Identity, Unity, and Individuality: Towards a Formal Toolkit for Ontological Analysis

Nicola Guarino[1] and **Christopher Welty**[2]

Abstract. We introduce here the notions of identity and unity as they have been discussed in Philosophy, and then provide additional clarifications needed to use these notions as fundamental tools in a methodology for ontology-driven conceptual analysis. We show how identity and unity complement each other under a general notion of individuality, and conclude with an example of how these tools can be used in analysis to help check the ontological consistency of taxonomies.

1 INTRODUCTION

Identity is one of the most fundamental notions in ontology, yet the related issues are very subtle, and isolating the most relevant ones is not an easy task; see [5] for an account of the identity problems of ordinary objects, and [11] for a collection of philosophical papers in this area. In particular, the relationship between *identity* and *unity* appears to be crucial for our interest in ontological analysis. These notions are different, albeit closely related and often confused under a generic notion of identity. Strictly speaking, identity is related to the problem of distinguishing a specific instance of a certain class from other instances by means of a *characteristic property*, which is unique for *it* (that *whole* instance). Unity, on the other hand, is related to the problem of distinguishing the *parts* of an instance from the rest of the world by means of a *unifying relation* that binds them together (not involving anything else).

For example, asking “Is that my dog?” would be a problem of identity, whereas asking “is the collar part of my dog?” would be a problem of unity. As we shall see, the two notions are complementary: when something can be both recognized as a whole and kept distinct from other wholes then we say that it is an *individual*, and can be counted as *one*.

The actual *conditions* we use to support our judgements concerning identity and unity for a certain class of things vary from case to case, depending on the properties holding for these things. If we find a condition that consistently supports identity or unity judgements for *all* instances of a certain property, then we say that property *carries* an identity or a unity condition. Of course, deciding this depends on the assumptions resulting from our *conceptualization* of the world, i.e. on our *ontology* [3]. For example, the decision as to whether cats remain the same after they lose their tails, or whether statues are identical with the marble they are constituted of, are ultimately the result of our sensory system, our culture, and so on. The aim of the present analysis is to clarify the formal tools needed to make such assumptions explicit, and to explore the logical consequences of them. These formal tools will form the foundation of a rigorous methodology for ontology driven conceptual modeling.

We begin by introducing the notions of identity and unity as they have been discussed in Philosophy, and then provide additional clarifications needed for our methodology. We show how identity and unity complement each other under a general notion of individuality, and conclude with an example of how these tools can be used in analysis to help check the ontological consistency of taxonomies.

2 PRELIMINARIES

Logically speaking, identity is a primitive equivalence relation, with the peculiar property of allowing the substitution of terms within logical formulas (Leibniz’s rule). In the following, we shall adopt a first order logic with identity. This will be occasionally extended to a simple temporal logic, where all predicates are temporally indexed by means of an extra argument. If the time argument is omitted for a certain predicate P, then the predicate is assumed to be time invariant, that is $t\mathrm{P}(x,t) \rightarrow t\mathrm{P}(x,t)$. Note that the identity relation will be assumed as time invariant: if two things are identical, they are identical forever. This means that we are assuming absolute identity, not relative identity, and Leibniz’s rule holds with no exceptions.

Our domain of quantification will be that of possibilia. That is, the extension of predicates will not be limited to what exists in the actual world, but to what exists in any possible world [7]. For example, a predicate like “Unicorn” will not be empty under this account. The special predicate $\mathrm{E}(x,t)$ will be used to express that x has actual existence at time t. Free variables appearing in formulas are assumed to be universally quantified. We also assume all properties to be *discriminating properties* that are not trivially false nor trivially true [3].

3 IDENTITY

Before discussing the formal structure of identity conditions (ICs), some clarifications about their intuitive meaning may be useful. If we say, “Two persons are the same if they have the same SSN,” we seem to create a puzzle: how can they be *two* if they are the same? The puzzle can be solved by recognizing that two (incomplete) descriptions of a person (like two records in different databases) can be different while referring to the same individual. The statement “two persons are the same” can be therefore rephrased as “two descriptions of a person refer to the same object”. A description can be seen as a set of properties that apply to a certain object. Our intuition is that two incomplete descriptions denote the same object if they have an identifying property in common.

