ARTEFACTS IN FORMAL ONTOLOGY

Stefano Borgo, Laure Vieu

ABSTRACT:

This chapter examines the field of applied ontology and the characterization of the notion of artefact within this same domain. It begins with a brief discussion of the basic issues that are of special relevance to formal ontology before then going on to present the main existing systems and emphasize their differences. The situation regarding the formalization of artefacts is discussed and the problems behind the few existing proposals are highlighted. The rest of the paper concentrates on our formal definition of artefact which is motivated by philosophical and application concerns. The grounds for the definition are introduced in conjunction with the presentation of the relevant formalization. This proposal is developed within the formal ontology DOLCE but it is designed to be of wider interest. The primitive notions on which the definition is based are constitution, capacity (modeled as a space of individual qualities), and intentionality (which is used to characterize both the act of selecting an entity and the attributing of capacities).

1 INTRODUCTION

Since the early 1990s, there has been an increasing interest, in the knowledge representation area, in formal systems which aim at describing general notions. Indeed, as the research community became aware of the limits of *ad hoc* approaches such as expert systems (that focus on very specialized domains and pay no attention to flexibility and reusability) and conceptual modeling frameworks (which are limited to capturing the distinctions explicitly needed in the target application or domain), the need to provide clear and unambiguous meaning to notions across knowledge domains became self-evident. Simultaneously, researchers started to look to the philosophical tradition for the characterization of general concepts and relations and to build formal systems based on ontological distinctions. All this led to a new area of research now widely, but perhaps naively, known as *applied ontology* or even simply ontology. Research in this area is both theoretical and application driven since it aims to reconstruct and to organize philosophical views into sophisticated formal systems whilst achieving the semantic integration of various information systems.

Although the term ontology has been endowed with several meanings in the knowledge representation field, it does not directly refer to the discipline that goes back at least to Aristotle and has influenced most of western philosophy. The two disciplines are clearly distinct since the general goal of applied ontology is to construct and apply knowledge structures in order to reliably and automatically manipulate information content, and is motivated by research in areas like information retrieval, data management and conceptual modeling. Nonetheless, applied ontology is strongly linked to the philosophical approach because it relies on general philosophical principles and considerations to justify the various adopted knowledge structures.

In this chapter we are interested in the ontological systems, hereafter referred to as *ontologies*, which satisfy the two main requirements of being *formal* and *foundational*. Roughly speaking, an ontology is formal if it is expressed in a logic language endowed with clear semantics (for instance in model-theoretic terms as first-order predicate logic [Hodges, 1983]). This choice is not determined by application concerns (at least not primarily), it emphasizes the relevance that semantic transparency has in this domain. By foundational ontologies we mean those knowledge systems that focus on very general and basic concepts (like object, event, state, quality) and relations (such as constitution, participation, dependence, parthood).

Often the term *formal ontology* is used to cover both the above requirements, thus reminding us of Husserl's distinction between formal logic and formal ontology. In this specific meaning, formal ontology is the study of the interconnections between entities, properties, parts, wholes and collectives. These are considered to be "formal" because they can be exemplified by objects in all domains of reality [Smith, 1998]. To take yet another perspective, one can say that formal ontology is the study of formal (logical) systems which are: *general*, since they include the most usable and widely applicable concepts; *reliable*, as they are logical theories with clear semantics, a rich axiomatization and carefully analysed formal consequences (theorems); and *well organized*, because they are based on philosophical principles the choice of which is explicitly motivated and remains independent from particular domains. In this work, when using the expression "formal ontology" we will be referring to this latter general characterization.

Among the ontological systems that have been used in applications, there are just a few that more or less satisfactorily present these properties. They are BFO, DOLCE, GFO, OPENCYC, and SUMO. They will be discussed in the following sections. We anticipate that unfortunately only a couple of these ontologies will include an explicit characterization of artefacts and that these existing characterizations will be quite shallow.

An ontological system that properly models artefacts would have large potentialities in applications where artefacts are central. For instance, information systems that control processes in the manufacturing industry must rely on a rich but coherent notion of artefact. The integration of artefact data is crucial to concurrent engineering scenarios and to product lifecycle management. Our goal is to show that a satisfactory characterization of the artefact category can be made within a system as complex as a foundational ontology and within the constraints of classical formal logic, in other words, via an axiomatization. In particular, this means formalizing the notion in such a way that it does justice to several of the crucial properties (derived from philosophical considerations as well as from practical usage) that we usually ascribe to artefacts. At the same time, the definition has to remain independent from particular application domains. We achieve this goal by working within a specific foundational ontology (DOLCE) which, in its present form, lacks a characterization of the notion of artefact. However, our work goes beyond the extension of this system since it provides a general analysis of the category of artefacts which is helpful when formalizing this and related notions in other formal systems as well.

It is important to note that in this enterprise we place ourselves within the framework of an ontology of *social reality*. Social reality [Smith, 1995; Searle, 1983] has to do with the part of reality that covers groups of agents and the social relationships therein, actions that are either collective or directed towards a social group, and the whole range of relevant resulting "social entities", such as contracts or companies. Such entities are often dependent on mental attitudes, either individual or collective. A formal ontology dedicated to social reality takes into account all such entities in its domain and attempts to characterize them by modeling general properties and facts. The ontology of social reality and the ontology of mind need to be separated from epistemological studies that would account for the ways in which an agent constructs his or her beliefs about reality, for instance by categorizing entities. Formal ontology takes for granted an objective[1] point of view on reality, that is, a point of view that is external to any particular agent. This stance justifies the introduction of notions like "intentional selection" or "social artefact", which are important to our approach as well as to the philosophical debates that inspired us.

2 FROM ONTOLOGICAL CHOICES TO FORMAL ONTOLOGIES

In this section, we shall begin by giving a brief description of the relevant foundational ontologies. This presentation, though admittedly brief and limited, provides information on the development and quality of the available ontological systems. Later we shall go on to discuss some of the ontological issues which help in the characterizing and comparing of these systems. Although the literature on foundational ontologies and their comparison remains scarce, the ontological topics we will review have largely been analyzed within the context of the philosophical tradition [Rea, 1997].

2.1 Existing formal ontologies

Basic Formal Ontology, BFO. The development of BFO[2] was initiated in 2002 by the Institute for Formal Ontology and Medical Information Science (IFOMIS, first at the University of Leipzig and later at Saarland University). What characterizes this ontology is the careful description of both the general philosophical

[1]Certain ontologies adopt a *cognitive* approach: the categories of entities and the relations used to represent reality are chosen for their compatibility with those arguably used by humans in their language structures and/or their conceptual notions. If a cognitive approach is adopted this does not necessarily mean that the represented facts have to be subjective.

[2]http://www.ifomis.org/bfo

viewpoint and the organization of its structure. The ontology is only partly axiomatized and is not aligned with other knowledge systems, e.g., lexical resources. BFO is actually a framework of sub-ontologies linked together by formal relations. Every sub-ontology must be conceived of as a particular perspective on reality: the user selects the sub-ontology that she finds most appropriate to capture the aspects of the world she is interested in. The most important ontologies in BFO are: SNAP (a series of time-indexed snapshot ontologies, these are ontologies of endurants which, roughly speaking, are objects) and SPAN (a single ontology of perdurants which are, in rough terms, events). SNAP-BFO provides a list of all the entities existing in time such as cars, animals and mountains. They can be seen as "a snapshot of reality" with no temporal extension. By contrast SPAN-BFO is a catalogue of events which necessarily occur over the course of time such as races, deaths and avalanches. SNAP and SPAN are intertwined via transontological relationships since SNAP entities participate in SPAN entities.

BFO (version 1.1) consists of about 40 classes (categories) and is formalized in the weak language known as OWL (Web Ontology Language [Antoniou and van Harmelen, 2004]). It is partially available in first-order logic as well [Masolo et al., 2003]. BFO, which is freely available, has so far been mainly applied in the biomedical domain.

Descriptive Ontology for Linguistic and Cognitive Engineering, DOLCE. DOLCE[3] has been developed at the Laboratory for Applied Ontology (LOA), which is part of the Italian ISTC-CNR, as a reference module for a library of ontologies (mainly within the context of the WonderWeb Project). DOLCE has a definite cognitive bias since it aims at capturing the ontological categories underlying natural language and human common sense. Among the primitive formal relations we find parthood, dependence and constitution.

All the various DOLCE terms and expressions are influenced by philosophy and linguistics. It has a tree-structure that is obtained by applying a "top-down" type of methodology. DOLCE provides a rich axiomatization of the different main categories and their relationships and it has been aligned to WordNet [Fellbaum, 1998; Prévot et al., 2005]. It is publicly distributed (see licence on the web site) and available in first-order logic (including KIF) and weaker languages like OWL, DAML+OIL and RDFS. It is also distributed as a software running in CASL, the Common Algebraic Specification Language[4], which makes available certain theorem provers and graphical devices. It is actively used in several projects in a variety of domains such as manufacturing, linguistics and the Semantic Web. Further information will be provided on this ontology in Section 4.

General Formal Ontology, GFO. GFO[5] was developed at the Onto-Med Re-

[3] http://www.loa-cnr.it/DOLCE.html

[4] http://www.brics.dk/Projects/CoFI/CASL.html

[5] http://www.onto-med.de/en/theories/gfo/. At the time when this paper was being written a new version of GFO had just been presented. We therefore regret it has not been possible to

search Group of the University of Leipzig. It consists of a taxonomy of entities, a taxonomy of relations, and an axiomatization in first-order logic (an axiomatization that is currently still under development). GFO is philosophically well motivated and its ontological choices clearly stated. A crucial guideline for GFO ontology resides in the notion of "levels of reality". These are the material level (biological, chemical and physical), the mental level (that of psychological phenomena) and the social level (where we find agents, organizations and societies). This structure of levels forms the very basis of the ontology from the point of view that every element in GFO is required to participate in at least one of these levels. Note that GFO is a component of a larger perspective since it forms part of the Integrated Framework for the Development and Application of Ontologies (IFDAO), an evolution of the General Ontological Language (GOL) project which dates back to 1999.

The OWL axiomatization of the previous GFO version consisted of about 80 classes, some 100 subclass-relations and around 70 properties. We expect the new version to be of a comparable size. The Onto-Med group is active in the biomedical domain which explains why GFO has been predominantly applied in that area but it has also been implemented in conceptual modeling. As far as we know, no link to WordNet [Fellbaum, 1998] has yet been developed. The ontology is available through a modified BSD Licence.

OpenCyc and Cyc. Cyc, owned by Cycorp Inc., is not a proper foundational ontology but rather a very large, multi-contextual knowledge base enriched with inference engines. It started being constructed in 1984 the aim being to develop a computer program "equipped" with a large amount of commonsense knowledge. The name "Cyc" derives from encyclopedia. The purpose of the project is not, however, to build an electronic encyclopedia, but rather to "complement" such information resources [Guha and Lenat, 1990]. The idea was to create a knowledge base with millions of everyday terms, concepts and rules that would capture the naïve human knowledge bank of reality. To ensure the consistency of the system, the information stored in Cyc is segregated into hundreds of contexts or "microtheories" which are ontologically primitive entities. Essentially, a microtheory is a bundle of assertions that share common assumptions about the world (microtheories are isolated by topics, granularity, culture, etc.). One basic assumption in Cyc is that everything is a member of some microtheory.

