A plea for complex categories in ontologies

Alexandra Arapinis, Laure Vieu

► To cite this version:

Alexandra Arapinis, Laure Vieu. A plea for complex categories in ontologies. Applied Ontology, 2015, vol. 10 (n° 3-4), pp. 285-296. 10.3233/AO-150156 . hal-01567016

HAL Id: hal-01567016

<https://hal.science/hal-01567016>

Submitted on 21 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Open Archive TOULOUSE Archive Ouverte (OATAO)

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible.

This is an author-deposited version published in : <http://oatao.univ-toulouse.fr/>
Eprints ID : 16894

To link to this article : DOI : 10.3233/AO-150156
URL : <http://dx.doi.org/10.3233/AO-150156>

<p>To cite this version : Arapinis, Alexandra and Vieu, Laure <i>A plea for complex categories in ontologies</i>. (2015) <i>Applied Ontology</i>, vol. 10 (n° 3-4). pp. 285-296. ISSN 1570-5838</p>
--

Any correspondence concerning this service should be sent to the repository administrator: staff-oatao@listes-diff.inp-toulouse.fr

A plea for complex categories in ontologies

Alexandra Arapinis ^{a,*} and Laure Vieu ^{a,b}

^a *Laboratory for Applied Ontology, Institute of Cognitive Science and Technologies, CNR, via alla Cascata 56/C, Povo, 38123 Trento, Italy*

E-mail: aarapinis@gmail.com

^b *Institut de Recherche en Informatique de Toulouse, CNRS, Toulouse, France*

E-mail: vieu@irit.fr

Abstract. This paper investigates an issue at the interface between language and ontology. It is argued that the phenomenon of ‘inherent polysemy’ observed in lexical semantics for nouns such as *book* or *country* actually is a deeper phenomenon grounded on specific ontological relations involving the entities referred to. It is shown that this phenomenon emerges not only in language but also in most available ontologies. Beyond the ‘dot types’ used in some linguistic theories to account for logical polysemy, it is proposed to introduce ‘complex categories’ in ontologies in order to solve incoherence and inconsistency issues appearing when this phenomenon is not acknowledged and to characterize complex categories on the basis of formal ontology relations.

Keywords: Inherent polysemy, dot types, complex categories, mereology

Accepted by: Nicola Guarino

1. Introduction

It is well known that polysemy is a bottleneck for systems managing semantic data. The problem of polysemy is in fact multiple, in the sense that different kinds of polysemy raise different conceptual issues and seem to call for different treatments. However, following the way polysemy has traditionally been dealt by lexicographers, many information systems simply assume that polysemy is a univocal phenomenon that can be addressed superficially (or rather can be by-passed) by enumerating the different senses of each polysem and relying on disambiguation technics to identify the relevant sense, that is, each separate concept that the polysem might denote, on each particular use. This strategy, known as ‘sense enumeration’, has both conceptual and practical limitations, in particular in dealing with more regular or systematic kinds of polysemy. Indeed, it is now widely admitted, at least in the cognitive linguistic and formal semantic communities, that systematic polysemy is very productive and pervasive, mobilizing general patterns of conceptual relatedness that structure our perception of the domain of reference (see Apresjan (1973) and subsequent work). Therefore, merely listing the senses inevitably fails to account for the conceptual/ontological mechanisms that trigger such multiple meaning phenomena.

In this respect, conceptual analysis and ontological methods can bring a valuable support, providing the means to express the conceptual relations involved in systematic polysemy at the right level of ab-

*Corresponding author. E-mail: aarapinis@gmail.com.

straction, instead of pursuing a time-consuming and much more error-prone local strategy that deals with each expression in isolation. The important role of ontologies in providing the conceptual foundations and formal means to express such relations, and thus gain access to the intended meaning, has recently been brought to attention by Gangemi et al. (2000) and Pisanelli et al. (2004). These authors focus on the effects of logical polysemy on medical terminologies arguing for the usefulness of considering ontological relations among entities in order to better grasp the metonymic mappings they trigger.

In this paper we will show that logical polysemy not only has an ontological base, but most crucially invites us to extend the standard schema of ontological categorization. In a nutshell, it points to the existence of entities that instantiate *complex categories* that ‘overlap’ with disjoint domains of entities. This idea builds on the linguistic hypothesis of ‘dot objects’ and ‘dot types’ introduced by Pustejovsky (1995) and Asher (2011) to account for the phenomena of nominal logical polysemy, which will be introduced in Section 2.2. Then we will explore the ontological foundations of the dot complex type constructor. We will argue for its close connection with the philosophical notion of constitution (Section 3), and outline a formal mereological characterization (Section 4). But before we enter in the heart of the discussion of complex ‘dot’ categories, we think it is important to go a bit deeper into the typology of the different kinds of polysemy, and situate the notions of systematic polysemy, logical polysemy, metonymy, and inherent polysemy with respect to each other. So doing, we will see that authors like Gangemi et al. (2000) and Pisanelli et al. (2004) mistakenly reduce logical polysemy to a kind of metonymy, thus missing the deep categorization issues at stake. We claim such issues are of significant importance for applied ontology. Indeed, we will first of all show how inherent polysemy affects not only lexical ontologies, but also ontologies with no linguistic purpose (Section 2.1).