Depending on whether the two descriptions hold at the same time, we distinguish between *synchronic* and *diachronic* ICs. The former are needed to tell, e.g., whether the statue is identical with the marble it is made of, or whether a hole is identical with its filler

[1] LADSEB-CNR, Corso Stati Uniti 4, I-35127 Padova, Italy, email: Nicola.Guarino@ladseb.pd.cnr.it
[2] On sabbatical at LADSEB-CNR from Vassar College, Poughkeepsie, NY, USA, email: welty@cs.vassar.edu.

[1], while the latter allow us to re-identify things over time.

In the philosophical literature, an *identity criterion* is generally defined as a condition that is both necessary and sufficient for identity. According to [8], a property φ carries an identity criterion iff the following formula holds for a suitable ρ:

$$\varphi(x) \quad \varphi(y) \rightarrow (\rho(x,y) \leftrightarrow x=y) \tag{1}$$

Since identity is an equivalence relation, it follows that ρ restricted to φ must also be an equivalence relation.

The above formulation has two main problems, in our opinion. First, the nature of the ρ relation remains mysterious: what makes it an IC, and how can we index it with respect to time to account for the difference between synchronic and diachronic identity? Second, deciding whether a property carries an identity criterion may be difficult, since finding a ρ that is both necessary *and* sufficient for identity is often hard, especially for natural kinds and artifacts.

3.1 A framework for identity

Our intuition is that the nature of the ρ relation in (1) is based on the "sameness" of a certain *property*, which is unique to a specific instance. Suppose we stipulate, e.g., that two persons are the same iff they have the same fingerprint: the reason why this *relation* can be used as an IC for persons lies in the fact that a property like "having this particular fingerprint" is an *identifying property*, since it holds exactly for one person. Fingerprints are then *identifying characteristics* of persons.

Identifying properties can be seen as relational properties, involving a *characteristic relation* between a class of individuals and their identifying characteristics. Such characteristics can be internal to individuals themselves (parts or qualities) or external to them (other "reference" entities). So two things can be the same because they have some parts or qualities in common, or because they are related in the same way to something else (for instance, we may say that two material objects are the same if they occupy the same spatial region). This means that, if χ denotes a suitable characteristic relation for φ, we can assume:

$$\rho(x,y)= \quad z(\chi(x,z) \leftrightarrow \chi(y,z)) \tag{2}$$

The scheme (1) becomes therefore:

$$\varphi(x) \quad \varphi(y) \rightarrow (\quad z(\chi(x,z) \leftrightarrow \chi(y,z)) \leftrightarrow x=y) \tag{3}$$

For instance, if we take φ as the property of being a set, and χ as the relation "has-member", this scheme tells us that two sets are identical iff they have the same members.

An important advantage of (3) over (1) is that it is based on a characteristic relation χ holding separately for x and y, rather than on a relation ρ holding between them. This allows us to take time into account more easily, clarifying the distinction between *synchronic* and *diachronic* identity:

$$E(x,t) \quad \varphi(x,t) \quad E(y,t') \quad \varphi(y,t') \rightarrow (\quad z(\chi(x,z,t) \leftrightarrow \chi(y,z,t')) \leftrightarrow x=y) \tag{4}$$

We shall have a synchronic criterion if $t=t'$, and a diachronic criterion otherwise. For the sake of simplicity, we are restricting our criteria to the times where the entities to be identified actually exist. Note that accounting for the difference between synchronic and diachronic identity would be difficult with (1): we may think of adding two temporal arguments to the ρ relation, but in this case its semantics would become quite unclear, being a relation that binds together two entities at different times. Note also that synchronic identity criteria are weaker than diachronic ones. For instance, the sameness of spatial location is usually adopted as a synchronic identity criterion for material objects, but of course it does not work as a diachronic criterion.

