Adaptive Multimedia Retrieval: Retrieval, User, and Semantics: 5th International Workshop, AMR 2007, Paris, France, July 5-6, 2007 Revised Selected Papers. Lecture Notes in Computer Science, Volumen 4918. Springer, 2008. 17-31.

DOI: http://dx.doi.org/10.1007/978-3-540-79860-6_2

Ontology: Use and Abuse

Simone Santini

Escuela Politécnica Superior, Universidad Autónoma de Madrid, Spain

Abstract. This paper is a critical analysis of the use of ontology as an instrument to specify the semantics of a document. The paper argue that not only is a logic of the type used in ontology insufficient for such a purpose, but that the very idea that meaning is a property of a document that can be expressed and stored independently of the interpretation activity is misguided.
The paper proposes, in very general lines, a possible alternative view of meaning as modification of context and shows that many current approaches to meaning, from ontology to emergent semantics, can be seen as spacial cases of this approach, and can be analyzed from a very general theoretical framework.

In his book *What do you care what other people think?* the physicist Richard Feynmann remembers the way his father used to teach him about birds on their hiking trips together:

> "See that bird?" he says "It's a Spencer's warbler." (I knew he didn't know the real name.) "Well, in Italian, it's a *Chutto lapittida*. In Portugese is a *Bom da Peida*. In Chinese is a *Chung-long-ta*, and in Japanese it's a *Katano Tekeda*. You can know the name of that bird in all the languages of the world, but when you are finished, you'll know absolutely nothing whatever about the bird. You'll only know about humans in different places, and what they call the bird."[1]

I am often reminded of this story when I think of ontology. The feeling I have is that one can encode a document in whatever formalism one can think of, as many times as one wants but, in the end, one will know nothing about the document, only something about the person who encoded it.

Generally speaking, computational ontology, which I will simply call *ontology*, is an attempt to encode the semantics of a text or of a document in a formal structure that can be manipulated following logical (syntactic) rules. A formula, really. In this paper, I will use the word *text* in the very extended sense in which semioticians use it: to me, for the purposes of this paper, a written document, an image, a video, fashion, a series of moves in a chess games, the etiquette on how to behave in a given circumstances, and many, many other things, are examples of texts.

Some of the texts for which one might want to formalize a meaning are themselves highly formalized, programming languages, for example. I am not interested in them here, although they will make a fleeting appearance later on. I am interested in the texts that people use to communicate with each other, of which one can ask with some

[1] (Feynman, 1988), p. 13-14

interest "What do they mean?". Ontology tries to answer this question by creating a logic theory meant to give the possibility to *process* the meaning of texts in an automatic fashion. Using an algorithm, as it is. The adjective "syntactic", in parentheses in the previous paragraph, is in this sense very important: logic systems manipulate data based on their *form*, not on their meaning; in this sense ontology can be seen as an attempt to reduce semantics to syntax. Note that ontology makes weaker (viz. less assailable) epistemological claims than artificial intelligence[2]: it never assumes that meaning can be extracted from a text and encoded in an automatic fashion. Ontologists are quite content with the idea that a person might have to be there to make the encoding. What they *do* assume is that a formal code exists in which meaning can be translated and that, once the translation is done, meaning can be accessed through an algorithm, that is, through syntactic means. They do not assume that the semantics of a document can be interpreted in the document itself, but they do assume that there is a code in which semantics reduces to syntax.

My purpose in this paper is to analyze the plausibility of these assumptions. I will present several reasons to believe that the theoretical assumptions of ontology, although weaker than those of "strong" artificial intelligence, are no more tenable than those. I will argue this point based on what we now about the semantics of natural language. It is worth noticing that the objections I will make do not cover the whole field of application of ontology. In Section 3 I will make a distinction between eminently *syntactic* applications, such as schema integration (Hakimpour & Geppert, 2001), and those that claim to be doing *semantics*. The prototype of the latter class of application are the rather extravagant claims of the *semantic web* (Berners-Lee, Hendler & Lassila, 2001). What the web semanticians are after is nothing less than the codification of the meaning of web sites, with their multimedia expressivity and cultural variability.

The first class of applications is, *mutatis mutandis* a slight variation of what was once known as *deductive databases*. While I don't see any major breakthrough coming from the renewed interest and the new terminology that are being used these days, there are certainly problems in this general area that can benefit from the "ontological" methods. The second class of problems is the focus of this paper and, I will argue, the area where the claims are more doubtful and less supported.