2. Why do we need complex categories in ontologies?

2.1. Multiple inheritance and the conceptual complexity behind inherent polysemy

Nouns are often assumed to refer to classes, and even when such a strong correspondance is not fully endorsed, concepts are in practice often accessed through language. For these reasons, ontology construction is unsurprisingly affected by polysemy, even when the ontology is not built from texts. In particular, unresolved ambiguities are often ontologically handled through multiple inheritance, each subsumption relation implicitly representing one of the concepts that are conflated in the ambiguous class. A natural remedy to such problematic instances of multiple inheritance, which often result in taxonomic confusion and incoherence,¹ is to resort to a multiplicative strategy to disambiguate the problematic classes. The undesired consequences of such erroneous use of multiple inheritance, which occurs alike in linguistic and non-linguistic ontologies, can indeed easily be avoided by a finer conceptual analysis and categorization, and a more rigorous use of the subsumption relation. It is a matter of good practice and does not fundamentally question the way we conceive of ontological categorization. Over the past years, significant efforts have in particular been made to improve existing resources, using sense enumeration to solve ambiguities. Kumar and Smith (2003) note, for instance, that, “the Gene Ontology [GO]² definition of the term ‘Molecular Function’ is: ‘the tasks performed by individual gene products; examples are transcription factor and DNA helicase’. The term ‘task’ is unfortunately ambiguous and GO, like the UMLS³ Semantic Types, correspondingly incorporate some confusions in the distinction

¹In the technical sense that the class inheriting from two disjoint classes cannot have any instance.

²<http://geneontology.org/>.

³<https://www.nlm.nih.gov/research/umls/>.

between functions and their functioning. [...] This problem was remedied by a policy change effective as of March 1, 2003 whereby GO molecular function term names are to be appended with the word ‘activity’. Because, however, the change was not applied to the parent term ‘Molecular Function’, and because associated definitions were not overhauled in the light of the new policy, some confusion as between ‘function’ and ‘activity’ still remains” (pp. 144–145).

However, as we will now see, polysemy-driven multiple inheritance cannot always be repaired in this way. Indeed, some cases of systematic polysemy (Apresjan, 1973) rather invite a deep revision of standard simple ontological categories. Within the range of systematic polysemy phenomena, which are admittedly revealing of general non-linguistic organization principles structuring our conceptualization of the world, inherent polysemy raises very particular referential challenges (Pustejovsky, 1995, 2008; Asher, 2011), with ramifications both at the linguistic and the ontological level. More precisely, the peculiarity of inherently polysemous nouns, and their ontological counterparts, lies in the fact that the disjoint senses/concepts are not antagonistic (viz. do not call for disambiguation), but are rather brought into play simultaneously, co-specifying a unique referent. For example, a term like *book* is inherently polysemous between a physical object reading and an abstract informational reading, which are both ‘aspects’ or ‘facets’ of a unified complex meaning/concept. There are no purely physical books on the one hand, and informational books on the other. Books are precisely complex ‘materialized informational contents’⁴ that do not correspond to the conjunction of disjoint ‘aspects’ (viz. multiple inheritance introduces incoherence or inconsistency), but neither to their disjunction (viz. subsumption does not allow to represent the fact that aspects are simultaneously present, not mutually exclusive), as illustrated in Fig. 1.

Note further that inherent polysemy spans through a wide range of domains, including, among others, all linguistic artifacts (<information content, physical object>), but also organizations (e.g. *bank*, *university*: <staff, institution, building>), or pathologies (e.g. *inflammation*, *neoplasm*: <process, physical object>).

While this phenomenon has not, up to now, attracted the due attention from the ontological community,⁵ a quick overview of some of the most widely used ontologies clearly shows that the issue of inherent polysemy cannot easily be bypassed.

Of course, this unsurprisingly happens in lexical ontologies that straightforwardly assimilate lexical entries to ontological categories. Thus, scanning through wide coverage lexical ontologies, like Word-

Fig. 1. What complex categories are not.

⁴While such materialized contents admit different kinds of physical supports, we take it that a material support is necessary for any kind of book to exist. This holds of paper books as well as electronic books. We further assume that books that only exist in the mind of the author have some material support in his brain. Conversely, we assume that a book-like physical object with no content (like an empty notebook) cannot properly be considered as a book *tout court*. The use of the term *book* to denote pages with no content rather seems to point to the future book that will come into existence once the pages are filled.