OPENCYC[6] was first presented as the "semantic heart" of Cyc but in reality it is just the open source version of the whole Cyc. Ontologically it appears to be deeply affected by cognitive assumptions since its categories try to capture naïve conceptions of the real world or, more simply, common sense knowledge.

OPENCYC adopts a dedicated language (CycL), an extension of first-order logic. It adopts a set-theoretical approach and makes available hundreds of thousands of terms (categories and relations) and millions of assertions (facts and rules).

fully analyze this ontology and compare it with all the others.

[6] http://www.opencyc.org

Clearly, only a mere fragment of more general categories may be qualified as foundational. This fragment is not explicitly singled out by the developers. If one looks at the inheritance relation OPENCYC has a graph-like structure revealing a "bottom-up" approach, that is to say, the organization of the ontology was developed at a later stage to systematize the data present in the knowledge base. OPENCYC provides connections with other formalisms and domains: (partial) translators to OWL and Lisp, a connection to WordNet's synset structure [Fellbaum, 1998; Prévot et al., 2005], and a Cyc-to-English generator.

It should be emphasized that the overall system suffers from its commercial targets. For example, the semantics of Cyc partly depends on the implemented inference engines which shows the relevance of performance at run-time and blurs the semantic transparency of categories and relations. In addition, the characterization of the ontological choices on which Cyc and OPENCYC are based seems to be a secondary task: the documentation is still incomplete and references to the established literature are lacking. OPENCYC does not claim to be a foundational ontology, nor is it a proper top-level ontology because of its size. Nevertheless, it formalizes top categories and their mutual relationships, provides fairly extended descriptions for most of the categories and the examples provided are helpful in conveying (at least in part) the intended meaning of the terms. Finally, the success of this long-standing project (it has been running for almost 25 years) is definitely questioned today.[7]

Suggested Upper Merged Ontology, SUMO. SUMO[8] was created by a private company in 2000 when publicly available specialized ontologies were merged into a single structure in an attempt to obtain a top-level ontology that would be suitable for interoperability, communication and search in the Semantic Web area. The project includes Sowa's upper level ontology [Sowa, 2000] and the work of Guarino and colleagues on theories of space, time and objects [Borgo et al., 1996].

SUMO is not directly influenced by a specific theoretical approach instead it selects from various ontological proposals the categories that seem to be of general use and are broadly accepted by the knowledge representation community. Thus, like OpenCyc, SUMO is not a proper foundational ontology. Nonetheless, it is often included among the others because it is a large ontology used in several applications and one that provides an axiomatization of its terms in a rich language.

SUMO is divided into "sections" or kinds of complementary ontologies that cluster about 1000 terms and relations, 4000 axioms, and 750 rules (but the volume increases considerably if we include all the related domain ontologies). The sections isolate relevant topics: the Mereotopology ontology, for example, contains concepts that deal with the formalization of a general part/whole relation, while the Unit of Measure ontology provides definitions for unit systems. The inheritance structure of this ontology forms a tree obtained through a top-down methodology.

[7]In specialized mailing lists like *SUO* (http://suo.ieee.org/) and *ontolog-forum* (http://ontolog.cim3.net/) a number of discussions have been conducted on this issue.

[8]http://www.ontologyportal.org/

It provides an axiomatization of the categories and their relationships in a version of first-order logic known as SUO-KIF as well as in OWL and can be exploited via several theorem provers. It is available in different natural languages and linked to WordNet [Fellbaum, 1998; Prévot et al., 2005]. SUMO has been implemented in several projects. The distribution of the ontology is regulated by a licence (see the web site).

2.2 Ontological choices

So, ontologies are knowledge systems. They provide a framework in which every entity and relation we want to talk about can be classified. The construction of such a general framework is not simple and it relies on various basic principles, principles that are studied in depth in metaphysics and the best ontologies do indeed refer to the philosophical literature.

Universals, Particulars and Tropes The ontological distinction between *universals* and *particulars* can be characterized by taking the primitive relation of instantiation: particulars are entities that cannot have instances; universals are entities that can have instances. Linguistically, proper nouns are normally considered to refer to particulars, while common nouns refer to universals. For example, Pavarotti, the Italian tenor, is an instance of "person", but he cannot himself be instantiated. (This characterization of the concept of universal is admittedly imprecise since it does not, for instance, clarify whether sets, predicates and abstract entities should be considered as universals or not. A complete presentation of the different notions demands an analysis of these other entities. Fortunately, we do not need to go into further details to understand the overall position of the ontologies described in Section 2.1.)

By adopting a different ontological perspective, one can reject universals and rely on other entities in the way done in the trope theory [Campbell, 1990]. There one claims that the "whiteness" of the specific piece of paper one is holding is a *trope* (a located property or individual quality) while the universal "white" does not in fact exist. Roughly speaking, tropes are properties of specific material entities upon which they ontologically depend because if the entity ceases to exist, so too does the trope. Tropes do not have instances and cannot be confused with universals.

The crucial ontological choice is the decision to include universals in the domain of the ontology, a necessary step if we want to refer to and classify them within the formalism. Some ontologies, such as DOLCE and SUMO, are examples of foundational ontologies of particulars that do not refer directly to universals. OPENCYC and BFO admit both particulars and universals into the domain. DOLCE and BFO include some forms of tropes as well.

Abstract and Concrete Entities Abstract entities (or abstracts) are entities that do not exist in space or time which means to say that they are not located.

In contrast, concrete entities (or concretes) are defined as entities that do exist, at least in time. Mathematical objects (like numbers and sets) are examples of abstracts, while ordinary objects (like cars and planets) and events (such as the 2008 Olympiad and the Second World War) are examples of concretes. The ontological formalization of abstracts seems to depend on negative properties (i.e. a lack of location) but that is not quite correct: one can take a different tuck by claiming that abstracts are eternal and immutable in that they exist at all times and are unchangeable. A third alternative definition is based on the "causal criterion": abstracts possess no causal power while concretes do. Note that in this way we have already switched to a different notion of abstract entity: if abstracts are "timeless", as in the first definition, then it seems awkward to include them in causal relations; conversely it is possible to individuate entities located in time and space that lack any causal power, like the center of mass of the solar system [Lowe, 1998].

Existing ontologies tend to focus on the first kind of characterization. In DOLCE, temporal and/or spatial locations are not defined for categories like `Abstract Quality` and `Abstract`, so it is roughly the "negative" perspective that is adopted.[9] DOLCE also distinguishes between "direct" and "indirect" location. Some entities do not have a direct location but they inherit their locations from entities on which they depend: tables inherit their temporal locations from the temporal locations of events of which they are participants. In OPENCYC, instances of the class `SetOrCollection` do not have spatial or temporal locations, thus OPENCYC adopts the same "negative" perspective on abstracts. Elements in `TemporalThing`, a subclass of `Individual`, are at least located in time while `SpatialThing`, also a subclass of `Individual`, are at least located in space. It is not clear if there are instances of `SpatialThing` that are not also instances of `TemporalThing`, that is to say, individuals that are located in space but not in time. If not, then all `SpatialThing` (like all `TemporalThing`) are concretes. In SUMO, the distinction between `Physical` and `Abstract` is very similar to the distinction between concretes and abstracts: elements in `Physical` are said to be entities "that have a location in space-time" and in `Abstract` they are entities that "cannot exist at any particular place and time without some physical encoding or embodiment". As far as we can see, the BFO ontology only takes into account entities existing in space and/or in time, that is to say, only concrete entities.

Endurants and Perdurants Classically, endurants (also sometimes called continuants or objects) are characterized as entities that "are" in time; they are wholly present (all their proper parts are present) at any given time of their existence. On the other hand, perdurants (also called occurrents or events) are entities that "happen" in time, they extend in time by accumulating different "temporal parts", so that, at any time t when they exist, only their temporal parts at t will be present.

[9] The `courier font` is used to denote the names of categories or classes of entities in the ontologies described. It is a notational system that is also adhered to in the quotations regardless of the authors' chosen system of notation.

For example, the car you now own can be viewed as an endurant because it is now entirely present, while "your driving to the office" is a perdurant because "your driving out of the garage" is not present when "your driving through the city centre" happens (assuming that these are events that actually occur when you drive to the office). Sometimes only perdurants are admitted in an ontology. It then becomes possible to distinguish between ordinary objects (like "a person") and events or processes (like "a person's life"), relying on properties that lie outside spatio-temporal aspects. (It should be noted that other ways of characterizing endurants and perdurants have also been proposed.)

DOLCE assumes a classical view which accepts both the concept of endurant and that of perdurant. OPENCYC has a similar view: the class `SomethingExisting` (i.e. entities that remain relatively stable throughout their lifetimes) corresponds fairly closely to the classical concept of endurant. Analogously, `SituationTemporal` (or the union of `Event` and `StaticSituation`) corresponds to the classical notion of perdurant. In the case of SUMO the distinction is between `Object` and `Process`. Here processes are characterized as "the class of things that happen and have temporal parts or stages", while for objects a less standard interpretation is accepted: "an `Object` is something whose spatio-temporal extent is thought to divide into spatial parts roughly parallel to the time-axis". In any case, note that in SUMO objects and processes are considered to be necessarily located in the space-time. In BFO the distinction endurant-perdurant forms the basis to the development of two separate sub-ontologies (in other words, the two types of entities do not coexist in the same ontology): SNAP-BFO contains only endurants, while SPAN-BFO contains only perdurants.

Co-located entities No matter what one decides about the ontological status of space and time, one can include spatially and/or temporally co-located objects. It is natural to accept objects that are temporally co-located (at least in part), like the moon and the earth or oneself and one's clothes but the embodiment of spatially or spatio-temporally co-located distinct objects can sometimes be questioned. This issue is addressed by posing questions like: Is a hole different from the region of space it occupies? Is a statue different from the matter which constitutes it? Is a person different from his or her body? The subject is complex and includes rather difficult relations like identity across time, material constitution, essentiality and modality.

DOLCE, which takes a multiplicative approach, uses spatial co-location and the relations of dependence and constitution in order to "stratify" co-located entities. For example, persons (elements of `Agentive Physical Object`) are constituted by their bodies (`Non-agentive Physical Object`), and elements of `Physical Object` are constituted by elements of `Amount of Matter`. OPENCYC has a weaker position. It sometimes adopts a genuine multiplication of co-located entities (e.g., it considers a statue and the matter that constitutes it as distinct co-located entities). On the other hand, it takes persons to be entities in the class `CompositeTangibleAndIntangibleObject`, so that Marilyn Monroe, for instance, has two components: a “body” and a “mind”. SUMO, as far as we can see, does not have an explicit position on this issue. It might be that this ontology suffers from the heterogeneity of the basic theories on which it is founded (see page 6). Finally, BFO approaches this issue by distinguishing between different SNAP ontologies: a statue would be an element of an ontology of art (or of social reality) while the material it is made of would fall into an ontology of physical reality.

3 ARTEFACTS IN EXISTING ONTOLOGIES

This section examines the present situation regarding the formalization of the notion of artefact in formal ontology. In this respect, it constitutes a preliminary step to our study that begins in Section 5.