1 What is an ontology?

A good answer to this question is surprisingly difficult to find. The definition that is usually quoted comes from (Uschold & Gruninger, 1996) and sounds something like this:

> [an ontology is] an explicit account of a *shared understanding*[3]

This definition is bizarre for at least three reasons. The first is the presence of the word *shared*. Now, it is true that in general one creates an ontology in order to share

[2] Weaker in this context means that it stands on a smaller number of hypotheses, that is, that it doesn't assume as much. In this sense of the term, a weaker assumption is harder to disprove.

[3] Emphasis in the original.

it but why on God's green earth shouldn't I be allowed to create an ontology only for myself? Looking at the definition it seems that, if I create a certain artifact (we'll see which) for myself it is not an ontology but, at the moment when I share it with my friend Bob from down the street, it becomes, *ipso facto*, one. I am not quite sure this is what Uschold had in mind but, as they say, *scripta manent*, and this is a logical consequence of his definition.

The second is the presence of the word *understanding*. The word sounds suspiciously related to the "offer he cannot refuse" from the *Godfather* but, apart from this, such a crucial point of the definition is no clearer than the thing that one is trying to define. If we don't know what an understanding is, this definition will tell us nothing about the substance of an ontology. So, what is an understanding? The ontologists don't tell us.

But the main problem of this definition is that it is not *structural* but *functional*: it doesn't tell us what an ontology *is* but, simply, what it is used *for*. Allow me to make a parallel. Take a well known computing concept, that of formal grammar. Had the same style of definition been used in that case, a formal grammar would be defined as:

> a formal grammar is the specification of a programming language.

Any computing scientist worth her salary knows that this is not a proper way to define a formal grammar, because the definition doesn't tell what a grammar is, but merely what one of its possible uses is. Rather, we define it as a quadruple N, T, P, s, where N is a set of non-terminals, T a set of terminals disjoint from N, P a set of pairs called productions (whose precise definition depends on the type of grammar), and $s \in N$ the initial non-terminal.

Can one find a similar definition of ontology? There doesn't seem to be a widespread agreement on any precise definition, but one can try to unify the common usages of the term. One relatively uncontroversial aspect of the matter is that an ontology is a system of axioms that spell out the rules that allow reasoning on the terms of a certain domain. Terms in the axioms are usually interpreted as universals in an extensional interpretation of meaning. In other words, they are interpreted as sets of individuals.

Also, an ontology contains more than just axioms: it also contains the definition of the terms in the domain of discourse and the relations that, subject to the axioms, hold between these terms: more than a set of axioms, it looks like a logical *theory*. We can therefore give a first, quite informal, definition of ontology:

An ontology is a theory that admits a model that associates a set to each term defined in the theory.

We still haven't specified what kind of logic system should be used to express the theory. To see what this entails, consider a simple system that defines three terms: *person*, *employee*, and *retiree*. Ontologists are fond of taxonomies, so the first thing an ontologist would do is to say that *retiree* and *employee* are *subsorts* of *person*. This entails defining, in our theory, a subsort relation $\preceq$, introducing axioms for it, and then declaring which terms are in that relation:

$$\forall X, Y, x \;\; X \preceq Y \;\wedge\; X(x) \;\Rightarrow\; Y(x) \tag{1}$$

(here I have used the convention that upper-case letters denote sets, and lower-case letter denote elements of the set. Also $X(x)$ is the unary predicate "x is an X".)

$$\begin{aligned} \text{employee} &\preceq \text{person} \\ \text{retiree} &\preceq \text{person} \end{aligned} \tag{2}$$

and similarly for the disjunction relation $\neq$:

$$\forall X, Y, x \;\; X \neq Y \;\wedge\; X(x) \;\Rightarrow\; \neg Y(x) \tag{3}$$

$$\text{employee} \neq \text{retiree} \tag{4}$$

These axioms are expressed using set variables and the unary predicates X and Y, which is a rather obvious thing to do if we deal with models in which the variables can be interpreted as set. Alas, this means that the logic that we are using to express them is *monadic second order logic*, which, containing first order logic, has the unpleasant characteristic of being undecidable.

There is no immediately obvious reason why undecidability should be a negative thing. On the one hand, it doesn't seem to generate any problem that, on a pragmatic plane, can't be solved by a good heuristics and common sense. On the other hand, at least for semantic web applications, undecidability seems to resonate with the general web attitude (a very positive one) of seeing imprecision and messiness as a positive fact of life: something to be welcomed as a source of expressive richness, not to be avoided as a source of potential problems[4]. In spite of this, ontologists are in general wary of undecidability, so they take great care to restrict the logic system in which they work, so as to make it decidable. With these considerations in mind, one can work, at least in the first approximation, with the following definition, which is not by all means above all suspicions but around which, I believe, the consent of the majority of the people involved could be gathered.