⁵With the exception of Nardi et al. (2015) who consider the difficulty to properly conceptualize the notion of *service* to be due to the inherent polysemy of the term.

Net,⁶ it is clear that some of the uses of multiple inheritance are meant to capture this type of inherent polysemy. A letter, for example, understood as “a written message addressed to a person or organization”, is in a sense both an abstract information object and physical object. As reported by Verdezoto and Vieu (2011), this is represented in WN 3.0, by allowing *letter%1:10:00* to inherit its abstract content from its hypernym *text%1:10:00* and its physical aspect from its hypernym *document%1:06:.* But even though the category of abstract and physical objects are clearly disjoint ontological categories which is a compelling reason to remove this multiple inheritance, one cannot simply split these two aspects of letters. In reading a letter, we do not merely read the information object, nor the mere physical object, but a complex entity composed of the two disjoint aspects. Reading is accessing the information object through the inspection of the physical object.

Most importantly for us, similar challenging examples of multiple inheritance can be found in non-lexical ontologies alike, such as DBpedia,⁷ Schema.org,⁸ or the UMLS ontology. For example, though the structure of DBpedia’s proper ontology is very poor, and in fact close to flat, it is interesting to note that it nevertheless resorts to multiple inheritance for categories like that of libraries, which are represented as subsumed by both ‘EducationalInstitution’, a subcategory of ‘Agent’, and ‘Building’, a subcategory of ‘Place’. But just separating the aspects of libraries that belong to disjoint ontological dimensions does not seem to do the job. Thus, when a new library opens in a city, what is it that opens? The building or the institution? Can there be one without the other? Is a library with no dedicated place really a library? Or is any place with stored books a library, even when it is not institutionally created and acknowledged as such? In this case, attributing the use of multiple inheritance to the practical choices motivated by the poor ontological backbone of DBpedia would simply be misleading. For a start, a very similar use of multiple inheritance happens in Schema.org with the type ‘LocalBusiness’ inheriting from both the types ‘Organization’ and ‘Place’. And looking at analogous examples in an ontology like UMLS is another good indicator for that. In fact, the biomedical domain has recently concentrated a lot of attention in the ontology design community, and is renowned for the significant efforts that have been made towards the design of conceptually sound ontologies. One would thus expect not to find inappropriate multiple inheritance as the ones found in the other ontologies. However, the biomedical domain is populated by classes that are simultaneously categorized as processes and (resulting) substances.

In principle, events and entities are clearly distinguished in UMLS, being the top nodes of disjoint trees, a fragment of which is represented in Fig. 2. No category should therefore belong to both of these

Fig. 2. UMLS Entity tree and Event tree.

⁶<https://wordnet.princeton.edu/>.

⁷<http://wiki.dbpedia.org/>.

⁸<http://schema.org>.

hierarchies at the same time. But looking at the very definitions of some of these categories, it appears that this clearcut partition between two realms of being is in some cases problematic, and is for instance explicitly transgressed with respect to the ‘Disease or Syndrome’ and the ‘Anatomical Abnormality’ categories. Indeed, the ‘Disease or Syndrome’ type is defined as follows: “A condition which alters or interferes with a normal process, state, or activity of an organism. It is usually characterized by the abnormal functioning of one or more of the host’s systems, parts, or organs. Included here is a complex of symptoms descriptive of a disorder”. This definition is however immediately followed by a usage note addressed to the annotators: “Any specific disease or syndrome that is modified by such modifiers as ‘acute’, ‘prolonged’, etc. will also be assigned to this type. If an anatomic abnormality has a pathologic manifestation, then it will be given this type as well as a type from the ‘Anatomical Abnormality’ hierarchy, e.g., ‘Diabetic Cataract’ will be double-typed for this reason”. UMLS thus allows for multiple inheritance crossing the entity/event domains, though this is done in a more principled way than in most other ontologies. Indeed, it is mentioned in an explicitly guideline and specified for high-level categories, thus resulting in an inadequate, more precisely incoherent because of the assumed disjointness, but at least uniform treatment of the most relevant kind of complex category for the biomedical domain.

Finally, if one chooses an enumerative approach to resolve such cases of multiple inheritance, this should be coherently echoed through all its hyponyms, and all the way down to its instances. But this is clearly not actually the case neither in WordNet nor DBpedia. So, while WordNet distinguishes countries as locations (*country%1:15:00*) and countries as groups of people, viz. citizens (*country%1:14:01*), individual countries are usually represented by a single synset which is categorized as an instance of one of the two senses of *country* (*ethiopia%1:15:00* for example is only an instance of *country%1:15:00*).⁹ Looking now at the ABox of DBpedia, it is not difficult to come up with examples of individuals being annotated as instances of multiple categories. In particular, while DBpedia reduces museums, unlike libraries, to their locational aspect, many individual museums are in fact found to be also instances of the category ‘Organisation’, like: the Long Beach Museum of Art, the Museum of Far Eastern Antiquities from Stockholm, the Pettaquamscutt Historical Society Museum, etc. Regarding individual museums, DBpedia comes much closer to the way libraries are represented. In fact, we have observed that the categorization of instances under multiple types is much more frequent than multiple inheritance. If, following general good practice principles, categories involving contradictory individuation criteria should simply be broken down into simple categories, there seems to be a serious reluctance in applying the same multiplicative strategy for individuals. Obviously, the resulting knowledge base is inconsistent in the case of multiple instantiation as soon as the disjointnesses (here, between the categories ‘Place’ and ‘Organisation’) are made explicit.