Generally speaking, the study of artefacts has attracted the attention of researchers from different domains ranging from engineering to philosophy and psychology to linguistics. Despite this wide interest, all attempts to either formally or informally characterize a shared notion of artefact have come up against serious problems. Existing formal ontologies indirectly register this fact. Indeed, of the five systems listed above, only OPENCYC and SUMO include a category of artefacts. This might be partly attributable to contingent aspects: certain formal ontologies are still strengthening their top-level concepts, while others focus primarily on domains where the role of artefacts is marginal, such as in the biomedical study of living organisms.[10] In other cases, it is all the result of specific choice: the notion of artefact may not be considered by some to be so general and basic that it warrants inclusion in foundational ontology.

Nonetheless, we suspect that were a clear and shared characterization of artefact available in the literature, all formal ontologies would happily make it part of their system. After all, it is indisputable that artefacts are omnipresent components of our social life.

These considerations highlight the need to extend and enrich the debate on the properties that distinguish artefacts from other entities, a topic that will be resumed later in the paper. For now we shall consider how the category of artefact is introduced to the two formal ontologies that deal with it. Since we are not interested in the particular formalization of these systems, we shall concentrate on the overall notion by looking at the inheritance structure for this category and at the explanations accompanying the relevant categories but we will not take into account the formal issues.

3.1 Artefacts in OPENCYC

In OPENCYC[11], the class `Artifact` is part of `UniversalVocabulary`, one of the most general microtheories of OPENCYC. From the given description, Cyc’s asser-

[10] See, for example, the Gene Ontology: http://www.geneontology.org/index.shtml

[11] Data and citations are from OPENCYC 1.0.2: http://www.opencyc.com

tions on this concept are "intrinsic to the [artefact] concept's nature and cannot be violated in any context." That is, no exceptions are possible, not even within other microtheories.

The top class is actually called `Artifact-Generic` and is described as "a collection of things created by agents" where an `Agent-Generic` is a "being that has desires or intentions, and the ability to act on those desires or intentions" (it includes social organizations like legal corporations and animals). Elements of `Artifact-Generic` like a hammer or a bird nest can be tangible or alternatively intangible like a set of laws. The category `Artifact-Generic` breaks down into `Artifact` and `Artifact-Intangible`. If we ignore the latter (which collects entities like computer languages and legal agreements), an element of `Artifact` may be said to be an inanimate thing which is "at least partially tangible" and "intentionally created by an agent (or group of agents working together) to serve some purpose or perform some function." The result of an assembling operation or of a modification of existing matter may not be an artefact unless the creating agent performs it intentionally and with a purpose.

In OPENCYC some amounts of matter are classified as artefacts under the class `ArtificialMaterial`, a subclass of `Artifact`. According to the informal description, this class contains "portion[s] of artificial stuff that was intentionally made by some agent(s), such as Plastic..." but excludes the byproducts of such activities. Note that the top category `Artifact-Generic` has a second (orthogonal) partition. The subcategories here are `Artifact-NonAgentive` and `Artifact-Agentive`. The first class collects artefacts which are not agents, like bicycles whilst the latter collects agents which are themselves created by agents, such as organizations. Unfortunately the OPENCYC documentation gives no information on the underlying view and no link with the specialized literature is provided. Some rationalisations and underlying intuitions are to be inferred from the category descriptions, when provided, and the formalization itself (for which the licence is needed).

Finally, the hierarchy of concepts below the `Artifact` category seems to be more driven by application goals (like the need to have a detailed and broad coverage of concepts of specific domains) than by ontological factors. Otherwise, it seems difficult to justify the presence (at the same level in the hierarchy) of `Artifact` subcategories like `InstrumentalArtifact` ("A sub-collection of `Artifact`. Each instance is an artifact (or system of artifacts) that is instrumental in accomplishing some end."), `ItalianCusine` ("the collection of instances of what many Americans tend to think of as Italian food"), and `StuffedToy` (no description provided).

3.2 Artefacts in SUMO

In SUMO[12], an artefact is described as a "`CorpuscularObject` that is the product of a `Making`" where an element of the `CorpuscularObject` is a "`SelfConnectedObject` whose parts have properties that are not shared by the whole". These descriptions

[12]Data and citations derive from the SUMO webpage: http://www.ontologyportal.org/ (Sept. 2007).

do not provide a clear view because SUMO adopts a very general notion of property which means that if one has two entities (e.g., an object and one of its proper parts) it will seem possible to find properties with which to distinguish them (for instance, properties relative to size). The description provided does not clarify this issue and precisely which objects are to be effectively considered artefacts in SUMO thus remains unclear. From the above definitions, it must also follow that SUMO artefacts are located in space-time and are self-connected. In other words, SUMO artefacts are non-scattered and physical. This confinement to physical entities is adopted in several discussions even in the philosophical literature but the constraint on self-connected objects is new and not explicitly justified: a jigsaw puzzle would count as an artefact in SUMO only when assembled. Another peculiarity resides in the description of the category `Making` seen above and said to characterize artefacts. `Making` is described as follows: "The subclass of `Creation` in which an individual `Artifact` or a type of `Artifact` is made". In turn, `Creation` is said to be "the subclass of `Process` in which something is created". These descriptions are hopelessly circular, they do not help us to understand the intended meaning of the categories. When tracing the links between categories, we did not find a direct relationship between the part of the hierarchy containing `Making` and the category `Agent`. This is surprising because in the relevant literature the direct and strong connection which exists between artefacts and agents is crucial to the ontological status of artefacts.

From the information collected, we can conclude that the class `Artifact` in SUMO captures the notion of physical, self-connected, and made (created) objects. Since the terms and relations used in SUMO are poorly characterized and documented, the true extent of this class of SUMO remains obscure.

3.3 Artefacts and the other ontologies

The foundational ontologies BFO, DOLCE and GFO do not introduce artefacts in their hierarchy. However, two of them seem to be in a position to provide a definition for the category.

From our previous description, it can be concluded that BFO has a limited number of categories and thus few expressive tools to introduce artefacts. If it is true that the ontology has a minimal notion of function, the lack of axiomatization and the limited number of classes makes the formalization of a category of artefacts problematic: one should introduce several preliminary notions simultaneously. We are not aware of any attempt to extend BFO with a category of artefacts.

DOLCE does not provide a notion of function but it has a very expressive framework to deal with qualities. However, some categories in the hierarchy are not fully formalized while some of these (e.g. `Social Object`) are crucial to model artefacts. To our knowledge, there has been no attempt to extend the ontology in this direction.

GFO has carried out an interesting study on the notion of function and it has a fairly rich hierarchy that may provide the tools to define the category of artefacts or, at least, a generalization of it in terms of functionalities. The developers of GFO have some idea of how to tackle the artefact category[13] but the ontology is still being developed and no extension of this is expected in the near future.

Since artefact is a notion that has direct consequences for applications one might assume that expanding our analysis to include non-foundational ontologies would lead to an interesting characterization of this notion. However, ontologies developed within certain application domains only rarely introduce categories of artefacts (but one often finds a subcategory for products). Even then their characterization of artefact (or of product, for that matter) is of little or no help. These ontologies are poorly characterized and the descriptions of their categories show that they are based on practical considerations and rely on the implicit knowledge shared in the community they address. Their proposals are therefore only valid when applied to their particular domain but remain, even in these cases, quite minimal.

4 THE DOLCE FOUNDATIONAL ONTOLOGY

Our next step is to elaborate a formal notion of artefact that is philosophically motivated. To surmount the problems shown above we will make explicit our modeling choices while also relating our approach to major philosophical positions in this area. This analysis will, we hope, be widely applicable. We want it to go beyond a philosophical discussion though. We also want to establish a formal characterization based on a specific formal framework. In the following pages we will therefore try to establish a balance between achieving a general analysis of the notion of artefact and recognizing the constraints imposed by the language and ontological choices behind the specific ontology used.

If we leave aside OPENCYC and SUMO which, as pointed out in their respective descriptions, are not proper foundational ontologies, we can choose between BFO, DOLCE and GFO. These systems all seem quite promising but each has its own particular drawbacks. If we bear in mind that BFO is only partly axiomatized and provides only a few categories, and that GFO is still working out the formalization of its new version, then it would seem that we would be better off working with the DOLCE ontology. A positive feature of this latter ontology is its rich and flexible framework for modeling qualities which provides an interesting theoretical tool for the capturing of formal and practical distinctions. It will therefore be exploited extensively in our work. On the negative side, the ontology only focuses on particulars (individuals), as it will become clear below.

4.1 An introduction to DOLCE

The Descriptive Ontology for Linguistic and Cognitive Engineering, DOLCE [Masolo et al., 2003] (www.loa-cnr.it/DOLCE), concentrates on particulars, that is,

[13]H. Herre, personal communication.

endurants, perdurants, qualities and abstract entities. It does not attempt to provide a taxonomy of properties and relations which are only included in the system if deemed crucial for characterizing particulars.

We mentioned above that DOLCE adopts a *multiplicative approach*: it assumes that different entities can be co-located in the same space-time. For example, a car and its matter are captured in DOLCE as two distinct entities (as opposed to being different aspects of the same entity). The reason for this lies in the different sets of properties that these entities enjoy: the car ceases to exist if a radical change of shape occurs (e.g., when it is crushed and cannot be repaired) while the amount of matter is not affected by this type of change (changes in spatial properties, like in shape or connectedness, is irrelevant for the identity of an amount of matter; only mereological properties can affect it). Going back to the classical example of the statue made of clay, for example, DOLCE might be said to model the statue and the amount of clay as different entities which share the same spatial (and possibly even temporal) location; the amount of clay used constitutes the statue. This allows the user to capture the strong intuition that a scratched statue is different (since it is scratched) while still remaining the same statue that it was before. In DOLCE this is possible because the identity of the statue itself might not be affected by minor scratches, but the identity of the clay is because scratches are the result of parts of the clay breaking off.

The category `Endurant` collects objects like cars and bits of matter like steel blocks, while events like the making of this car and the moving of that steel block fall into the category of `Perdurant`. The term "object" itself is used in the ontology to capture a notion of unity or wholeness as suggested by the partition of the category `Physical Endurant` into `Amount of Matter` whose elements are (amounts of) gold, air, etc.; `Feature` (a hole, a boundary); and `Physical Object` (a hammer, a human body). See Figure 1. Some of the categories are informally described in Section 4.2.

Every subcategory of `Endurant` and `Perdurant` is associated with a group of *qualities*. Qualities and their values (qualia) form distinct categories of entities in DOLCE and the distinction between *individual quality*, *quale*, and *quality space* has been established in order to capture several common sense intuitions in a coherent and consistent way.