Definition 1. *Given a set V (of elements called* terms*) and a collection R of relations on V, an ontology for V and R (or* (V, R)-ontology*) is a decidable logical theory O on V and R that admits a model M such that $M \models O$ and, for each $v \in V$, $M(v)$ is a set.*

Consider again the *subclass* relation $X \prec Y$. In an extensional semantics the basis for this relation is subset-hood, so we might want to include two statements in the theory:

i) every X is also an Y;
ii) every statement about the properties of the term Y applies, *ipso facto* to X.

We can formulate these points with the following expressions, the first of which we have already encountered in (1):

[4] This attitude notwithstanding, ontology doesn't seem to prepare for the task of dealing with this messiness. For example, all the current ontological approaches use some form of monotonic logic, which is not the best model to use in an environment in which contradicting facts are to be expected with a certain frequency.

i) $\forall x, X, Y\ X(x) \wedge X \prec Y\ \Rightarrow\ Y(x)$;
ii) $\forall X, Y, P\ \ P(Y) \wedge X \prec Y\ \Rightarrow\ P(X)$.

As I have already mentioned, in so doing, we have introduced enough machinery from monadic second order logic that the logic is no longer decidable.

The problem in this case is that the machinery that we have introduced is very expressive—it can express much more than we need—and this expressivity generates undecidability. If we insist on decidability, we must find a way to limit the expressive power of monadic second order logic so that it will express the properties that one needs but, at the same time, it won't express too much, so to speak. One way in which this reduction can be achieved is by restricting the use of quantifiers to pre-specified patterns through the use of *operators* with defined semantics[5]. There seems to be two major classes of logics that ontologists use for this purpose: the first is the family of *description logics* in its various flavors, the second is *Horn logic* and its various extensions. Neither description logics nor Horn logic allow the unrestricted use of quantifiers but, for instance, in standard description logic an expression such as $X \sqsubseteq Y$ is equivalent to (1), with the quantification used therein, while in Horn logic the expression

$$A(x, z), B(z, y)\ \Rightarrow\ C(x, y) \tag{5}$$

is implicitly quantified as

$$\forall x, y\, \exists z\ A(x, z)\ \wedge\ B(z, y)\ \Rightarrow\ C(x, y) \tag{6}$$

All these syntactic restrictions entail, of course, corresponding restrictions in the modeling possibilities of the logic. Description logic, for instance, can't offer a model for the proposition "a resident student is a student who lives in the same city where he studies", while Horn logic (which can model the previous statement) can't model "a person is either a man or a woman, and no person is at the same time a man and a woman" (a statement that can be modeled in description logic).

In this paper, I am not quite interested in the technicalities of these logic systems, since they are irrelevant for the arguments that will follow: the point of contention here is not whether this or that specific logic system is expressive enough to encode the meaning of a document, but whether meaning is something that can be encoded at all, let alone encoded in a logic system, whatever its expressive power. So, from the point of view of the issue at hand, we don't need to be too specific on the type of logic that is being used: we may even assume the use of the full monadic second order logic: it won't make any difference.

For our purposes, it will suffice to know that an ontological approach to meaning consists of two parts:

i) a set of terms and relations specific to a document which contains the meaning of *that* document;

[5] In this context, when I speak of semantics, I mean the formal (e.g. denotational) semantics of the operators, nothing to do with the semantics of text, which is the subject of this paper. I have observed this confusion to happen, a fact that contributes greatly to muddle the ontological waters.

ii) an ontology, which may be divided in several chunks more or less coherently connected, which defines the classes from which these terms are drawn, and the general relations among these classes, and a number of axioms, in a suitable logic system.

2 Ontology and meaning

Is ontology a viable way of representing meaning? Let us begin by noting that, by posing the problem in these terms we are already begging the question of whether meaning can be represented at all, that is, if it can be reified as a property of a document. Ontology says that it can, and that it can be represented as a collection of axioms on terms and relations.

Since relations and their axioms are an important part of any ontology, one obvious way to start our analysis is to ask is whether they are constitutive of meaning or not, that is, once we have represented a text by referring its elements to an ontology, whether the meaning resides in the terms themselves or in the relations. We shall see that none of the two alternative is ultimately satisfactory.

Let us consider the first option first. This point of view is expressed quite well in Jerry Fodor's *informational semantics*:

> Informational semantics denies that "dog" means *dog* because of the way it is related to other linguistic expressions [...]. Correspondingly, informational semantics denies that the concept DOG has its content in virtue of its position in a network of conceptual relations

The "correspondingly" here does a lot of work, and requires a fairly important metaphysical investment since it maps conceptual structures to linguistic ones. This, *passim,* is the same investment that ontology requires when it takes a linguistic structure (composed of words and relations) and calls it a conceptual model.