A possible criticism of (4) is that it looks circular, since it defines the identity between x and y in terms of the identity between something else (in this case, the identifying characteristics z common to x and y). However, as observed by Lowe ([9], p. 45), we must take in mind that *ICs are not definition*s, as identity is a primitive.This means that the circularity of identity criteria with respect to the very notion of identity is just a fact of life: identity can't be defined. Rather, we may ask ICs to be *informative*, in the sense that identity conditions must be non-circular with respect to the properties involved in their definition. For instance, Lowe points out that Davidson's identity criterion for events, stating that two events are the same if they have the same causes and they are originated by the same causes [2], is circular in this sense since it presupposes the identity of causes, which are themselves events. In many cases, however, even this requirement cannot easily be met, and we must regard ICs as simple constraints.

There is another objection we can raise against (4), namely that it is not general enough. In particular, the notion of "sameness" of characterizing properties does not capture the case where spatio-temporal continuity is taken as a criterion for diachronic identity. This criterion is of course not valid in general [5], although we believe it may hold for certain entities, like atoms of matter. If we want to state a general scheme for identity criteria, overcoming the problems of (1) and the restrictions of (4), we generalize the relation ρ into a generic *formula* containing x, y, t, t' as the only free variables, of course excluding the trivial cases (as discussed in section 3.2):

$$E(x,t) \quad \varphi(x,t) \quad E(y,t') \quad \varphi(y,t') \rightarrow (\Psi(x,y,t,t') \leftrightarrow x=y) \tag{5}$$

3.2 Weak identity conditions

In the philosophical literature, properties carrying an IC are called *sortals* [14]. In general, their linguistic counterparts are *nouns* (e.g., Apple), while non-sortals correspond to *adjectives* (e.g., Red). Distinguishing sortals from non-sortals is of high practical relevance for conceptual modeling, as we tend to naturally organize knowledge around nouns. Unfortunately, recognizing that a property carries a *specific* IC is often difficult in practice. However, in many cases it suffices to recognize that a property carries *some* (kind of) IC, without telling exactly *which* IC it is. To achieve this goal, we can introduce *weak* ICs, which are (only) necessary or (only) sufficient for identity. In other words, assuming that an unknown ρ satisfying (1) exists, we can look for generic relations γ, not necessarily equivalence relations, satisfying either:

$$\gamma(x,y) \rightarrow \rho(x,y) \tag{6}$$

$$\rho(x,y) \rightarrow \gamma(x,y) \tag{7}$$

in a non-trivial way [15]. γ will correspond to a sufficient IC in (6), and to a necessary IC in (7). We shall say that a property is a *sortal* if we find a γ such that at least one of these two conditions is satisfied.

In the same way we generalized ρ into Ψ in the previous section, we generalize the relation γ into a formula Γ (depending in general on φ) containing x, y, t, t' as the only free variables, and define the sufficient and necessary cases as follows:

Definition 1 A *sufficient identity condition* for a property φ is a formula Γ, such that:

$$E(x,t) \quad \varphi(x,t) \quad E(y,t') \quad \varphi(y,t') \quad \Gamma(x,y,t,t') \rightarrow x=y \tag{8}$$

$$\neg \quad xy(\Gamma(x,y,t,t') \leftrightarrow x=y) \tag{9}$$

$$xytt' \, \Gamma(x,y,t,t') \tag{10}$$

Definition 2 A *necessary identity condition* for a property φ is a formula Γ, such that:

$$E(x,t) \quad \varphi(x,t) \quad E(y,t') \quad \varphi(y,t') \quad x=y \rightarrow \Gamma(x,y,t,t') \tag{11}$$

$$\neg \quad xy(E(x,t) \quad \varphi(x,t) \quad E(y,t') \quad \varphi(y,t') \rightarrow \Gamma(x,y,t,t')) \tag{12}$$

$$\neg \quad xy(\Gamma(x,y,t,t') \leftrightarrow x=y) \tag{13}$$

(8) and (11) come from (5), each considering only one sense of the double implication. (9) and (13) guarantee that Γ is bound to identity under a certain sortal (φ), and not to arbitrary identity. (10) ensures that Γ is not trivially false. (12) is needed to guarantee that the last conjunct in (11) is relevant and not tautological.