Individual qualities, like the colour of this pen, *inhere in* specific individuals meaning that the colour of this pen is different from the colour of that pen no matter how similar the two pens may be. These qualities can change over the course of time since the colour of this pen can match the colour red today and the colour pink tomorrow. In contrast to individual qualities, qualia are not entity dependent. An example of a quale is a specific colour, like, for instance, red. Intuitively, these entities are obtained by abstracting individual qualities from time and from their hosts (see the discussion on tropes in Section 2.2). If the colour of this pen and the colour of that pen match the same shade of red, then they have the *same* (colour-)quale. In this respect, qualia represent perfect and *objective* similarities between (aspects of) objects. Quality spaces correspond to

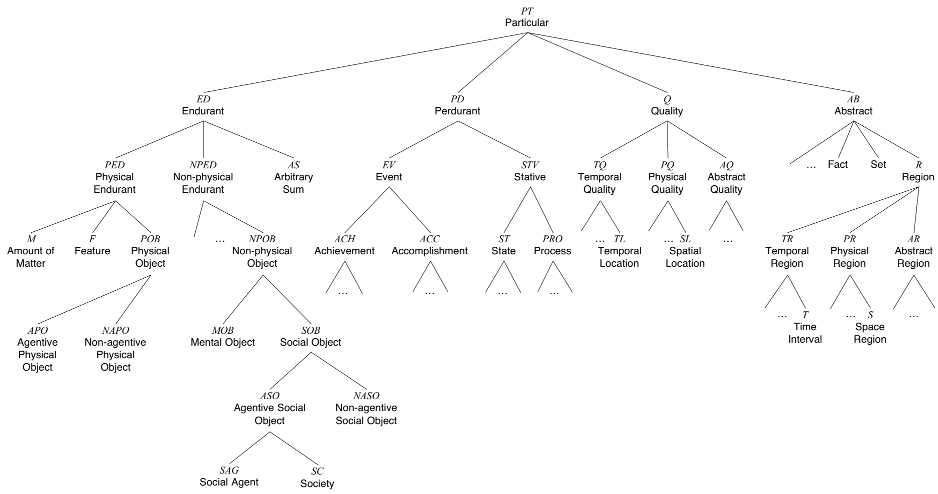

Figure 1. Taxonomy of DOLCE basic categories. (From [Masolo et al., 2003])

different ways of arranging qualia. They are motivated by similarities between objects or aspects of objects. By allowing different spaces for the same quality, different structures can be imposed on qualia (for example, a geometry, a metric, or just a topology) and this makes it possible to differentiate several quantitative and qualitative degrees of similarity (consider, for instance, the different ways of classifying and measuring colours).

The actual list of qualities associated with an entity depends on the user. Standard examples of qualities are shape and weight (usually taken to be qualities of endurants) and duration and direction (which are usually qualities of perdurants). However, these examples are not enforced by the ontology itself which is indeed neutral on the topic.

4.2 Some categories and relations in DOLCE

Several of the categories given in Figure 1 will be used to characterize artefacts. Here we shall just consider a few of them (including their relations) by way of informal introduction to the DOLCE terminology. Their formal names as used in the next sections are given in italicised parenthesis. The interested reader can find in [Masolo et al., 2003] the formal system together with a more detailed discussion.

Let us first recall the general category `Endurant` (*ED*) which collects entities that are *wholly* present at any time when they are present like, for instance, Bush, the first car built by Ferrari and the steel of the Eiffel tower. The elements of `Physical Endurant` (*PED*) are the endurants located in space-time, e.g., Gandhi's glasses as opposed to La Divina Commedia poem. This latter entity is classified as a `Non-physical Endurant` ($NPED$). `Amount of Matter` (M), e.g., some oxygen, `Feature` (F), e.g., a curve, and `Physical Object` (POB), e.g., a car have already been mentioned. Regarding agency, `Non-agentive Physical Object` ($NAPO$) pertains to the physical objects to which one cannot ascribe intentions, beliefs or desires (like a product or a ticket). A person falls into the category `Agentive Physical Object` (APO) which is different from the category of social entities, called `Social Object` (SOB), where we find things like organizations, companies, and their institutional artefacts such as constitutions and cheques. Social systems, such as a linguistic community, the people of a village or western society, form a subcategory denoted as `Society` (SC).

Entities that happen in time fall into the `Perdurant` (PD) category. In informal terms a perdurant is an entity that is only partially present whenever it is present. In this category we find happenings like football games and productions. These entities have temporal parts (like the first half of the game) as well as spatial parts, i.e. parts that are spatially identified (like the event being restricted to half of the football field during the game). Note that endurants are not parts of perdurants but that they *participate* in them instead (this relation is labelled PC). Some perdurants (like finishing a race or reaching the top of a mountain) are further classified in the subcategory `Achievement` (ACH). They are distinguished according to two properties: they have no temporal parts (e.g. instantaneous events) and their type is not preserved by sum: if we add together two consecutive events consisting in, say, finishing a book we get a new complex event which does not add up to the finishing of a book. Contrast this with events like drilling or walking: if we add together two consecutive walking events we still have a (possibly complex) walking event. Perdurants that have temporal parts but behave in the same way regarding their sum, like football games, fall into the `Accomplishment` (ACC) category. Note the distinction between finishing a book (an achievement) and reading a book (an accomplishment). Achievements and accomplishments form the category of eventive perdurants (EV).

Entities of a different type are found in the `Quality` category (Q), which covers all individual qualities. As discussed above, individual qualities can be seen as instantiations of basic properties of endurants or perdurants (shape, weight, duration, electric charge; usually qualities can be perceived or measured). The term "individual" is used to mark the fundamental role of the inherence relationship between an entity and its own qualities. Every endurant (or perdurant) comes with its physical (or temporal) qualities. Note that qualities are particulars in DOLCE that are not to be confused with properties (universals).

Expression $qt(q, x)$ stands for "q is an individual quality of x". Qualities are associated with quality spaces and the position an individual quality has in a space is called a quale. We write $ql(r, q, t)$ to indicate that "r is the quale of the endurant's quality q during time t" while $ql_T(t, x)$ stands for "t is the quale of the temporal quality of x". (Note the temporal parameter in $ql(r, q, t)$. If we want to evaluate "John is 5 feet tall", we have to be explicit when this sentence is stated as John's height changes over the course of time. Instead, relation $ql_T(t, x)$ describes temporal location and it is used to formalize, e.g., "the party last Sunday lasted from sunrise to sunset".) Each quale informally identifies a class of equivalence with respect to some individual quality, that is, with respect to an aspect of the entities. For instance, the same weight quale is associated with all the weight individual qualities that are ontologically indistinguishable (i.e., independently of any measuring instrument we have). From the remaining relations we will make use of *parthood* as in "x is part of y", written as $P(x, y)$. The relationship *being present* states when an entity exists in the world, thus one writes $PRE(x, t)$ to mean that "x is present in the world at time t". Earlier, we mentioned *participation*: expression $PC(x, y, t)$ stands for "endurant x participates in perdurant y during time t". *Constitution*, another crucial relation in DOLCE denotes a strong form of dependence: $K(x, y, t)$ stands for "x constitutes y during t". That is the relationship that holds between an amount of matter and a statue so that the statue cannot be present unless the material it is made of is also present.

Finally, we will use an extension of DOLCE proposed in [Masolo et al., 2004], and consider the category of *concepts* (CN) together with the relationship *classification*, written CF, that relates concepts and their "instances" at any one time. One writes $CF(x, y, t)$ to state that "at time t, x satisfies the concept y". Concepts are not standard universals in the sense of being individuals that depend on agents who create them or possibly on societies that adopt them. To account for their dependence, concepts are classified as particulars in this extension of DOLCE. Above all else, they are endurants, not abstracts, since they exist in time. Creating a concept means among other things providing a definition for it; the satisfaction of a concept is characterized by the constraints stated in the *description* defining a concept. For instance, the concept of Italian President has been created and defined by the Italian constitution which has been adopted by the Italian people.

5 EXTENDING DOLCE TO ARTEFACTS

We now turn to examining how the formal ontology DOLCE can be extended to include a category of artefacts. As explained above, this endeavour aims at showing that philosophical findings in this little explored domain can be successfully incorporated into an axiomatic first-order theory. This is not to say that we provide a general definition of artefact *tout court*. Indeed, in the literature the term "artefact" has been associated with a variety of meanings depending on the research domain and on the specific viewpoint of the authors. Here, we formally develop a coherent view that is compatible with the basic DOLCE choices.

5.1 Artefacts in the taxonomy

Most authors acknowledge that the notion of artefact seems to cover entities in a large variety of basic categories. Following the DOLCE taxonomy of basic categories depicted in Fig.1, it can be easily argued that artefacts may be either endurants (bottles and laws) or perdurants (judgements, performances and wars).[14] Endurants can be physical (bottles, glass and robots) or non-physical (pieces of music, laws and social institutions). For some authors, the whole `Social Object` category appears to fall under the larger class of non-physical artefacts [Searle, 1995]. Included in physical endurants are amounts of matter (pieces of glass or plastic), physical objects, which can be non-agentive (bottles, pens and paperweights) as well as agentive (robots, and arguably, bred animals and perhaps intended babies), and features (folds in a skirt, tunnels in mountains).

So it would seem that artefacts are not a separate category in the ontology, but rather a class of entities overlapping a variety of categories. As we will see below though, the *identity criteria* for artefacts, that is to say, their intentional nature, force us to regard artefacts as entities which actually are in a separate category.

As a starting point for providing a more general notion, we will only focus here on a subclass of the larger category of artefacts. This paper considers artefacts that correspond to physical endurants only, and among them, amounts of matter and non-agentive physical objects only.[15] For this first step, we will simply add to the DOLCE categories the category of `Physical Artefact` which falls under the category of `Physical Endurant` as a new sibling of `Amount of Matter`, `Physical Object`, and `Feature`. This category contains the most prototypical artefacts (e.g., tools like knives and pens) or, in other words, the least controversial ones so that we can be confident that it is covered by any specific view on artefacts. It arguably is the most studied category in the literature [Baker, 2004; Kroes and Meijers, 2006; Elder, 2007; Thomasson, 2007]. It is also the easiest to grasp in an ontology that is particularly well developed in the domain of material entities and the related fundamental relations, as is DOLCE.

The category of artefacts considered is quite big and presents a variety of interesting subclasses. We will not go into more specialized notions though; that is, we are not going to provide definitions to distinguish, for instance, “technical artefacts” or “works of art” within this class. These are crucial subclasses but a justification and presentation of their specific distinctions would be too detailed for the purposes of the present chapter.

Before proceeding further, we should however make it clear that we are dealing primarily with specific *tokens*, e.g., with the telephone that sits on Mary’s desk, and not with artefact types like *the* telephone. This implies that we are ignoring here the important process of designing (possible or impossible) artefacts, a process that often precedes the actual creation of any technical artefact token. The focus on tokens is natural within the DOLCE framework since, as pointed out in Section 4, this ontology is about particulars. Nonetheless, one sees that from the formalization, a notion of artefact type does emerge. We shall introduce and dis-

[14]There are no abstract artefacts because according to DOLCE, all abstract entities are out of time (cf. p.7). All artefacts are created, even non-physical ones like logical theories or novels. They all have a creation time before which they don’t exist and after which they do, i.e., they are in time, and thus non abstract.

[15]We actually even exclude living entities from `Non-agentive Physical Objects` (e.g., plants, viruses), but the remaining subclass is not identified as a category in DOLCE at present.

cuss this notion in Section 7.1. Focusing on token artefacts in an ontology theory means being concerned with the nature of these objects, or in other words with the essential properties that make the difference between artefacts and non-artefacts, and with the relationships between artefacts and other entities. As stated earlier, we are not concerned with the epistemological processes of recognition and categorization by an agent in the presence of a new object which may possibly be an artefact, nor by the process of designing a new artefact for a given purpose.

5.2 The approach

In accordance with the quite limited, although recently significantly increasing, literature on the ontology of artefacts [Dipert, 1993; Baker, 2004; Houkes and Meijers, 2006; Kroes and Meijers, 2006; Elder, 2007; Thomasson, 2007], our approach seeks to do justice to the complex nature of artefacts, which blends a physical substrate (regarding the category of physical artefacts we are concerned with here) with intentional aspects as well as social aspects.