Let us get rid immediately of the idea that "dog" means DOG because of the three letters of which it is composed. There is absolutely nothing in the sequence /d/, /o/, and /g/ that is in any way connected to dogness. The fact that I can read it and understand that we are talking about a dog tells something about me and the linguistic community in which I function, but absolutely nothing about dogs or dogness. If you don's speak Italian, to you the sequence /c/, /a/, /n/, and /e/ doesn't mean anything, but to me it means the same thing (with some important distinguos that I will consider shortly). As a matter of fact, the sequence

0 1 0 0 0 1 0 0 0 1 0 0 1 1 1 1 0 1 0 0 0 1 1 1

will probably appear obscure to most people but, if one assumes an 8 bit ASCII code, it means exactly the same thing.

The idea that the three symbols /d/, /o/, and /g/ are somehow related to the *concept* dog is, indeed, quite naïve, and the fact that in ontology not only it is never openly denied but, many times, it appears to be tacitly assumed as obvious, doesn't increase our confidence in the soundness of the approach.

But if the letters themselves do not create any connection between the symbol "dog" and the meaning of the word, where does this connection come from? What is left of the symbol once you take away the elements that constitute it? Where does its identity lie? The only way one can save the symbol is to say that its identity derives from its relations of opposition with the other symbols of the system. Dog is dog not because of the letters that make it up, but because they allow us to distinguish it from *dot*, from *hog*, from *god*. We are led, in other words, to a position that might oscillate between some form of cognitive functionalism (Stich, 1983) and structural semantics (Greimas, 1966), depending on the degree to which we want to rely on logic formulas in order to define meaning. Both these positions, in spite of their fundamental differences, will agree that the meaning of a symbol is not in the symbol itself, but in the whole system, and in the relation of the symbols with the other symbols.

In mathematical terms, one could say that a system of signification must be invariant to any isomorphic transformation of its terms: if we change dog in hog, hog in bog, and so on, in such a way that the differences between symbols are maintained, the ontology that we get must be exactly equivalent to the original one. Of course, we, as English speaking people, will be completely unable to read it, but here we are talking about algorithms, and they do not care if we write dog or bog, as long as they can distinguish one name from the other.

An isomorphism of this type will leave the relations between symbols unchanged so, if we take the second position outlined above—namely that the relations are constitutive of meaning—we obtain the necessary invariance. This position also entails that, whenever this relational invariance is not in force, meaning is not preserved. In other words: any transformations that is not an isomorphism of the terms of an ontology will not preserve meaning. A good way to test the plausibility of this assumption is to look at the relations between different languages. Different languages break the semantic field in different ways, and concepts arise at the fissures of these divisions. Consider, for example, the way in which adjectives of old age are constituted in Italian, Spanish and French[6]. The basic adjective, *vecchio/viejo/vieux* is applied both to things and to persons. There are specific forms, however: in Spanish, *añejo* is an appreciative form used mainly for liquors *(un ron añejo)*. The Italian adjective *anziano* is applied mainly to people, and the correspondence is roughly *anziano/anciano/âgé*, but *anziano* has a broader meaning than the other two adjectives, being used in expressions such as "il sergente anziano" to denote seniority in a function, a situation in which the Spanish would use *antiguo* and the French *ancien*. Note that Spanish also has the possibility of using the word *mayor* as a softer and more respectful form of denoting a person of old age, while the corresponding Italian and French words are never used in this sense. The correspondence is, in other words, according to the schema that follows. The differences are not just in the way different languages divide the same semantic axes, but also in the choice of semantic axes along which concepts are divided. In English, for instance, the two most widely used words that indicate moving bodies of sweet water are *river* and *stream*, while in Italian they are *fiume* and *affluente*.

[6] This example is an extension of a similar one in Geckeler (1976).

<table>
<tr><th>Italian</th><th>Spanish</th><th>French</th></tr>
<tr><td rowspan="2">vecchio</td><td>añejo</td><td rowspan="2">vieux</td></tr>
<tr><td>viejo</td></tr>
<tr><td rowspan="2">anziano</td><td>anciano</td><td rowspan="2">âgé</td></tr>
<tr><td>mayor</td></tr>
<tr><td rowspan="2"></td><td rowspan="2">antiguo</td><td>ancien</td></tr>
<tr><td>antico</td><td>antique</td></tr>
</table>

The semantic field in English is organized by size: streams are smaller than rivers and have a more irregular course. In Italian the semantic field is organized by destination: fiumi end up in the sea, while affluenti end up in other rivers.