2.2. Selectional restriction arguments for complex categories

Probably the most compelling argument for the existence of complex categories comes from looking at selectional restrictions of predicates, and correlatively of properties and relations. Co-predication tests, combining predicates with contradictory ‘selectional restrictions’ are thus standardly taken to show that certain names express complex faceted concepts (Cruse, 1986), and refer to entities of a complex type (Pustejovsky, 1995; Asher, 2011).

⁹The unicity of individual countries creates further difficulties when combined with meronymic relations. Ethiopian citizens are thus related to Ethiopia through the member holonym relation, which leads to the non-sensical conclusion that they are members of an entity that is assumed to be a location and not a group of people (Verdezoto and Vieu, 2011).

- (1) *This thick* <physical object> **book** *is incomprehensible* <information object>.
- (2) *The inflammation is acute* <process> *and visible to the naked eye* <physical object>.
- (3) *The university in the city center* <building> *specializes in humanities* <institution, staff>.¹⁰
- (4) **Brazil** *is a large* <land> *two-century-old* <institution> *portuguese-speaking* <people> *country*.
- (5) *Dr. Smith's dental service is expensive* <service offering> *but unreliable* <service delivery> (taken from Nardi et al., 2015).

What counts as a proper copredication test is not always clear in the literature though, and the acceptability of various constructions varies. This is why Ježek and Vieu (2014) suggested to focus these tests on constructions with a single clause as in (1), (3) and (4). This move avoids the interference of co-reference and discursive phenomena, which may result in puzzling or zeugmatic cases as in (6).

- (6) *They took the door off its hinges and went through it* (taken from Cruse, 1986).

A further conceptual issue that is revealed by looking at predicates (analogously, properties and relations) and the categories they selectively apply to, is that they sometimes appear to mobilize for several aspects at once. As already evoked, reading a book involves both the information and the physical artifact, and treating a tumor both the physical abnormality and the pathological process.

To solve these semantic puzzles raised by inherent polysemy, typed theories of the lexicon (where taxonomically ordered types specify the kind of entities terms can apply to) first introduced the notion of ‘dot-types’ (and dot-objects) to account for complex selectional restrictions and the possibility of copredications selecting for different aspects of the same referent. Such complex types enter the taxonomy of types as illustrated in Fig. 3.

Following an analogous line of argument, it seems that a proper treatment of relations and properties militates in favor of the existence of complex categories at the ontological level. In fact, it is noteworthy that most ontology constructions generally focus on domain independent relations (part-whole, dependence, subsumption, etc.) which carry no selectional restrictions, and have thus failed to appreciate the issue of the selectional restrictions that are carried by domain relations and properties.

Once more, looking at DBpedia and the relations ranging over the categories just discussed clearly militates for a proper treatment of their complexity. For instance, DBpedia includes ‘being the curator of’ among the properties of museums, which are, as mentioned above, buildings, ‘building’ being itself a subcategory of ‘location’. But certainly, this is a relation that people do not bear to simple buildings or locations. Curators are the administrative directors or managers of museums and their collections. We suspect this is why ‘being the curator of’ is tagged as ranging over the specific category of museums, while most other properties are defined as ranging over the super-categories of buildings or locations. The same kind of tagging strategy is in fact found with other complex categories which are reduced to a single dimension in DBpedia, but for which properties selecting for the other facets are given. The

Fig. 3. The idea of dot-types.

¹⁰As in the case of blank books, we sometimes use names of complex institutional entities to designate a building that is not yet in use for institutional purposes, pointing to its future role in relation to the corresponding institution and its staff: *The new large bank downtown is not yet completed*.

category of countries, i.e. a subclass of ‘PopulatedPlace’, is one such example. Most of the properties ranging over this category are tagged as bearing over the domain of populated places, with few notable exceptions such as ‘European Union Entrance Date’, which is said to range over countries. This makes sense since a place cannot as such enter into the European Union. A country qua geographical stretch *is* part of the European continent, but entering into the Union rather focalizes on the institutional (governmental) aspect of countries.