We adopt the view that artefacts have an *ontological status*, in other words that they are full citizens of our ontology, and that artefacts are essentially the result of an *intentional act* of their creator. As we shall see below, two entities are therefore distinguished, the purely physical object or amount of matter that constitutes the artefact and the artefact itself, which emerges when it is created with both its physical and intention-based properties. Among these artefacts, generated by the private intention of their creators, we can further distinguish *social artefacts* (or artefacts proper in Dipert's terminology [Dipert, 1993]), whose intention-based properties take on a social dimension, more specifically, they are artefacts which are recognized as such by members of a specific society, e.g., cars.[16]

In accounting for these complex aspects, we will insist on developing the minimal formal apparatus required to characterize artefacts. For instance, although we need to model intentional aspects, we will avoid as much as possible the direct reference to theories of mental attitudes, which are not *per se* the subject of this paper and are by no means consensual. Similarly, we will not deal directly with the extremely elusive and much debated notion of purpose or function but will refer instead to a space of *capacities*, by making use of a general formal tool for describing the qualities of entities. We will not analyze in this paper the space of capacities itself: the structure of such space and the relationship between capacities and functions are issues that warrant further analysis. Here we shall merely introduce part of the formal machinery.

[16]The social character of artefacts is intended here in the broader sense. Social artefacts are not confined to the more restricted class of the physical artefacts that have a marked social purpose and use identity (like money and schools) nor to non-physical artefacts which fall in the category of social objects (e.g., laws and organizations).

6 MODELING ARTEFACTS

6.1 Intentional selection

Let us consider first the matter of agent intention which underlies the creation of an artefact. Two aspects need to be distinguished: the intention of obtaining an entity with some desired property (that perhaps makes it suitable for a particular purpose) and the intention of physically modifying or processing some pre-existing entity or entities in order to produce the artefact. We focus here on the first one only because we do not see the action of physical modification as an essential aspect of the creation of an artefact. In other words, artefacts do not need to be artificial entities: a pebble can make a paperweight, and a fallen tree a bench.[17] In addition, *residues*, such as sawdust, are intentionally processed but not intentionally selected for having certain properties and use. Artefacts are, *in nuce*, created entities in which "created" refers to a mental event, not to a physical modification.

One might remark that we do not base artefacts on actual use. The paperweight exists as soon as the agent selects it on the beach, and not just at the moment when he places it on a pile of loose papers. Similarly, knives coming out of a factory already exist as such; they are sold as knives, even though they have never at that stage cut anything. This approach therefore departs somewhat from Dipert's views [Dipert, 1993]. Our artefacts are what Dipert terms "contemplated instruments", in that they may still lack a use to be called "instruments". In addition, as just explained, we do not restrict artefacts to physically modified entities, so our artefacts may not be termed "tools" (intentionally modified instruments) in Dipert's terminology.

Contemplated instruments, and even only once-used tools, may be far too many for some people's tastes, but we believe this is not really an issue as one could define and focus on a *relevant* subclass of artefacts, "stable artefacts", based on the creator's repeated use according to her original purpose and/or her maintaining of some specific mental attitudes (memory of the creation, intention to use in the future), regardless of one of the remaining subclasses, the multitude of one-time-used or even one-time-contemplated artefacts. The class of social artefacts is another relevant subclass which will be described in this paper, in Section 7.2. This class arguably overlaps the class of stable artefacts, though one cannot ignore the matter of the many manufactured tools lying, yet unused, in stores and warehouses.

So the essence of any artefact lies in the creator's intention. It is certainly possible to explicitly represent the intentions or goals of the creator in an ontology of mental attitudes [Ferrario and Oltramari, 2004] and to reason about them adopting some dedicated logical formalism, for instance a so-called "Belief, Desire, Intention" logic [Rao and Georgeff, 1991]. As indicated above, for reasons of brevity we prefer to focus on the intentional creation event only and on the

[17]Even though one could argue that transporting the pebble to one's desk forms some kind of modification, the tree trunk can come to serve as a bench in the very place where it fell. It may furthermore be argued that physical modification has to be restricted to change in intrinsic physical properties, thus disregarding spatial location [Geach, 1969].

product of such an event, the artefact itself.

The paperweight is the result of some agent *intentionally selecting* a pebble and *attributing* to it certain *capacities*. The artefact itself is the new entity whose physical realization is the selected object and which has *attributed capacities*. In particular, the paperweight is a selected pebble together with the attributed capacity to stand firm and hold down paper without damaging it. Ultimately, the artefact might prove not to have the capacities the agent attributed to it, as it could be flawed or malfunctioning. More will be stated on this point later.

6.2 Constitution

We then suppose that the paperweight *is not* the pebble. The paperweight only starts to exist when it is created, usually well after the pebble has come into existence; the two objects, although co-located when both present, may have different lifetimes and are therefore different. The alternative solution which would be to consider artefactuality as a property that physical endurants may or may not have would avoid the multiplication of entities. But, by not granting artefacts an ontological status, it would not do justice to the notion of creation, intended as a notion distinct from physical modification, nor to the common view that artefacts are sortals [Elder, 2007]. We are thus confronted with two co-located entities, the artefact and the underlying physical object. In addition, the former is dependent on the latter, as the paperweight cannot exist without the pebble. In short, the pebble *constitutes* the paperweight [Rea, 1997].

The same physical object can constitute two different artefacts, for example the same pebble can constitute both a paperweight and a pestle. Only physical objects or amounts of matter may constitute an artefact, as it is only material artefacts that we are here considering here. So, when an artefact is apparently selected from another artefact, e.g., when a coffee-grinder is chosen to be used as a spice-grinder [Scheele, 2005], it is in fact the physical object constituting the first artefact which is selected again.

Although we do not dwell here on the special cases of artefacts constituted by aggregates and those which are copies of previously existing models, we agree with Baker and Elder [Baker, 2004; Elder, 2007] that constitution is, in this instance, a powerful tool. As pointed out above, DOLCE already adopts the corresponding multiplicative approach, in particular to distinguish the statue from the amount of matter that constitutes it. However, in this extension, it is important to note that what directly constitutes the paperweight here is the physical object pebble, and not simply the amount of (rock) matter that in turn constitutes the pebble.[18] The pebble is not an amount of rock because it is shape-dependent: the amount of rock persists after crushing, but the pebble does not —we obtain small stones or sand grains. Artefacts therefore bring yet another layer, an intentional level, to the constitution hierarchy. As a result, since the statue is an artefact, we actually need to distinguish three co-located entities, and not simply two as argued in DOLCE

[18]The amount of matter also constitutes the artefact, as constitution is transitive in DOLCE.

and more generally in the literature on material constitution: the intentionally created statue, the specifically-shaped and structured physical object, and the mereologically determined amount of matter.[19] We will see below that this further stratification is useful in understanding what happens when an artefact is repaired.

6.3 Capacity

What, then, are capacities? Our notion of capacity builds on Cummins's work on functions [Cummins, 1975]. His behavior-based approach avoids both the etiological account of function often given in philosophy of biology theories and the intentional approach adequate only for artefacts. We do take into account the intention of the agent in the creation event but we characterize, as does Cummins, the function of the artefact in agent-independent ways. To this end, we use the notion of quality in DOLCE, and assume that all physical endurants, restricted here to the categories `Amount of Matter` (M), and `Non-agentive Physical Object` ($NAPO$), have a single individual quality named *capacity* that characterizes all the capacities the physical endurant has. The capacity of an entity is an individual, just as is its colour. This quality maps into a quale that is a region (possibly a sum of atomic qualia) in the *capacity space*, which can be seen as some sort of functional *conceptual space* [Gärdenfors, 2000]. The quale corresponding to the capacity of an entity at a given time collects all the various *dispositions* [McLaughlin, 1995; Mumford, 1998] or behaviours the entity is able to express at that time. For instance, the capacity of this pen now has the quale of writing finely in black when drawn over paper, fitting in one's hand when grasped, and making a certain noise when it contacts the table. The capacity space is certainly complex, possibly founded on more elementary spaces of quality that provide the "bases" to such dispositions. We also assume that this space extends beyond standard (conditional) dispositions to also include structural properties, that is, the internal arrangement of parts. Our purpose here is not to analyse and describe this space in detail but to give the overall architecture of a possible formal ontological view of artefacts. Further study will certainly be required if we are to understand the structure of capacity space. It may in particular assess the need to use several distinct such spaces instead of a single one, and accordingly, several capacity qualities instead of a single one.

6.4 Attributed capacity

In addition to the capacity possessed by any physical endurant,[20] artefacts also have an *attributed capacity*, another quality associated with qualia in the *same* space. The fact that actual dispositions and intended functions are elements of a same space has a number of advantages. It first of all permits the definition of

[19] When the artefact is not selected from a physical object but from an amount of matter, as with a piece of glass, there are of course two layers only.

[20] We assume that an artefact's capacity is inherited from its constituting entity.

malfunctioning (see Section 7.3 below). It furthermore demonstrates how the dual nature of artefacts that are physical endurants can be reconciled as it forms the interface for the mental and the physical aspects of artefacts [Kroes and Meijers, 2006]. This unique space also makes possible a future account of the design process. As the capacity space also covers structural properties, the very design plans of technical artefacts could be conceived as part of the attributed capacity.

Although capacity and attributed capacity map into the same space of qualia, the former remains a physical quality whereas the latter is an intentional quality as it is dependent on the intentions of the creator at the time of the creation event.[21] Capacity and attributed capacity also differ in the following way: the quale associated with the attributed capacity does not change in time as it is fixed by the creation event; moreover, this quale is a *set* of regions of the capacity space because the intended behaviour of the artefact need not be specified in a precise way, and may present vagueness. For instance, when one is looking for something to write on a board and one selects a piece of coal the value of the attributed capacity is only to write on the board and fit in the hand. Therefore, the attributed capacity maps into the qualia space in a variety of alternative regions corresponding to different possible writing behaviours, e.g., writing finely in black on the board and fitting in one's hand, writing thickly in black on the board and fitting in one's hand, writing finely in brown on the board and fitting in one's hand, etc. On the other hand, the creator of a commercial pen has probably designed it precisely and has therefore chosen an attributed capacity which maps to a reduced number of regions or even to a unique one. This region, though, is *a priori* smaller than the capacity's quale region, e.g., the creator of the pen did not design it for the noise it produces when crushed with a rock. So typically the quale of the artefact's capacity strictly includes one member of the quale of its attributed capacity. This obviously does not hold in the case of malfunctioning or faulty artefacts: one can create an artefact with an attributed capacity's quale that the selected physical endurant's capacity (and the artefact's capacity) will not always have or may not have from the start and perhaps ever (see Section 7.3).

6.5 Identity criteria

If we are to grant an ontological status to artefacts, a delicate point now needs to be addressed. We need to examine their identity criteria. We have seen that artefacts are distinct from the physical objects (or amounts of matter, in the case of artefactual matter) that constitute them. They should therefore have distinct identity criteria. Indeed, artefacts can be repaired and have some parts substituted, thus changing the entity that constitutes them for another without losing their identity. Such change comes at the cost of the former constituting entity disappearing simultaneously with the newer constituting entity coming into existence, though maintaining a certain degree of spatio-temporal continuity between

[21] This dependence will not, however, be formally expressed here, as we deliberately refrained from introducing intentions.

the two. In fact, no artefact can "jump" from one material entity to a separate preexisting one at will. If Theseus's ship [Rea, 1997, Introduction], an artefact, does not disappear when a plank is substituted, the physical object that constitutes it, the planks-and-nails assembly, changes so that the former assembly ceases to exist and a new assembly comes into existence.[22]

By pointing out the property that an artefact cannot jump from one physical object to another, we can shed some light on the important distinction between artefacts and artefact *roles*. Roles, in general, can be played by different entities (e.g., different persons at different times can play the role of president of the US) [Masolo et al., 2004] and the change between players can be seen as a "jump", as the previous player usually survives the change and the successor often already exists. Physical artefacts are more stable. They are not roles. This distinction is evident, for instance, in the house/home contrast. A house is an artefact which can play the role of being someone's home. One's home changes, there is a jump from a house to another when one moves house, so "home" is not a type of artefact subsumed by "house", but rather a role.