One could build many examples of this, and even better ones if one considers languages that are culturally very different, such as Chinese and French, or Urdu and Italian. European languages are all (with notable exceptions: Welsh, Euskera, Hungarian, etc.), Indo-European languages, and have the same basic structure. Moreover, Europeans have been trading with each other and killing each other for centuries, activities that, of course, tend to mix up the language and to have them influence one another. People whose first language is not Indo-European, or who are proficient speakers of a non Indo-European language are in a better position to determine the extent to which different languages differ in their organization of the semantic field.

To the extent to which a functional translation from Chinese to English, or from Hungarian to Quechua are possible, then, we must admit that a meaning-preserving morphism is not required to be an isomorphism of terms that preserves relations[7]. Meaning, in other words, is a more abstract entity than a mere structural correspondence: depending on the global organization of the semantic field operated by a language, one can introduce considerable structural distortion and still end up with documents that "mean the same thing". Of course, this doesn't mean that all transformations are admissible; to make a trivial example, our previous consideration on the constitutive nature of relations tell us that, since a symbol is identified only by differentiation with other symbols, one can't have a signification system composed of a single symbol: a morphism that maps all the terms of a document or an ontology to a single one would not only destroy the meaning of that document, it would destroy the very idea of meaning even though, mathematically, we are in the presence of a homomorphism.

These example show or, at least, hint that terms and relations are simply not enough to determine meaning. Both the nature of the terms and of the relations can change quite dramatically, and we still have signification systems that can be considered roughly equivalent, at least to the extent that it is possible to translate an Milorad Pavic novel, written in Serbian, into English.

But is translation simply a linguistic problem? Language differences are very relevant to the problem of signification, and I believe that focusing on a single language one will not be able to place the problem of encoding meaning in its proper light. So, it might be useful to look at translation in a little more depth. All modern theories of

[7] As a matter of fact, it is not required to be a function at all: the idea of *one* correct translation has since long disappeared from translation theory (Moya, 2004). Rather, different translations are possible depending on the rôle that the translation will play in the receiving culture.

translation deny that translation is simply, or even mainly, a linguistic fact. If this were so, automatic translation would be relatively easy while we know that, declarations of success notwithstanding, it is an unsolved problem.

Eugene Nida, an American theorist considers translation as an act of cultural replacement (Nida, 1964). The work of a translator consists in studying the effect of a text in its original culture and translate it into a text that will achieve the same effect in the *target* culture. The emphasis for the translation of meaning here is not much on the content of the original linguistic expression, as much as on the *effect* that this expression has on the culture to which it is directed. In other words, the meaning of a text can't be separated from the act of interpretation, an act that is always cultural and contextual.

This orientation is even more pronounced in the successive developments of the theory of translation. The *skopos* theory (Vermeer, 1983) emphasizes that the primary force behind translation is the function assigned to the translated text by the translator, as an independent reader and interpreter of the text. This theory incorporates the opinions of reception theory of reading as a contextualized act in which meaning is created. The translator, as a reader of the original text, translates not the text itself, but his own specific reading of the text.

How is this relevant for ontology? Well, the transformation of a document into a formal text is a form of translation and, if we follow the finding of translation theorists, it has much less to do with a phantomatic inherent meaning of the text than with the contextualized reading of whoever did the encoding.

3 The ontology view of meaning

The perspective on meaning given by ontology is very different from the contextual, interpretative process that emerges from the previous considerations, and here lies, I believe, its main limitation. This limitation goes beyond the use of a specific logic system, and even beyond the limitations of any conceivable logic system: it derives from the disregard of interpretation as a creator of meaning and, consequently, to the idea that meaning is a *thing* rather than a process. In order for ontology to work, it is necessary that meaning be a *property* of a document, something that can be reified, formalized, and attached as a property to a document.

I have already argued that I see the idea of formalizing meaning in a set of symbols and relations between them as highly problematic, but I want to make you notice how the observations that we are about to make will lead us even further: the very idea that the meaning of a document is *in* the document, that it can somehow be attached to the document in such a way that it can be revealed to a un-contextualized reading, is quite wrong. But let us proceed in an orderly fashion.

An ontology encodes an absolute and immutable meaning of a text[8]. Where does it come from? For such an hypothesis to work, meaning must exist prior to text and

[8] This doesn't exclude the possibility that different encodings may give different, possibly conflicting, accounts of the meaning of a document, among which it may be necessary to negotiate. But every encoding will give one account of meaning, in absolute terms, that is, independently of the circumstances of interpretation.

independently of the language in which it is expressed. The scheme is pretty much that of a communication channel.