4 UNITY

The notion of unity is closely tied to that of parthood, and our formalization requires some basic definitions. We adopt a time-indexed mereological relation $P(x,y,t)$, meaning that x is a (proper or improper) part of y at time t, satisfying the minimal set of axioms and definitions (adapted from [13], p. 362) shown in Table 1. Differently from Simons, this mereological relation will be taken as completely general, holding on a domain which includes individuals, collections, and amounts of matter. Based on these definitions, it should be clear that our definitions of unity below are synchronic, and therefore hold only at one time. The notion of what parts can change over time is tied to identity, not to unity.

Table 1. Axioms and definitions for the part-of relation.

$PP(x,y,t) =_{def} P(x,y,t) \quad \neg x=y$	(proper part)
$O(x,y,t) =_{def} \quad z(P(z,x,t) \quad P(z,y,t))$	(overlap)
$P(x,y,t) \rightarrow E(x,t) \quad E(y,t)$	(actual existence of parts)
$P(x,y,t) \quad P(y,x,t) \rightarrow x=y$	(antisymmetry)
$P(x,y,t) \quad P(y,z,t) \rightarrow P(x,z,t)$	(transitivity)
$PP(x,y,t) \rightarrow \quad z(PP(z,y,t) \quad \neg O(z,x,t))$	(weak supplementation)

4.1 Contingent unity

Before addressing what it means for a certain property to *carry* a unity condition (UC), we must first clarify what it means for a certain object to *have* a UC, that is to be a *whole*. A general and informal definition for wholeness was proposed by Peter Simons [13]:

> "Every member of some division of the object stands in a certain relation to every other member, and no member bears this relation to anything other than members of the division." (p. 327)

Here, a division of a certain object is assumed to be a class of parts (not necessarily disjoint from each other) completely exhausting it. The "unifying" relation binding the members of a division together must be an equivalence relation. An example could be "sharing both parents".

In his discussion, Simons emphasizes two aspects. First, the unifying relation may not hold between arbitrary parts of the whole, but just between those parts that are members of a division: "sharing both parents" does not involve parts of persons (at least not directly). Second, such a unifying relation can be constructed from an arbitrary "base" relation R by taking the transitive closure of the union of R with its converse $(R \quad R^{-})^{*}$, forming an equivalence relation. For example, starting with the base relation "being parent of", we can form an equivalence relation that binds together all elements of a biological species, say all humans.

From these general intuitions, Simons presents two definitions, which we have adapted to our terminology:

Definition 3 A class a is a *closed system under R* if

$$x \quad a \rightarrow (y \quad a \leftrightarrow R(x,y) \quad R(y,x)) \tag{14}$$

Definition 4 An object w is a *(contingent) R-integrated whole* if there exists a division a of w such that a is a closed system under $(R \quad R^{-})^{*}$. R will be called a *base unifying relation* for w.

We now generalize Simons' definition by saying that an object is an integrated whole if it has a division that is a closed system under a suitable equivalence relation, which will be called its *unifying relation*. As we have underlined, such a relation (let's call it υ) holds between the elements of a certain division, not between all the parts of an object. However, if υ is a unifying relation for a certain object, we can always construct a *deep unifying relation* ω that binds together its parts, by assuming

$$\omega(x,y,t) =_{def} \quad zz' \, (P(x,z,t) \quad P(y,z',t) \quad \upsilon(z,z',t)) \tag{15}$$

We can easily check that, since υ is an equivalence relation, ω will also be an equivalence relation. We are therefore in the position to state the following (still preliminary) general definition, which avoids mentioning a suitable division:

Definition 5 Let ω be an equivalence relation. At a given time t, an object x is a *contingent whole under* ω if:

$$y(P(y,x,t) \rightarrow \quad z(P(z,x,t) \leftrightarrow \omega(z,y,t))) \tag{16}$$

$$\neg \quad xyztP(y,x,t) \quad P(z,y,t) \leftrightarrow \omega(y,z,t) \tag{17}$$

We can read the above definition as follows: *at time t, each part of x must be bound by* ω *to all other parts and to nothing else.* (17) is a non-triviality condition on ω, that avoids considering any mereological sum as a contingent whole. (16) expresses a condition of *maximal self-connectedness* according to a suitable relation of "generalized connection," ω. The intuition is to exclude connection relations that are trivially constructed from the part-of relation.

Depending on the ontological nature of this relation, we may have different kinds of unity. For example, we may distinguish *topological unity* (a piece of coal, a lump of coal), *morphological unity* (a ball, a constellation), *functional unity* (a hammer, a bikini). As these examples show, nothing prevents a whole from having parts that are themselves wholes (with a different UC). Indeed, a *plural whole* can be defined as a whole which is a sum of wholes.

4.2 Essential Unity

Simons touches only briefly on the temporal issues related to integrity. Nothing prevents a physical object from having a certain UC only at a single time, being therefore only a *contingent* whole. Consider for instance an isolated piece of clay; which certainly has a certain topological unity; what happens if it is attached to a much larger piece? We can hardly say it is still a topological whole. Yet we may invent another relation: supposing that our piece occupies a certain spatial region r within the larger piece, a relation like *being-materially-connected-but-confined-within-r* could do. Our piece of clay will again be a whole, but only in a contingent way, as its spatial location may change as soon as the object moves.

We define a stronger notion of whole by assuming that a UC must hold for an object throughout its existence, i.e. by assuming unity as an *essential* property:

Definition 6 An object x is an *intrinsic whole under* ω if, at any time where x exists, it is a contingent whole under ω.

An important remark is that, if an object is always atomic (i.e., it has no proper parts), then it is an intrinsic whole under the identity relation. We are now in the position to state the following:

Definition 7 A property φ *carries a unity condition* if there is a relation ω such that instances of φ are intrinsic wholes under ω.

It is important to make clear that carrying a UC does not imply carrying a necessary IC. This is due to the way Definition 2 is formulated. To see that, suppose that φ carries a UC. We may think that the *persistence* of such condition across time could be a good candidate for a necessary IC for φ, since it satisfies (11). However, it fails to satisfy (12), and does not qualify as a necessary *identity* condition: thus, unity is just a persistence condition.

5 INDIVIDUATION AND COUNTABILITY

We have seen how identity and unity, though related, are independent from each other. In fact, properties may carry only identity, only unity, both, or neither of them. When something is an instance of a property *carrying* identity, it *can be identified*. If something is an instance of a property that carries unity, it is a *whole*. If something can be identified *and* is a whole, then we say it is an *individual.* The notion of *individuality* can be seen therefore as the sum of identity plus unity. This is not to say, however, that all individuals are instances of properties that carry unity and identity, since it is possible for some instance to be an individual for only part of its existence (as discussed in section 4.2). Moreover, we must be careful in not confusing individuality with *singularity*: an individual can count as *one* even if it is a plural whole.

With the help of the examples reported in Table 2, we now discuss various combinations of identity and unity.

Table 2. Examples of properties carrying different IC/UC combinations.

	Property	Identity	Unity
1	Apple	√	√
2	Apple piece	√	√
3	Apple food	√	—
4	Intrinsic whole	—	√
5	Red	—	—

Case 1 of Table 2 shows the prototypical example of a *countable property,* whose instances are all and always individuals:

Definition 8 A *countable property* is a property that carries both an identity and a unity condition.

Countability can be used therefore as a practical test to check whether a property carries an IC and a UC, even if the specific IC or UC may be not clearly determined.