The gradual change in the constituting material entity may only occur with artefacts selected from physical objects and not with those selected from amounts of matter. It is reasonable to assume that pieces of plastic or of glass cannot switch over just as quantities of matter cannot interchange. Indeed, amounts of matter in DOLCE have purely mereological identity criteria.[23] Non-agentive physical objects have more complex identity criteria, which vary from sortal to sortal. It is not the purpose of this paper to establish those criteria, but as a general guideline, we will take shape and internal structure to be part of these criteria. We assume though that minor changes in shape and in the constituting amount of matter, like those induced by a scratch, are allowed. Granularity is certainly an issue here.

With artefacts, an obvious characteristic for determining their identity criteria is their intentional aspects, that is, their attributed capacity. The identity criteria should among other things determine when an artefact disappears all together. Ordinary malfunctioning does not make an artefact disappear, so its identity criteria cannot be simply based on a match between attributed capacity and capacity. Nor is the artefact's disappearance simply based on its constituting entity's disappearance, since that can be substituted, as we have just seen. So, the loss of much of the attributed capacity must be involved. We do not intend to solve here the infamous ship-of-Theseus puzzle [Rea, 1997, Introduction], but we believe that we can nevertheless safely assume that the identity criteria of artefacts are based on a combination of *significant* degree of spatio-temporal continuity of the constituting entities, the existence of all specific essential parts if any (e.g., for a car, its frame), and the actuality of a *significant* amount of attributed capacity, i.e., a significant overlap between one region member of the quale of the attributed capacity and

[22]The term "assembly" denotes here an aggregate in a specific arrangement.

[23]DOLCE does not take into account the nature of the substance of which the amount of matter is made. As a result, it does not consider homogeneity conditions. A different choice would not not affect the present discussion.

the region quale of the capacity. Note that since the attributed capacity is not restricted to the overall or main function of the artefact and since it covers structural specifications like size, shape, weight and composition, a malfunctioning artefact does possess most of its attributed capacity. Even a badly designed artefact, like a medieval flying machine, possesses most of its attributed capacity.

6.6 Axiomatics

We now turn to the formal theory that corresponds to the above choices.

A *physical artefact* or artefact for short is an element of `Physical Artefact` category. It is dependent on a constituting entity of category `Amount of Matter` (M) $\cup$ `Non-agentive Physical Object` ($NAPO$), called for short "material entity" in the remainder. For instance, an amount of glass is an artefact constituted by an entity of M category, while a paperweight is an artefact constituted by an entity of $NAPO$ category.

An artefact x is created by an intentional association of a material entity y and a quality q which is of the type $\mathsf{AttributedCap}$, a new primitive predicate denoting attributed capacities.[24] The intentional association that generates the artefact is a special event of type $\mathsf{CreationEv}$ (creation event). To characterize this we use the $\mathsf{IntentionalSel}$ (intentional selection) primitive relation which takes as its arguments an event (EV) e, an agent p, physical (APO) or not (ASO, e.g., a company), a physical artefact ($\mathsf{PhysArt}$) x, a material entity y, and a quality (Q) q. $\mathsf{IntentionalSel}(e, p, x, y, q)$ should be read as "e is the event of p obtaining the artefact x by intentionally selecting y and attributing to it capacity q." Our first axiom states that artefacts, that is, the elements of the category dubbed `Physical Artefact` and represented by the primitive predicate $\mathsf{PhysArt}$, are the product of some intentional selection event:

(A1) $\mathsf{PhysArt}(x) \leftrightarrow \exists e, p, y, q\ \mathsf{IntentionalSel}(e, p, x, y, q)$

Next we constrain the primitive $\mathsf{IntentionalSel}$ as indicated above:

(A2) $\mathsf{IntentionalSel}(e, p, x, y, q) \rightarrow EV(e) \wedge (APO(p) \vee ASO(p)) \wedge (M(y) \vee NAPO(y)) \wedge \mathsf{AttributedCap}(q) \wedge qt(q, x) \wedge \exists t\ (ql_T(t, e) \wedge PC(y, e, t) \wedge PC(x, e, t) \wedge PC(p, e, t) \wedge K(y, x, t)))$

Axiom (A2), in addition to restricting the arguments of `IntentionalSel`, specifies a number of assumptions. The quality q is a quality (qt) of the artefact x. The agent, the artefact and the material entity all participate (PC) in the selection event for the time of the event. One consequence is that these three entities are present (PRE), that is, exist, during the event. For non-instantaneous events, in other words accomplishments (ACC) as opposed to achievements (ACH), it is nevertheless somehow arbitrary to say that the artefact x exists during the event;

[24] The new —primitive or defined— predicates introduced to characterize artefacts are given in sans-serif font to distinguish them from the predicates denoting DOLCE categories and relations.

one could assume it exists only immediately after the event instead. This decision is not crucial to our approach. Lastly, the artefact x is constituted (K) by the material entity y during the selection event.

Axiom (A5) below will guarantee that this constitution relation lasts until the artefact disappears or until its constituting entity is replaced. The constituting entity y does not need to exist before e, since e could co-occur with a physical creation event, in which case both x and y would be created simultaneously.

As far as the existence persistence of the artefact after the creation event is concerned, this can be at best a default rule because nothing prevents its destruction. We assume though that while the original constituting entity is present, the artefact will also be present. This partially underlines the fact that for an artefact to disappear it needs to undergo a major change which cannot happen without altering the identity of the constituting entity. So, as long as the original constituting entity is present, the artefact will also be present. The opposite is not true, though, because the artefact can change its constituting entity as in the case of repairing by substituting a component.

While the artefact exists, there is a unique entity of the same category of the material entity originally selected, `Amount of Matter` or `Non-agentive Physical Object`, that constitutes it.[25] As explained above, if the selected material entity is an amount of matter, this cannot change:

(A3) $(\mathsf{IntentionalSel}(e,p,x,y,q) \wedge M(y) \wedge PRE(x,t)) \rightarrow (K(y,x,t) \wedge \forall z(\neg z = y \rightarrow \neg K(z,x,t)))$

(A4) $(\mathsf{IntentionalSel}(e,p,x,y,q) \wedge NAPO(y) \wedge PRE(x,t)) \rightarrow \exists! z\,(K(z,x,t) \wedge NAPO(z))$

While the constituting material entity exists (which may or may not be the original entity), it continually constitutes the artefact and so, as a consequence, the artefact still exists:

(A5) $(\mathsf{PhysArt}(x) \wedge K(y,x,t) \wedge (M(y) \rightarrow \exists e,p,q\ \mathsf{IntentionalSel}(e,p,x,y,q)) \wedge PRE(y,t') \wedge t < t') \rightarrow K(y,x,t')$[26]

A consequence of the previous axioms is that if the artefact is constituted by different physical objects at different times, these physical objects will not exist simultaneously (thus ruling out "jumps"):

$$(\mathsf{PhysArt}(x) \wedge K(y,x,t) \wedge NAPO(y) \wedge K(y',x,t') \wedge NAPO(y') \wedge \neg y = y') \rightarrow (\neg PRE(y',t) \wedge \neg PRE(y,t'))$$

[25]It is impossible to simply assert that the entity that constitutes an artefact is unique since an artefact constituted by a physical object is also constituted by the amount of matter that constitutes it. All physical objects are constituted by some amount of matter, and constitution is taken to be transitive in DOLCE. DOLCE assumes that `Amount of Matter` is the lowest substrate, in other words that nothing constitutes an amount of matter.

[26]This very partial account of continuity ignores the case of assemblies having an intermittent existence, as in the Theseus's ship puzzle, in which the original constituting physical object is reassembled. We leave this issue for further developments.

Since a material entity may constitute several different artefacts, the formula $(\mathsf{IntentionalSel}(e, p, x, y, q) \ \wedge \ \mathsf{IntentionalSel}(e', p', x', y, q')) \rightarrow x = x'$ is *neither* a theorem *nor* an axiom.

Further constraints are needed. First, for a given artefact, the attributed capacity is unique and the quale of the attributed capacity does not change in time:

(A6) $(\mathsf{IntentionalSel}(e, p, x, y, q) \ \wedge \ \mathsf{IntentionalSel}(e', p', x, y', q')) \rightarrow q = q'$

(A7) $(\mathsf{IntentionalSel}(e, p, x, y, q) \ \wedge \ ql(v, q, t) \ \wedge \ ql(v', q, t')) \rightarrow v = v'$

As asserted above, we assume that the same artefact can be selected several times, by possibly different agents (or societies) like, for example, the same "tree trunk bench" in the woods, so the event and the selector are not necessarily unique. This choice is not essential to the approach.

However, for a given intentional selection event, the artefact and the selector must be unique (as well as the attributed capacity quality because of the axiom above):

(A8) $(\mathsf{IntentionalSel}(e, p, x, y, q) \wedge \mathsf{IntentionalSel}(e, p', x', y', q')) \rightarrow (x = x' \wedge p = p')$

We also make sure that attributed capacities are only qualities of artefacts:[27]

(A9) $(\mathsf{AttributedCap}(q) \ \wedge \ qt(q, x)) \rightarrow \mathsf{PhysArt}(x)$

On the other hand, all physical endurants, including artefacts, have a capacity (A10). The capacity of physical endurants is inherited through constitution, in the sense that the quale of the capacity of the constituted entity includes that of the constituting entity (A11). Note that the two qualia need not be identical. The capacity of a physical object may include shape-based dispositions, while the capacity of the amount of matter constituting it cannot. Similarly, when creating a socially relevant artefact, like a cheque, new capacities are created that the constituent itself, i.e., the rectangular piece of paper, does not possess.

(A10) $PED(x) \rightarrow \exists q \, (\mathsf{Capacity}(q) \ \wedge \ qt(q, x))$

(A11) $(\mathsf{PhysArt}(x) \ \wedge \ K(y, x, t) \ \wedge \ \mathsf{Capacity}(q) \ \wedge \ qt(q, x) \ \wedge \ ql(v, q, t) \ \wedge \ \mathsf{Capacity}(q') \ \wedge \ qt(q', y) \ \wedge \ ql(v', q', t)) \rightarrow P(v', v)$

To ensure that capacities and attributed capacities map to qualia in the same space of capacities, we use a unary predicate CR, for capacity region, to characterize those qualia. However, we need to allow the attributed capacity of an artefact to have a *set* — or some sort of collection — of such regions for quale, as mentioned above. Sets, collections and aggregates are not yet formalized in DOLCE. We will nevertheless use a fairly intuitive primitive predicate IN to denote "membership" (its characterization does not concern us here):

(A12) $(\mathsf{Capacity}(q) \ \wedge \ ql(v, q, t)) \rightarrow \mathsf{CR}(v)$

[27] In DOLCEa quality inheres in a unique entity so given q there is a unique x such that $qt(q, x)$.