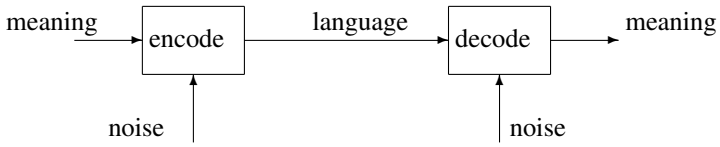

The origin of the communicative act is a *meaning* that resides with the author, and that the author wishes to express in a permanent text. This meaning is a-historical, immutable, and pre-linguistic. In order to communicate meaning, the author translates it into the shared code of language, and sends it to the receiver. This translation may be imperfect, as indicated by the "noise" arrow entering the translation box; a contingency due to the accidental imperfections of human languages. A perfect language (ontology acknowledges that this might be an unattainable theoretical limit) would be the perfect mirror of the *essential* meaning as it appears in the mind of the author and would allow a perfect translation. Once meaning is translated into language, it can be delivered to the reader, who can then proceed to decode it (possibly with the insertion of some more noise) obtaining a reasonable approximation of the original meaning as *intended* by the author.

This model of signification is necessary for the ontological enterprise because it is the only one that allows meaning to be *assigned* to a text, and recorded in a formal language other than the natural language, from which it can be extracted through automatic means following a schema like this (I have omitted the noise for the sake of simplicity):

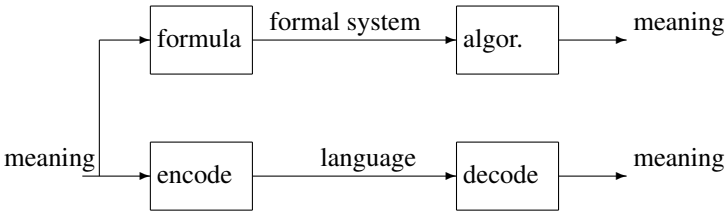

The conclusions of much of the linguistics and philosophy of language of the XX century, however, point in a different direction. There can be no meaning before language and independent of it: meaning can only exist within the categories and the strictures of language (Jameson, 1972). Not only meaning, but the signifying subject as well are a product of language (Lacan, 1982). There can be no pre-linguistic signification experience that belongs only to the author, because meaning can only be expressed in language, and language is a social instrument.

It is the act of reading, contextual and situated, that gives a text its meaning. The reader plays an active rôle in this meaning-creating activity: reading is not a one-directional activity in which a reader is imbued with meaning; it is a dynamic two-way process. It is an infinite process through which a frame of reference is created in which part of the text is interpreted, a text that changes the frame of reference and leads to

a different interpretation of the text, which changes the frame of reference and so on... This process of framing and interpretation is what reception theorists call the *hermeneutic circle* (Gadamer, 1975; Eco, 1979).

Mathematically, one could say that, according to this point of view, the meaning of a text is created by a complex interpretation act of which the text is only a boundary condition. And not even the only one, for that matter. Consider the following text, from (Eagleton, 1996):

Dogs must be carried at all time on the escalator.

Quite a puzzling bit of text and, if one were to see it out of any possible context one might have a hard time understanding what it means. What escalator? What dogs? Does it mean that you can't use this particular escalator unless you are carrying a dog? Does it mean that you can only carry dogs on an escalator, but not on a regular stairway?

Now let me give you some context: this sentence is printed on a plastic sign placed near an escalator on the London subway. I imagine that the sign now will be clearer. The sign is not a piece of reading material, but a warning directed to you, the average traveller. It tells you that, if you have a dog, you must carry it in your arms while you ride the escalator. It doesn't specify that you are allowed to ride the escalator even if you don't have a dog with you: the sign relies for that on your general knowledge about what you can and can't do while being on an escalator.

The sign doesn't tell you what will happen to you if you don't follow the directive, but you can well imagine that you will be punished in some way by some appointed authority and, given your general knowledge of these things, you can imagine that the punishment will not be harsh: you will not be sentenced to years in jail because of this, although you might receive a fine.

Notice that nothing of this rich context is implied in the text itself, simply because the text has to function in different ways in different context. You could use the same text, for instance, in a paper on semantics, where it would absolve a completely different function—that of an example. For, let's face it, is it really important for my example whether this text was ever really placed near an escalator in the London subway (although, as a matter of fact, it was)? In the context of clarifying a point this text would work, and its meaning would be quite clear, even if I were lying to you regarding the placement of the text.

But if the meaning of a text depends so crucially on the context in which it is read, then the general plan of ontology, to attach meaning to a text so that a simple algorithm can decode is in quite a bit of trouble. It should be stressed again that the limitations of ontology that we have highlighted are not a limitation of the particular logic that one might use to implement an ontology, nor of logic *per se*: the limitations are at a much more fundamental level. The discussion in this section problematizes the very possibility of representing meaning as an attribute of a text. According to this view, meaning is not *in* the text: a text is merely the boundary condition of a process that depends on the interpreter, his context, the linguistic community of which the interpreter is part, its discursive practices, etc. This doesn't necessarily imply that, for the purpose of meaning formation, the text can't be usefully represented using alternative means, including formal ones. As computing scientists, we are interested, pragmatically, in

situations in which reading and interpretation are somehow mediated by a computer, and alternative representations of the text may favor this mediation. What can't in any case be assumed is that representation is a representation of the meaning of the text, a representation from which meaning can be extracted in an a-contextual way by an algorithm.

