
- Ooi V (1998) Computer corpus lexicography. Edinburgh University Press, Edinburgh
- Paiva VD (1993) Types and constraints in the LKB. In: Briscoe T, Copestake A, Paiva VD (eds) Inheritance, defaults and the lexicon. Cambridge University Press, Cambridge
- Pustejovsky J (1995) The generative lexicon. MIT Press, Cambridge
- Pustejovsky J (1998) Lexical semantics and formal ontologies. In: Guarino N (ed) Formal ontology in information systems. IOS Press, Amsterdam
- Pustejovsky J (2001) Type construction and the logic of concepts. In: Bouillon P, Busa F (eds) The syntax of word meaning. Cambridge University Press, Cambridge
- Smith B (1998) Basic concepts of formal ontology. In: Guarino N (ed) Formal ontology in information systems. IOS Press, Amsterdam
- Varzi A (1996) Parts, wholes and part-whole relations: the prospects of mereotopology. *Data Knowl Eng* 20:259–286
- Varzi A (1998) Basic problems of mereotopology. In: Guarino N (ed) Formal ontology in information systems. IOS Press, Amsterdam