This being said, it is of course not surprising to find, in DBpedia, equivalent incoherences at the level of property attributions to individuals. For instance, while the Ethnological Museum of Berlin is categorized as an architectural structure and therefore a place, one finds among its properties the fact that it was established in 1960, or that it belongs to the type of archeological museums, viz. two properties of abstract institutions. The same with individual countries like the German New Guinea, which has the property of having been dissolved in 1919-06-28. Indeed, it is certainly not the location that was dissolved, nor the population.

Beyond the difficulty of handling the selectional restrictions of properties when a complex category is artificially reduced to one of its aspects, and where a multiplicative approach could in fact superficially solve the issues, other properties clearly range over the whole complex category, involving simultaneously all of its constituents. As evoked above, this is for instance the case of medical relations like ‘treat’ or ‘diagnose’ which involve diseases but at the same time their visible symptoms. In UMLS, these are in fact tagged as ranging over multiple and disjoint (according to UMLS’s own taxonomic structure) categories. While not satisfactory, this strategy is coherent with the decision to treat sub-categories of anatomical abnormalities as sub-categories of diseases and vice versa. It is further motivated by the intuition that a treatment acts necessarily both on the process of the disease and on the abnormality it generates. Analogously, in the cases where a disease is essentially producing anatomical abnormalities, the diagnoses will necessarily bear on both. Diagnosing the disease means observing the presence of abnormalities, and vice versa, a diagnosis of the anatomical abnormalities is a diagnosis of a process producing them. This does not mean that the process and the result are reducible to one another, but distinguishing them without further acknowledging that they form a sort of unity would lead to an absurd and, most importantly for ontology design, impracticable multiplication of properties and relations. If, following the suggestion of Kumar and Smith (2003), and the recommendations of BFO, widely endorsed in the biomedical domain, we were to completely isolate for example the physical abnormality called diabetic cataract from the pathological process called cataract, and from the pathological function called cataract, yet a third aspect of the disease, then one would further have to assume that there are three distinct relations involved in treating or diagnosing diabetic cataract. The authors did not take this into consideration at all. But it seems highly implausible that in diagnosing (or treating) a diabetic cataract, a medical doctor could in any rational way be seen as performing three distinct kinds of diagnosis (or treatment).

Gangemi et al. (2000) and Pisanelli et al. (2004) on the other hand, do mention such relations, but focus on their ambiguous interpretation depending on the domain of the relation (viz. physicians vs. therapies vs. devices) and not on their range. They propose a metonymic account of such ambiguities based on “the basic [unambiguous] meaning of notions like treat”. But their analysis misses precisely the fact that this basic meaning ranges over complex categories with taxonomically incompatible aspects.

In a word, the above mentioned difficulties, raised by tagging properties of complex ontological categories, point to the fact that the (linguistic) motivation for dot types, having to do with semantic selectional restrictions, has a direct ontological counterpart. It is not just that the predicates select for an

argument of complex type, but, more fundamentally, that the very actions and properties jointly mobilize the different aspects of a complex entity.

3. The ontological constitution behind complex categories: Dimensions of coincidence

3.1. *Material constitution*

The strong connection between logical polysemy and ontological relations structuring objects is clearly in the background of type-oriented theories of logical polysemy, and the Generative Lexicon (Pustejovsky, 1995) especially. Indeed, the qualia structure was explicitly designed to reflect the four Aristotelian causes¹¹ which are assumed to inform the lexicon. But very little is in fact said about the nature of the dot constructor, and its ontological foundation in complex categories is not even considered. In this section we will defend that the underlying ontological relation is a kind of constitution, a generalization of the relation of material constitution which also traces back to Aristotle and his hylemorphic theory of primary substances (i.e. roughly stated, physical objects are compounds of matter and form).

In a nutshell, theories of material constitution reject the view that physical objects are reducible to pure matter (Rea, 1997). One and the same amount of matter can in particular constitute more than one object, the different constituted objects being co-located but of distinct kinds or forms. A now classical example put forward in favor of material constitution, and of the proliferation of objects it implies, is that of the statue, the piece of clay out of which it is made, and its particles of matter. These three things share of course an important number of properties, and most significantly their spatial location. But they also differ in important respects. For example, the lump of clay existed before the statue, and would still exist if its shape came to change completely, while the statue would then be destroyed. More generally, they have different persistence conditions and modal properties, viz. they differ in the way they were, will, would or must be. A parallel reasoning shows that the lump of clay and the set of its particles are distinct, since the latter would still exist if its elements were scattered around, but the lump would not, since it requires connectedness.

In view of Leibniz's law, many philosophers have thus felt compelled to conclude that the statue, the lump of clay, and the set of its particles are non-identical, but are related in terms of constitution, i.e. a form of unity without identity. The set of particles materially constitutes the lump, which materially constitutes the statue for a certain period of time. As Baker puts it: "We need constitution to be similar to identity in order to account for the fact that if x constitutes y , then x and y are spatially coincident and share many properties; but we also need constitution to differ from identity in order to account for the fact that if x constitutes y , then x and y are of different kinds and can survive different sorts of changes" (Baker, 1999, p. 145).