In case 2 we still have countability, at least if we intend a "piece" as an undetached self-connected part of something (otherwise we cannot admit unity any more). Notice that in this case the UC is different from before: for the apple we may rely on a notion of biological unity, while for the piece we adopt maximal self-connectedness.

In case 3 we have the classical example of a mass-sortal, which carries no UCs (assuming that the parts of something that is "apple food" can be arbitrarily scattered) while carrying an IC. The IC assumed here is based on the *mereological extensionality* of food: two amounts of food are the same iff they have the same parts.

Case 4 shows an example of a property that trivially carries a UC without carrying an IC. In general, we may have practical cases where we want to avoid committing on ICs for a certain property while admitting UCs. For instance, within a certain application we may be interested in counting tokens without caring about the possibility of distinguishing one token from another. Lowe cites electrons as an example of entities that are whole (because of their mass and charge) that cannot be identified (because of Pauli's indeterminacy principle). While this example is debatable, it is important to admit at least in principle this possibility.

Finally, case 5 is an example of a property carrying neither identity nor unity (assuming that instances of "Red" are material objects, not particular color patches). Notice that this does not mean that red objects cannot be identified, since their IC can be supplied by other properties they are instances of. Indeed, we assume (as per Quine [12]) that every element of our quantification domain must be identifiable, although not necessarily a whole and therefore not every entity is an individual.

6 TAXONOMIC CONSTRAINTS

After the above clarifications, let us see how ICs and UCs may affect a taxonomic organization. One result of this work, upon which the methodology we are developing is based, is that the presence of such conditions imposes some constraints (which are *theorems* descending immediately from our definitions) on the IS-A relation. Due to space limitations, we shall avoid logical notation here, focusing on explaining the relevant concepts.

In the case of ICs, the following principle has been proposed by Lowe:

> "No individual can instantiate both of two sorts if they have different criteria of identity associated with them." ([8] p. 19)

We believe that this principle is illuminating, but its formulation is not accurate enough. Consider the domain of abstract geometrical figures, for example, where the property "Polygon" subsumes "Triangle". A necessary and sufficient IC for polygons is "Having the same edges and the same angles". On the other hand, an *additional* necessary and sufficient IC for triangles is "Having two edges and their internal angle in common" (note that this condition is only-necessary for polygons). So the two properties have *different* ICs (although they have one IC in common), but their extensions are not disjoint. The point, then, is not having different ICs, but having *incompatible* ICs. For example, "Amount of matter" must be disjoint from "Person" if we admit mereological extensionality for the former but not for the latter (since persons can replace their parts).

In the case of UCs, an obvious constraint is that a property carrying a UC cannot subsume one carrying no UCs. A more complicated situation arises when a property carrying no unity subsumes one that does. Suppose for instance we wonder if "Vase" is subsumed by "Amount of clay". We may think of a vase as an amount of clay c that has the property of being a whole, satisfying a suitable UC for vases. This is no problem at first, however since the vase must be an *intrinsic* whole for the corresponding property to carry a UC, we must verify whether this is compatible with our ontological assumptions for amounts of clay. This forces a modeler to consider what they mean by “an amount of clay.” In one account, certainly all amounts of clay must have the possibility of *not* forming a whole, therefore the subsumption is not valid. This analysis of UCs brings to light a very common misuse of the subsumption relation, the fact is that vases are *constituted* of amounts of clay, not subsumed by them.

7 A MOTIVATING EXAMPLE

We shall briefly discuss now how the notions we have defined can help in the task of ontological analysis. The scenario we are considering for our motivating example is that of a simple taxonomy of concepts, resulting from a preliminary study of a certain domain. We focus on a deceptively simple example, concerning the subsumption relationship between two concepts: physical objects and amounts of matter. Assuming that such concepts are not co-extensional, two well-known ontologies take opposite positions:

- A physical object is an amount of matter (Pangloss) [6]
- An amount of matter is a physical object (WordNet) [10]

This example illustrates how the lack of rigorous formal tools can lead to drastic inconsistencies even for experienced modelers.