(A13) $(\mathsf{AttributedCap}(q) \ \wedge \ ql(v,q,t) \ \wedge \ \mathsf{IN}(x,v)) \rightarrow \mathsf{CR}(x)$

Finally, an intentional selection is a creation event unless the *same* artefact has already been selected. CreationEv is a defined predicate:

(D1) $\mathsf{CreationEv}(e,x) =_{def} \exists p,y,q \ \mathsf{IntentionalSel}(e,p,x,y,q) \ \wedge \ \exists t \ (qt_T(t,e) \ \wedge \ \forall t'(t' < t \rightarrow \neg PRE(x,t')))$

We ensure that there is a creation event for each artefact and, as a consequence, that it does not already exist before the first intentional selection event:

(A14) $\mathsf{PhysArt}(x) \rightarrow \exists e \ \mathsf{CreationEv}(e,x)$

The creator of an artefact is the first selector; it is unique as long as there are no simultaneous first intentional selections of the same artefact:

(D2) $\mathsf{Creator}(p,x) =_{def} \exists e,y,q(\mathsf{CreationEv}(e,x) \wedge \mathsf{IntentionalSel}(e,p,x,y,q))$

No additional axiom is introduced to model the conditions in which a given artefact disappears, essentially through lack of means to adequately model the difficult vagueness issues involved in the identity criteria of artefacts as described above.

7 BEYOND THE BASICS

7.1 Artefact Types

The approach developed so far allows us to characterize a notion of *artefact type* within DOLCE. Since agents and societies develop concepts to discriminate between types of entities, it is natural to view concepts about artefacts as providing definitions of artefact types. Artefacts are endurants with a particular quality known as attributed capacity. It suffices for the concept to discriminate between the attributed capacities of the artefacts to coherently collect artefacts "of the same type". We thus assume that the definition of a concept classifying artefacts, i.e., an artefact type, is based on the comparison of the qualia of these attributed capacities with the attributed capacity of what we would call a prototype. So, we suppose that the definition of an artefact type, say *Hammer*, isolates the collection of the artefacts whose attributed capacities are such that all the regions in their quale include one of those of a generic or prototypical hammer. *The* prototypical hammer does not need to exist, but there must be a specific attributed capacity value v, i.e., a set of capacity regions, that characterizes what would count as a prototypical hammer. Evidently, the existence of an artefact type, a concept, is independent of the creation of any artefact token of this type.

Formally, we use the *classification* (CF) relation. In Section 4.2 we asserted that $CF(x,y,t)$ stands for "at t, x satisfies concept y", which we sometimes write "x is classified by y at t".

(D3) $\mathsf{ArtefactType}(a) =_{def} CN(a) \wedge \exists v\,(\exists u\,\mathsf{IN}(u,v) \wedge \forall w(\mathsf{IN}(w,v) \rightarrow \mathsf{CR}(w)) \wedge \forall x,t\,(CF(x,a,t) \leftrightarrow \exists q_x, v_x\,(\mathsf{PhysArt}(x) \wedge \mathsf{AttributedCap}(q_x) \wedge qt(q_x,x) \wedge ql(v_x,q_x,t) \wedge PRE(x,t) \wedge \forall w(\mathsf{IN}(w,v_x) \rightarrow \exists w'(\mathsf{IN}(w',v) \wedge P(w',w))))))$

This definition states that all and only the artefacts classified by concept a have an attributed capacity's quale whose member regions all include a region of the (non-empty) set of regions v characterizing a. Note that the artefacts of a given type do not need to be present at the same time, since, as we have seen, the quale of an attributed capacity cannot change in terms of time. For the same reason, we obtain a theorem which claims that an artefact is classified by the same type or types throughout its existence:

$$\forall x,a,t,t'\,((\mathsf{ArtefactType}(a) \wedge CF(x,a,t) \wedge PRE(x,t')) \rightarrow CF(x,a,t'))$$

Clearly, using mereological relations between the elements of the sets of capacity regions used to characterize artefact types generates a network of types so that we can relate subtypes to types, e.g., *Carpentry hammer* or *Physician's hammer* to *Hammer*.

As for all other concepts, artefact types are dependent on the agents or societies that define them [Masolo et al., 2004]. In fact, artefact types often apply to social artefacts, a subject to which we can now turn.

7.2 Social Artefacts

Informally, a social artefact is an artefact whose type is *recognizable* by the members of a certain society. A once-used object such as the bench-and-table that you selected from a fallen tree for your last picnic in the woods is not what we commonly call an artefact, nor is it a stable artefact repeatedly but privately used, such as the spice-grinder that was selected from the physical object constituting a coffee-grinder [Scheele, 2005]. Often, artefacts are recognized as such by agents other than their creators: we buy knives assuming that someone has made them suitable for cutting when used in a certain way. Societies share the knowledge of recognizing many different artefacts, that is of recognizing part of the attributed capacity of a given entity through its type: pens and knives, glass and flour, and so on so that most of the time there is no need for the creator to explain purpose. As Dipert puts it, a *proper artefact* is in effect an entity for which the attributed capacity (Dipert calls it the creator's intention) is recognizable [Dipert, 1993]. To emphasize their dependence on a given society we call these items *social artefacts*. We depart somewhat from Dipert's proposal by requiring that only the part of the quale of the attributed capacity defining an artefact type, in other words, the attributed capacity value characterizing a prototype, be recognized. This makes it possible that certain non-central or idiosyncratic aspects of the attributed capacity to be ignored, avoiding the assumption that members of society have unrealistic mind-reading abilities.

A recognition event has to be distinguished from an intentional selection, as the intention (recognized in the event) is attributed to someone else, even if that someone is unknown. Of course, one may assume that someone who has attributed a capacity to this entity exists thereby believing that one recognizes an artefact, but one might just be wrong. This is an epistemological issue and as such it is not a subject for ontological analysis. Our aim is to account for the ontological import of the fact that a given society shares knowledge about some kinds of artefacts. As mentioned above, the formal ontology of social reality does not seek to analyze how and on what grounds an agent of society is able to fulfil the recognition of a given artefact. It has to deal with the fact that some artefacts acquire a social status. Note, though, that the *existence* of the artefact is not affected by having or not having the property of being recognizable: social artefacts are essentially artefacts. So general consensus among a community of archaeologists on the artefactual nature of a collection of entities that are actually only residues makes them neither simple nor social artefacts.[28]

The recognition of a social artefact relies on at least one of two distinct elements: the object structure and the context. In the first case, the act of recognition is intrinsically related to the material entity constituting the artefact — its structure, its physical qualities, its actual capacity, etc. — and to the structure and properties of artefacts previously encountered. The recognition of knives and cars falls into this category. In the second case, recognition is based on the broader context in which the entity is observed. For instance, if we see a pebble (of a certain size) on someone's desk on top of a heap of papers, we will assume that the pebble constitutes a paperweight, while generally we do not identify paperweights on the seashore (though we may intentionally select some). Similarly, if we see in a shop a pile of pebbles labelled "paperweight", we all assume the shop is actually selling paperweights. There are obvious limits to this: a label "paperweight" on a heap of sand will not be sufficient to make us recognize there artefacts because certain constraints on the capacity of the material entity need to be satisfied to convince us that someone did select that material entity and attributed it a certain capacity.

We thus distinguish between (intrinsic) social artefacts and contextual social artefacts. To ontologically represent such notions, we shall introduce the new primitive predicate $\mathsf{Recognizable}(a, x, s, t)$, that reads as "the type a of artefact x is recognizable by society s at time t". The basic constraints on this relation are:

(A15) $\mathsf{Recognizable}(a, x, s, t) \rightarrow (\mathsf{ArtefactType}(a) \wedge \mathsf{PhysArt}(x) \wedge CF(x, a, t) \wedge SC(s) \wedge PRE(s, t))$

To constrain $\mathsf{Recognizable}$ further is not an easy matter if one wants to avoid referring explicitly to the mental attitudes of agents. To gain a better grasp of what this predicate is supposed to mean, we can sketch what could be a definition, if we were using a theory that allows for doxastic modalities (the modal *belief* operator [Hintikka, 1964] $\mathsf{Bel}_{x,t}\, \phi$ standing for "x believes proposition ϕ at t"), as well as the arguably simpler primitives Member (between an agent (APO) and a

[28] The gap between ontology and epistemology is a particularly difficult one to bridge when there is societal discontinuity with knowledge loss, viz. the "Nineveh lens" cf. http://www.badarchaeology.net/data/ooparts/nineveh.php

society (SC) at some time) and Perceives (describing the event of an agent (APO) perceiving a physical endurant (PED)). The type a of artefact x is recognizable by society s at time t if the artefact and the society are present at t and any member[29] of the society s believes the artefact x is an artefact of type a whenever during t the agent perceives[30] it:

$$\mathsf{Recognizable}(a,x,s,t) =_{def} \mathsf{ArtefactType}(a) \wedge \mathsf{PhysArt}(x) \wedge CF(x,a,t) \wedge SC(s) \wedge PRE(s,t) \wedge \forall e,y,t' ((\mathsf{Member}(y,s,t') \wedge \mathsf{Perceives}(e,y,x) \wedge ql_T(t',e) \wedge P(t',t)) \rightarrow \mathsf{Bel}_{y,t'} (\mathsf{PhysArt}(x) \wedge CF(x,a,t')))$$

A(n) (intrinsic) social artefact for society s is then an artefact whose type is recognizable by society s at all times when the artefact and the society are present:

(D4) $\mathsf{SocialArt}(x,s) =_{def} \mathsf{PhysArt}(x) \wedge \exists a \, \forall t \, ((PRE(x,t) \wedge PRE(s,t)) \rightarrow \mathsf{Recognizable}(a,x,s,t))$

Let us now turn to contextual social artefacts. For the sake of presentation, contexts are taken here as a category of entities[31] and we introduce a new predicate InContext to relate an endurant (ED) to a context at a time, assuming that $\mathsf{InContext}(x,c,t)$ entails $PRE(x,t)$. With these tools, we can define a contextual social artefact (ContextualSocialArt) for context c and society s to be an artefact whose type is recognizable by s whenever the artefact is in the context c:

(D5) $\mathsf{ContextualSocialArt}(x,s,c) =_{def} \mathsf{PhysArt}(x) \wedge \exists a \, \forall t \, ((\mathsf{InContext}(x,c,t) \wedge PRE(s,t)) \rightarrow \mathsf{Recognizable}(a,x,s,t))$

7.3 Malfunction

As mentioned above, our notion of artefact includes malfunctioning or even failed artefacts. With our approach, it is rather straightforward to state that an artefact is malfunctioning at t. It simply does not possess all the capacities attributed to it:

(A16) $\mathsf{MalFunctioningArt}(x,t) =_{def} \mathsf{PhysArt}(x) \wedge PRE(x,t) \wedge \forall q,q',v,v',w \, ((\mathsf{AttributedCap}(q) \wedge qt(q,x) \wedge \mathsf{Capacity}(q') \wedge qt(q',x) \wedge ql(v,q,t) \wedge ql(v',q',t) \wedge \mathsf{IN}(w,v)) \rightarrow \neg P(w,v'))$

This definition is based on the fact that "possessing the attributed capacities" means that at least one of the alternative regions in the quale of the artefact's

[29] There is surely a need here to restrict the conditional antecedent to qualified members of the society and to thus disregard babies, drunk people and so on.