At first sight, this strong ontological dependence seems to come close to the sort of dependence involved in complex categories lexicalized as dots, allowing for incompatible categories to form a kind of unity without postulating incoherent subsumption relations. And in fact, some of the complex categories discussed above are also invoked in discussions about material constitution. Fine (2000), for instance, takes letters (written messages addressed to someone) as paradigmatic examples of material constitution. So a letter inherits its material properties from the paper and the ink that constitute it (e.g. being stained), but is not reducible to them. In the same article, Fine further considers the interesting case

¹¹The material cause is encoded in the constitutive quale of lexical items, the formal cause in the formal quale, the efficient cause in the agentive quale, and the final cause in the telic quale.

where one and the same piece of paper and set of ink-traces constitute two letters, interpretable in two languages, and being addressed to two people (each with one of the required linguistic competences). The problem is that he does not go beyond material constitution to explain how the two informational contents (expressed in the two languages) contribute to constitute the two letters.

In fact, the central figures¹² of the theory of constitution only consider it from the perspective of materially coincident entities and do not even mention the possibility of other forms of constitution. This notion thus needs to be generalized in order to take into consideration other dimensions of coincidence than the material one. One should also be able to make sense of the fact that entities with distinct identity conditions can nevertheless co-constitute a further complex entity, allowing for combinations of constitution or more generally dependence relations, not only along the vertical dimension (as with the statue, the lump of clay and the particles) but also the horizontal one as illustrated by the case of dot types.

3.2. *Extending constitution beyond material coincidence*

In the past five years or so, scholars have tried to see how theories of constitution could apply to institutional reality. In so doing, they were faced with the problem of accounting for people and abstract rules constituting institutions such as universities, banks, etc. Obviously, these had to involve a different form of constitution from the material kind. Indeed, institutions are not necessarily located where their affiliates are, and it does not even make sense to speak of material coincidence of an institution with the set of abstract statuses and rules that make it up.

These questions thus motivated Wilson (2005) and more recently Hindriks (2012) to extend the notion of constitution to agency, keeping the idea that constitution is a sort of dependence involving coincidence, but extending the latter constraint from spatial coincidence to action. The basic idea was that institutions are agents, in the sense that they undertake certain actions, and that these actions coincide with the actions of their affiliates. For instance, if a bank makes an investment, it is in virtue of some person or group of people making this investment *on behalf of* the bank. This *on behalf of* needs of course to be further qualified to properly grasp what is the constitution relation involved in this particular domain, and has been elaborated differently by different scholars. But what is most important for us is that these works point to a core notion of what can be still called constitution, from which a family of more specific kinds of constitution can be derived by further specifying the dimensions of coincidence.

A similar objection to mainstream theories of constitution was also made by Jansen (2009), further suggesting that social constitution should be extended to immaterial entities. Beyond material and agency constitution, Jansen thus points to the fact that “deontic structures (rights, obligations, commitments to act in certain ways in given circumstances) as typically reported in legal documents that testify of the creation of organizations immaterially constitute the organization which perdures as long as the deontic structure is in force”. Though not further worked out by the author, one could qualify this sort of constitution as formal, in the Aristotelian sense, viz. the form specifying what it is to be an entity of a particular kind. Thus understood, formal constitution is typically involved in the way abstract statuses make up institutions like organizations, but also in the way informational contents make up informational artifacts like books or letters. Indeed, the meaning of, say, a token word (written or spoken) is its form in that it is precisely what makes it the word it is.

Understood in this broad sense, Arapinis (2013) thus argues that constitution can furnish the ontological counterpart of the semantic relation between each single type and the complex or dot type formed

¹²Let us mention, among others, Baker (1999), Fine (2000, 2003), Koslicki (2008), Lowe (1995), Simons (1987).

on them. The underlying notion of coincidence naturally accounts for many metaphysical but also linguistic issues raised by dot types, such as property inheritance, or the fact that the PROCESS/RESULT pattern of polysemy can give rise to a dot, as briefly mentioned in Section 2.2. In particular, the fact that *inflammation* denotes a complex object of type PROCESS•PHYSOBJ, as attested by co-predication tests, is explained by the fact that the resulting anatomical structure is co-located with the process for as long as the process goes on. But this is not case of all processes with resulting physical objects. For example, the word *construction* is also polysemous between a process or a result reading but does not support co-predications (see Ježek & Melloni, 2001), since the resulting object only comes into being after the process is over. In the same vein, the coincidence condition involved in the extended notion of constitution naturally explains why example (6) seen above is zeugmatic, even though *door* admittedly denotes objects of type APERTURE•PHYSOBJ. In this case, the spatio-temporal coincidence between the aperture and the frame is destroyed, and what was a single complex object is now broken down into two distinct entities. As mentioned in Section 2.2, not all copredications really allow to test for dottiness. In particular, the use of verbal conjunctions and pronouns (as in (6)) may induce discourse relations between the predications, resulting in complex spatio-temporal relations that affect the coincidence of the denoted entities and thus the acceptability of copredications.