4 Beyond ontology

Is there a way we can use these observations to deal with the semantics of documents, images, videos, and whatnot? I think there is, and doing so will give us a valuable alternative to the normative staticism of ontology.

Let us start with a fairly general theoretical model. We have said that the context in which a document is interpreted is essential to determine its meaning, that is, that the context changes the meaning of a text. We can also see things going in the opposite direction: the function of the semantics of a text is to change the context of the reader. If you are interested in novel, the context in which you look at American literature will not be the same after reading *Moby Dick*; if you travel on the London subway, your context will no longer be the same after you read that "dogs must be carried at all times". A document that doesn't change the context in which you act is, by definition, meaningless. We can express this situation with the following expression:

$$C_1 \xrightarrow{\mu(t)} C_2$$

where C_1 and C_2 are the contexts of the reader before and after interpreting the text, t is the text, and $\mu(t)$ is its meaning.

This is, as I have said, a very generic model, but we can use to start answering some questions. For one thing, *is it possible to formalize meaning?* The answer of our model is that it is possible only to the extent that it is possible to formalize context. If C_1 and C_2 are formally defined in mathematical terms, then, and only then, it will be possible to give a formal definition of the function $\mu(t)$.

At one extremum, we have the situation in which the context can be completely formalized. This is the case, for instance, in programming languages: here the context can be reduced to the *state* of a computer on which the program is run. The meaning of a program, from our point of view, is a function that transforms an initial state of the computer to a final one. In other words, if the text is a program and the context of its interpretation is a computer system, meaning reduces to the usual denotational semantics of a program.

At the other extremum we have the general semiotic context, which we know can't be formalized completely in symbols, that is, given that a computer is a symbol manipulation machine, it can't be formalized in a computer. Again, this doesn't entail that any attempt to use a computer (which, because of the characteristics of the device, requires a formalization of the context) is useless, but it does imply that no computing system can be semantically complete, so to speak, and that each computer system will require user interaction to contextualize access and allow signification to happen.

The properties of the "space of contexts" depend crucially on the properties of the representation of the context that we have chosen, and it is therefore difficult to say

something more about meaning is we don't impose some additional restriction. A reasonable one seems to be that we be capable of measuring the degree by which two contexts differ by means of an operation $\Delta(C_1, C_2) \geq 0$ such that, for each context C, it is $\Delta(C, C) = 0$. We don't require, for the time being, that Δ be a distance. Now the meaning of a document d in a context C can be defined as the difference that d causes to C:

$$\mu_C(d) = \Delta(\mu(d)(C), C) \tag{7}$$

Within this theoretical framework we can analyze, at least in the first approximation, various existing approaches, and devise ways to extend them. In this general scheme, the ontological approach to meaning can be synthesized as a constant function:

$$\perp \xrightarrow{\mu(d)} C \tag{8}$$

that is, ontology assigns a meaning to a document independently of the context in which the document is interpreted. This fact results, in our model, in the creation of a constant context, which depends only on the document and not on what was there before.

A very different point of view is that of *emergent semantics* (Santini & Jain, 1999; Santini, Gupta & Jain, 2001): in this approach, a highly interactive system allows the user and the system to organize the data in a way that highlights their contextual relations. The meaning of the data emerges as an epiphenomenon of this interaction. Emergent semantics does not work with one document at the time, but always with set of documents, since meaning always emerges from relations. Therefore, the meaning function μ will take as argument a suitable configuration D of documents. The user action is represented as an operator u, and the schema is the following:

$$C \underset{u}{\overset{\mu(D)}{\rightleftarrows}} C' \tag{9}$$

The context oscillates between C, which is the new contextual situation in which the user wants to end, and C', which is the context proposed by the computer with the access to the new documents. The semantic function is, in this case, the equilibrium of that cycle or, in other terms, the fix-point of the function $\mu(D) \circ u$.

These examples show how a pair of current—and very different—approaches to signification can be placed in this general theoretical framework. The full use of it, however, entails the use of some prior context that can specify, through the specification of the desired final context, the sought meaning, or that can be used as the starting point of a emergent semantic cycle.