[30] The nature of the perceptive events involved (seeing, hearing, touching, etc.) may depend on both the artefact and the society; we can safely assume that perceiving an artefact is equivalent to perceiving the material entity that constitutes it.

[31] The real nature of contexts is by no means obvious and the very issue of their reification raises some criticism. Contexts are not currently included in DOLCE and we could manage without them by using "descriptions" as introduced in [Masolo et al., 2004]. However, as it would be impossible to introduce here the notion of description, we shall rely for the present purposes on the intuitive notion of context.

attributed capacity is part of the region corresponding to the current quale of the artefact's capacity. Note that we talk of the capacity of the artefact itself and not of the capacity of its constituent. As posited above in (A11), the capacity of the constituting entity is inherited by the artefact.

7.4 Limitations of the framework

We conclude this section by highlighting a couple of open-ended problems that we think should be addressed if we are to understand the advantages and the limitations of this framework.

First, our formalization is not compatible (as it stands here) with the intuition that an artefact may gain or lose attributed capacities (more precisely, that the regions in the quale of its attributed capacity might grow or shrink) while it remains the same artefact. We have instead exploited the multiplicative approach of the DOLCE ontology, assuming that the same material entity can simultaneously constitute different artefacts, e.g., the paperweight and the pestle (both private or contextual social), the (intrinsic social) coffee grinder and the (private) spice grinder, or the (intrinsic social) anvil and the (contextual social) doorstop. Still, it might be possible to adopt a notion of attributed capacity that depends on time and so render, even in this approach, the intuition of artefact evolution. We have not studied that option here.

Similarly, concepts like artefact types may evolve (a nice example relates to the evolution of Aspirin from painkiller to painkiller and blood-thinner and is detailed in [Houkes and Meijers, 2006]). The evolution of concepts has not been fully addressed in the theory developed in [Masolo et al., 2004] which is exploited here. It is implicitly assumed, however, that we need to distinguish each change in concept as a creation of a new concept historically dependent on the previous one. This appears to be a reasonable solution here too.

We have just noticed that the theory developed makes extensive use of the multiplicative approach of DOLCE, the formal ontology chosen here. As explained on p.9, this feature is rather a specificity of DOLCE though it is not incompatible with other foundational ontologies. On the other hand, the multiplicative approach developed here is related to the "constitution view" developed by Baker [Baker, 2004] which has been criticized by Houkes and Meijers in [Houkes and Meijers, 2006]. Let us then examine the reasons for rejecting such an approach.

Houkes and Meijers's first criticism of the constitution view is that this approach leads to unnecessary "ontological stacking". We note that the multiplicative approach has many other applications, giving for instance a straightforward answer to the puzzle of multiple event descriptions which have different causal explanation power [Pianesi and Varzi, 2000]. The multiplicative approach does not convince all philosophers, though. There are arguments in favour of and against the reductionist and multiplicative points of view, in all domains. We believe that the artefact domain is not essentially different in this respect.

Their second criticism concerns the impossibility of Baker's proposal to account for the *Realizability Constraints* (RC), the idea that an ontology of artefacts must provide grounds for constraining the possible entities that may constitute an artefact — given its function — as well as providing grounds to constrain the artefacts an entity could constitute, given its structure. In our opinion, the multiplicative approach does not in itself provide any explanation for this, nor does it prevent a further formal account of RC. Constitution is a generic dependence relation which constrains the existence of the related entity, but that does not *explain* why such a constitution may or may not hold. As mentioned above, although we have left the designing process out of the picture, we believe that RC could be accounted for by comparing the qualia of the actual and attributed capacities of an artefact, something which is facilitated by the use of a single capacity space. This is not at all straightforward, though. As argued above, the conditions of existence of an artefact are indeed based on a match between its actual and its attributed capacities but such a match is of necessity vague to allow for the existence of malfunctioning artefacts. And vagueness is a notoriously difficult issue, especially when ontological matters are at stake.

Finally, we point out once more that our notion of artefact relies on one important quality space, namely the capacity space. This space has not been analyzed here and it is not yet well understood. It includes both functional and structural aspects in an interesting setting that certainly does deserve more attention. Furthermore, it seems necessary to study the dimensions of this space, the relationship to the other quality spaces such as weight, shape and colour, and the overall structure if we are to properly formalize other crucial notions such as that of technical artefact.

8 CONCLUSION

In this chapter we analyzed ontology research on the notion of artefact. After looking at existing ontologies and highlighting several shortcomings, we presented and discussed a new formalization that defines artefacts to be endurants with a special quality known as attributed capacity, which justifies their special status with respect to other endurants. This new quality allowed us to formalize a series of notions which were justified on the basis of philosophical distinctions as well as commonsense intuitions.

The theory proposed, although not self-contained and still requiring further development, shows the feasibility of extending a foundational ontology, namely DOLCE, to grasp the non-trivial notion of artefact. The theory of course reflects certain philosophical choices which will not be palatable for all researchers in the field. Similarly, some of its technical aspects strongly rely on the multiplicative structure of DOLCE, a feature which might not be easily transposed in all other formal frameworks. Nevertheless, we believe that this work has much potential for concrete applications where artefacts are central and semantic integration is an issue. It also illustrates the fecundity of applying philosophical studies to knowledge representation in computer science.

BIBLIOGRAPHY

[Antoniou and van Harmelen, 2004] Antoniou, G. and van Harmelen, F. Web Ontology Language: OWL. In S. Staab and R. Studer, editors, *Handbook on Ontologies*, pages 67–92. Springer, 2004.

[Baker, 2004] Baker, L.R. The ontology of artifacts. *Philosophical explorations*, 7(2):99–111, 2004.

[Borgo et al., 1996] Borgo, S., Guarino, N. and Masolo, C. A Pointless Theory of Space based on Strong Connection and Congruence. In L. Carlucci Aiello, J. Doyle, and S. C. Shapiro, editors, *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR'96)*, pages 220–229. Morgan Kaufmann, 1996.

[Campbell, 1990] Campbell, K. *Abstract Particulars.* Basil Blackwell, Oxford, 1990.

[Cummins, 1975] Cummins, R. Functional analysis. *Journal of Philosophy*, 72:741–765, 1975.

[Dipert, 1993] Dipert, R. *Artifacts, Art Works, and Agency.* Temple University Press, Philadelphia, 1993.

[Elder, 2007] Elder, C. On the place of artifacts in ontology. In E. Margolis and S. Laurence, editors, *Creations of the Mind: Theories of Artifacts and Their Representation*, pages 33–51. Oxford University Press, Oxford, 2007.

[Fellbaum, 1998] Fellbaum, C. editor. *WordNet. An Electronic Lexical Database.* MIT Press, Cambridge (MA), 1998.

[Ferrario and Oltramari, 2004] Ferrario, R. and Oltramari, A. Towards a computational ontology of mind. In A. C. Varzi and L. Vieu, editors, *Formal Ontology in Information Systems, Proceedings of the Intl. Conf. FOIS 2004*, pages 287–297. IOS Press, 2004.

[Gärdenfors, 2000] Gärdenfors, P. *Conceptual Spaces: the Geometry of Thought.* MIT Press, 2000.

[Geach, 1969] Geach, P. *God and the Soul.* Routledge and Kegan Paul, London, 1969.

[Guha and Lenat, 1990] Guha, R.V. and Lenat, D. Cyc: A mid-term report. *AI Magazine*, 11(3):32–59, 1990.

[Hintikka, 1964] Hintikka, J. *Knowledge and Belief.* Cornell Univ. Press, 1964.

[Hodges, 1983] Hodges, W. Elementary predicate logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume I, pages 1–131. Dordrecht: Reidel, 1983.

[Houkes and Meijers, 2006] Houkes, W. and Meijers, A. The ontology of artefacts: the hard problem. *Studies in History and Philosophy of Science*, 37(1):118–131, 2006.

[Kroes and Meijers, 2006] Kroes, P. and Meijers, A. The dual nature of technical artefacts. *Studies in History and Philosophy of Science*, 37(1):1–4, 2006.

[Lowe, 1998] Lowe, E. *The possibility of metaphysics.* Clarendon Press, Oxford, 1998.

[Lowe, 2002] Lowe, E. *A survey of metaphysics.* Oxford University Press, Oxford, 2002.

[Masolo et al., 2003] Masolo, C., Borgo, S., Gangemi, A., Guarino, N. and Oltramari, A. Ontology Library (Wonder-Web Deliverable D18). Available at http://wonderweb.semanticweb.org/deliverables/documents/D18.pdf, 2003.

[Masolo et al., 2004] Masolo, C., Vieu, L., Bottazzi, E., Catenacci, C., Ferrario, R., Gangemi, A. and Guarino, N. Social roles and their descriptions. In D. Dubois, C. Welty, and M. Williams, editors, *Proceedings of the 9th International Conference on the Principles of Knowledge Representation and Reasoning (KR)*, pages 267–277, 2004.

[McLaughlin, 1995] McLaughlin, B.P. *Dispositions.* In J. Kim and E. Sosa, editors, *A Companion to Metaphysics*, pages 121–124. Oxford: Blackwell Publishers, 1995.

[Mumford, 1998] Mumford, S. *Dispositions.* Oxford University Press, Oxford, 1998.

[Pianesi and Varzi, 2000] Pianesi, F. and Varzi, A. Events and Event Talk: An Introduction. In J. Higginbotham, F. Pianesi and A. Varzi, editors, *Speaking of Events*, pages 3–47. New York: Oxford University Press, 2000.

[Prévot et al., 2005] Prévot, L., Borgo, S. and Oltramari, A. Interfacing ontologies and lexical resources. In *Ontologies and Lexical Resources: IJCNLP-05 Workshop*, pages 1–12, 2005.

[Rao and Georgeff, 1991] Rao, A. and Georgeff, M. Modeling rational agents within a bdi-architecture. In J. F. Allen, R. Fikes, and E. Sandewall, editors, *Principles of Knowledge Representation and Reasoning (KR'91)*, pages 473–484. Morgan Kaufmann, 1991.

[Rea, 1997] Rea, M. editor. *Material Constitution. A Reader.* Rowman and Littlefield Publishers, Lanham, MD, 1997.

[Scheele, 2005] Scheele, M. *The proper use of artefacts: A philosophical theory of the social constitution of artefact functions.* PhD dissertation, Technical University of Delft, 2005.

[Searle, 1983] Searle, J.R. *Intentionality.* Cambridge University Press, 1983.

[Searle, 1995] Searle, J.R. *The Construction of Social Reality.* The Free Press, New York, 1995.

[Smith, 1995] Smith, B. Formal ontology, common sense and cognitive science. *International Journal of Human Computer Studies*, 43(5/6):626–640, 1995.

[Smith, 1998] Smith, B. Basic concepts of formal ontology. In N. Guarino, editor, *Proceedings of the First International Conference FOIS 1998*, pages 19–28. IOS Press, 1998.

[Sowa, 2000] Sowa, J. Knowledge Representation: Logical, Philosophical, and Computational Foundations. Brooks/Cole, Pacific Grove, CA, 2000.

[Thomasson, 2007] Thomasson, A. Artifacts and human concepts. In E. Margolis and S. Laurence, editors, *Creations of the Mind: Theories of Artifacts and Their Representation*, pages 52–73. Oxford University Press, 2007.