The question that remains to be addressed is that of the relation between the constituents that simultaneously enter into a complex constitution relation. Indeed, the literature on constitution only considers the vertical relation between an entity x and the entity y it constitutes. But where multiple entities are so to speak united to form a complex one, there has to be a horizontal relation, viz. the dot relation defined at the level of types, that holds between the constituents themselves. In the next and last section we turn to this issue, and make the first steps towards a formal characterization of the dot relation in a mereological framework enriched by the notion of coincidence.

4. Towards a mereological characterization of dot-objects

What are dot-objects, that is, what sort of entities are instances of complex categories and what is the ontological counterpart of the dot operator in lexical theories? Asher (2011) shows that dot-objects cannot be sets or tuples whose elements would be the different facet-entities united in a dot-object, as dot-objects are not necessarily abstract. It is rather natural to turn to mereology instead, for mereological sums do not abstract over the entities summed. We will now then explore how such a strategy could be pursued, and then show why Asher's objections against a mereological approach do not apply.

Mereology is well known and widely used in formal ontology as the theory of general parthood, yielding in its extensional version an account of identity (Simons, 1987; Varzi, 1996). Basically, mereology theories are constituted of a parthood relation forming a partial order and a sum operator. Proposing to make use of this sum operator to account for the dot operator is of course not enough: summing any two entities will not always yield a dot-object.

Actually, it is already suggested in the literature that although technically the mereological parthood and sum should be allowed to apply universally, conceptual parthood and parthood relations in language do not (Masolo et al., 2003; Vieu, 2006; Mellor, 2006), as they do not occur across fundamental categories: no physical object is part of an event, no abstract entity is part of a physical object etc. So in practice, to model parthood, the sum operator in mereology is filtered to apply only within separate categories.

For modelling dot objects, it is just the opposite: the mereological sum operator has to be filtered to apply only across some fundamental categories. A book will therefore be the sum of a physical object

and an informational object. This is still not enough though, the ‘glue’ or the ‘coherence’ relation in Pustejovsky’s terms, is still missing. The physical object constituting the book I am currently reading cannot be summed with the informational object constituting the book I read last month to yield a third book. The physical object is specifically existentially dependent on a single text, the one which is the form in the Aristotelician sense of the sequence of words printed on the pages, the one I am accessing to through the ink patterns appearing on the physical object when reading.¹³ This dependence implies a coincidence, extending what is introduced in Section 3.1 above, between the informational object and the physical object; in some specific sense, the text is present on the paper. So sums will be filtered both according to the category of the entities summed and according to the presence of a coincidence relation between them, itself consequence of dependence relations.

In a nutshell, if $+$ is the general mereological sum operator, C the general coincidence relation, and A and B two categories, dot-objects of the complex category $A \bullet B$ are so characterized:

$$\forall x(A \bullet B(x) \leftrightarrow \exists yz(A(y) \wedge B(z) \wedge x = y + z \wedge (C(y, z) \vee C(z, y)))).$$

The sum x will also be coincident with each of y and z provided we adopt this simple axiom on coincidence:

$$\forall xyz((x = y + z \wedge C(y, z)) \rightarrow (C(x, y) \wedge C(x, z))).$$

We cannot end here without commenting the fact that before proposing the most worked-out formal account of lexical dot types available Asher (2011) examines the possibility of a mereological approach to model dot-objects to conclude that “the mereological conception, like the pair conception of objects of complex type, seems fatally flawed”. What Asher has missed in the picture is the ‘glue’, the coincidence relation that filters the sums. He does not pretend that the mereological approach will produce books out of any information object-book and any physical object-book, yet, he disregards the exact coincidence between the two entities summed. His main argument is a counting argument based on the example of some author’s collected works in one volume. Supposing that the collected works gather 3 individual works, he argues that the mereological approach would yield 3 different books where there is only one. General mereology of course yields any possible sum, so, the whole volume (physical object) summed with any individual work (informational object) is indeed assumed to exist. But these sums do not make up a dot-object of the complex category *book*, as each individual work coincides only with a proper part of the whole volume, and this part is not an instance of physical object-book. The only book that there is in this situation is obtained summing the collected works, that is the mereological sum of the individual works, considered as an instance of information object-book, to the whole volume, indeed a physical object-book, which is possible since these two entities coincide.

5. Conclusion

In this paper we have shown that the particular polysemy phenomenon called inherent polysemy affects the coherence and consistency of extant ontologies, even carefully crafted ones such as UMLS. We have argued that this is not due to the fact that lexical polysemy as a whole often affects ontology

¹³We will leave aside here Fine’s letter constituted of two different informational objects in two different languages evoked above.

construction, but rather to the fact that this very specific polysemy phenomenon is a linguistic manifestation of complex categories, an ontological construct based on coincidence relations between entities of disjoint categories.