Applying our analysis to these properties, we see that the usual account of amounts of matter is that they have an extensional IC (two amounts of matter are the same iff they have the same parts), and no UC. The usual account of physical objects is that they are not extensional, since two physical objects may be the same while having different parts (e.g. a car with new tires). In a another account, physical objects may be extensional. In either case, physical objects are normally considered to have unity (because they are countable).

Turning to our taxonomic choices, the presence of a UC on physical objects prevents them from subsuming amounts of matter. The normal account of physical objects (as non-extensional) prevents them from being subsumed by amounts of matter. We are left with the case where extensional physical objects may be subsumed by amounts of matter, however considering the analysis of clay and vases in the previous section, this also violates the essentiality of the lack of unity for amounts of matter.

This analysis tells us that, in the normal account of these concepts, amounts of matter and physical objects are disjoint.

8 CONCLUSION

We have attempted a compact and rigorous formalization of the subtle notions behind identity, by assembling, clarifying, and adapting philosophical insights in a way useful for practical knowledge engineering. We have clarified in particular the importance of dealing with separate only-sufficient and only-necessary identity conditions, and the constraints they impose on a taxonomy. We have shown how the mix of unity and identity can help formalizing the notion of a countable property, improving previous accounts such as [3].

Our claim is that a rigorous analysis of identity and unity assumptions can offer two main advantages to the knowledge engineer:

- It results in a cleaner taxonomy, due to the semantic constraints imposed on the IS-A relation;
- It forces the analyst to make ontological commitments explicit, clarifying the intended meaning of the concepts used, exposing hidden assumptions, and producing therefore a more reusable ontology.

ACKNOWLEDGMENTS

We are indebted to Claudio Masolo, Pierdaniele Giaretta, Dario Maguolo, and the anonymous reviewers for helpful comments and feedback on earlier versions of this paper.

REFERENCES

[1] Casati, R. and Varzi, A. C. 1994. *Holes and Other Superficialities.* MIT Press/Bradford Books, Cambridge (MA) and London (UK).

[2] Davidson, D. 1980. *Essays on Actions and Events.* Clarendon Press.

[3] Guarino, N., Carrara, M., and Giaretta, P. 1994. An Ontology of Meta-Level Categories. In E. Sandewall and P. Torasso (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourth International Conference (KR94).* Morgan Kaufmann, San Mateo, CA: 270-280.

[4] Guarino, N. 1998. Formal Ontology in Information Systems. In N. Guarino (ed.) *Formal Ontology in Information Systems. Proceedings of FOIS'98, Trento, Italy, 6-8 June 1998.* IOS Press, Amsterdam: 3-15.

[5] Hirsch, E. 1982. *The Concept of Identity.* Oxford University Press, New York, Oxford.

[6] Knight, K. and Luk, S. 1994. Building a Large Knowledge Base for Machine Translation. In *Proceedings of American Association of Artificial Intelligence Conference (AAAI-94).* Seattle, WA: 773-778.

[7] Lewis, D. 1983. New Work for a Theory of Universals. *Australasian Journal of Philosophy*, **61**(4).

[8] Lowe, E. J. 1989. What is a Criterion of Identity? *The Philosophical Quarterly*, **39**: 1-21.

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

[10] Miller, G. A. 1995. WORDNET: A Lexical Database for English. *Communications of ACM*, 2(11): 39-41.

[11] Noonan, H. (ed.) 1993. *Identity.* Dartmouth, Aldershot, USA.

[12] Quine, W. V. O. 1969. Ontological Relativity and Other Essays. Columbia University Press, New-York, London.

[13] Simons, P. 1987. *Parts: a Study in Ontology.* Clarendon Press, Oxford.

[14] Strawson, P. F. 1959. Individuals. An Essay in Descriptive Metaphysics. Routledge, London and New York.

[15] Williamson, T. 1986. Criteria of Identity and the Axiom of Choice. *The Journal of Philosophy*, **83**: 380-94.