We would thereby like to call the attention of ontology engineers to acknowledge this issue and learn to recognize when complex categories need to be introduced in ontologies. We have described the formal relations founding complex categories, and sketched a mereological account. Of course, a full-fledged theory of complex categories would require further work, but we hope to have also raised the interest of formal ontologists in this respect.

Acknowledgements

This research was partly funded by the FP7-PEOPLE Project 331196. We wish to thank an anonymous reviewer for detailed comments.

References

- Apresjan, J. (1973). Regular polysemy. *Linguistics*, 142, 5–32.
- Arapinis, A. (2013). Referring to institutional entities: Semantic and ontological perspectives. *Applied Ontology*, 8, 31–57.
- Asher, N. (2011). *Lexical Meaning in Context: A Web of Words*. Cambridge Univ. Press.
- Baker, L.R. (1999). Unity without identity: A new look at material constitution. *Midwest Studies in Philosophy*, 23, 144–165.
- Cruse, D.A. (1986). *Lexical Semantics*. CUP.
- Fine, K. (2000). A counter-example to Locke's thesis. *Monist*, 83(3), 357–361.
- Fine, K. (2003). The non-identity of a thing and its matter. *Mind*, 112, 195–234.
- Gangemi, A., Pisanelli, D.M. & Steve, G. (2000). Understanding systematic conceptual structures in polysemous medical terms. In *Proceedings of the AMIA Symposium* (pp. 285–289).
- Hindriks, F. (2012). But where is the university? *Dialectica*, 66(1), 93–113.
- Jansen, L. (2009). Unity and constitution of social entities. In *Unity and Time in Metaphysics* (pp. 15–45). Walter De Gruyter Inc.
- Ježek, E. & Melloni, C. (2001). Nominals, polysemy, and co-predication. *Journal of Cognitive Science*, 22, 1–31.
- Ježek, E. & Vieu, L. (2014). Distributional analysis of copredication: Towards distinguishing systematic polysemy from coercion. In R. Basili, A. Lenci and B. Magnini (Eds.), *Proceedings of the First Italian Conference on Computational Linguistics CLiC-it 2014 & the Fourth International Workshop EVALITA 2014* (Vol. 1, pp. 219–223). Pisa Univ. Press.
- Koslicki, K. (2008). *The Structure of Objects*. Oxford: Oxford Univ. Press.
- Kumar, A. & Smith, B. (2003). The universal medical language system and the gene ontology: Some critical reflections. *Advances in Artificial Intelligence*, 2821, 135–148.
- Lowe, E.J. (1995). Coinciding objects: In defense of the 'standard account'. *Analysis*, 55, 171–178.
- Masolo, C., Borgo, S., Gangemi, A., Guarino, N. & Oltramari, A. (2003). The WonderWeb library of foundational ontologies and the DOLCE ontology, WonderWeb (EU IST project 2001-33052) deliverable D18, LOA-ISTC-CNR.
- Mellor, D.H. (2006). Wholes and parts: The limits of composition. *South African Journal of Philosophy*, 25(2), 138–145.
- Nardi, J.C., Pires, L.F., Falbo, R.A., van Sinderen, M.J., Almeida, J.P.A., Guarino, N., Guizzardi, G. & Fonseca, C.M. (2015). A commitment-based reference ontology for services. *Information Systems*, 54, 263–288.
- Pisanelli, D.M., Gangemi, A., Battaglia, M. & Catenacci, C. (2004). Coping with medical polysemy in the semantic web: The role of ontologies. *Studies in Health Technology and Informatics*, 107, 416–419.
- Pustejovsky, J. (1995). *The Generative Lexicon*. MIT Press.
- Pustejovsky, J. (2008). *From Concepts to Meaning*. Cambridge, MA: MIT Press.
- Rea, M. (Ed.) (1997). *Material Constitution: A Reader*. Lanham, MD: Rowman & Littlefield Publishers.
- Simons, P. (1987). *Parts: A Study in Ontology*. Oxford: Clarendon.
- Varzi, A. (1996). Parts, wholes, and part-whole relations: The prospects of mereotopology. *Data and Knowledge Engineering*, 20(3), 259–286.
- Verdezoto, N. & Vieu, L. (2011). Towards semi-automatic methods for improving WordNet. In *IWCS'11 Proceedings of the Ninth International Conference on Computational Semantics* (pp. 275–284).
- Vieu, L. (2006). On the transitivity of functional parthood. *Applied Ontology*, 1(2), 147–155.
- Wilson, R.Q. (2005). Persons, social agency, and constitution. *Social Philosophy and Policy*, 22(2), 49–69.