Where can we obtain such a context? The interactions in which we are interested take place through a computer, but the computer is much more than just the computing instrument of this interaction: it is a repository of documents, pictures, videos, and in general of all those elements that are the support and the product of the activity in which the interaction with the data is placed. Most interactions with a data repository take place as part of an activity that was developed on a computer, and whose witnesses are the files that are stored and organized in the computer. The contents of the files related to the current activity, any additional information about such activity, and the organization of these files are an ideal place to look for the context that we need.

Note that, consistently with our theoretical framework, we can no longer talk of "the meaning" of a document: independently of the way in which the context and the document are represented, the formation of meaning only takes place through the interaction between the two (mediated, always, by the user); it is the change in the context operated by the document that constitutes meaning. If the representation chosen for the document and the context is a logic formalization, the result is not too different from Jain's "personal ontology" (Scherp & Jain, 2007).

4.1 A demise of search?

In the previous pages, I made an effort to avoid the use of terms like "query" or "search", preferring more neutral locutions such as "interaction with the data". The reason for this is that the use of context to define semantics problematizes the notion of a generic "search" operation. The context in which the data assume a meaning is always given by a human activity, not only because of the collateral data that are the support and the product of the activity but also (an aspect that I haven't considered in this paper) because of the discursive practices of that activity, an aspect that Jain and Scherp called *expressive semantics* (Scherp & Jain, 2007).

The context, and especially the discursive practices in it, is so important that one can't quite pin down a single generic activity called "search" that can be applied to all human activities: looking for a text to elucidate a point to ourselves while writing a scientific paper is not the same activity as looking for something to read before going to bed. Looking for a picture to reinforce a point in one's presentation is not the same thing as arranging one's vacation photographs. The activity of "search", which in data bases is unique and well defined, breaks down here in a constellation of different modalities of interaction with data structures. The techniques that one uses in these different activities are largely the same, but the dynamics of the search, the type of interactions that takes place, the notion of what counts as an acceptable result, are specific of the human activity that surrounds the access to the data.

5 Conclusions

My purpose in this paper was unapologetically polemic. Ontologies are being hailed as the cornerstone of the constituenda *semantic web* and as the only viable solution to the "semantic gap" (Santini & Jain, 1998) based on justification that—when they are offered at all—are, in my opinion, mislead and insufficient. Most of the solutions of the semantic web come from the area of knowledge engineering, whose success in the representation of meaning in unrestricted domains has been questionable, but there is the unspoken assumption that, simply by replicating these techniques on a larger scale, one can make them successful. I believe that this is an illusion.

I have argued in this paper that meaning is not "in" a document, nor is it a property of a document, but that it comes from an interpretative act, contextually and culturally situated. If these arguments are valid, then there are serious doubts on the feasibility of the ontological programme, at least in its semantic incarnations. This, as I have shown, doesn't imply that the instruments and the techniques of ontology can't have a place in

a system that deals with semantics, but it does mean that an ontology can't in any way be considered as the repository of the meaning of a document or of the semantics of a community.

References

[Berners-Lee, Hendler & Lassila, 2001]Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American*, *284*.

[Eagleton, 1996]Eagleton, T. (1996). *Literary theory, an introduction*. Minneapolis:Minnesota University Press.

[Eco, 1979]Eco, U. (1979). *The role of the reader*. Bloomington: University of Indiana Press.

[Feynman, 1988]Feynman, R. (1988). *What do* you *care what other people think?* New York: Norton.

[Gadamer, 1975]Gadamer, H.-G. (1975). *Truth and method*. London.

[Geckeler, 1976]Geckeler, H. (1976). *Semántica estructural y teoría del campo léxico*. Madrid:Gredos. Spanish translation of *Strukturelle Semantik und Wortfeldtheorie* by Marcos Martínez Hernández.

[Greimas, 1966]Greimas, A. J. (1966). *Sémantique structurale*. Paris: larousse.

[Hakimpour & Geppert, 2001]Hakimpour, F. & Geppert, A. (2001). Resolving semantic heterogeneity in schema integration: an ontology based approach. In *Proceedings of the International Conference on Formal Ontologies in Information systems*.

[Jameson, 1972]Jameson, F. (1972). *The prison-house of language*. Princeton, NJ.

[Lacan, 1982]Lacan, J. (1982). *Escrits: a selection*. New York: Norton.

[Moya, 2004]Moya, V. (2004). *La selva de la traducción*. Madrid:Cátedra.

[Nida, 1964]Nida, E. (1964). *Towards a science of translating*. Leiden: E.J. Brill.

[Santini, Gupta & Jain, 2001]Santini, S., Gupta, A., & Jain, R. (2001). Emergent semantics through interaction in image databases. *IEEE Transaction on Knowledge and Data Engineering*.

[Santini & Jain, 1998]Santini, S. & Jain, R. (1998). Beyond query by example. In *Proceedings of ACM multimedia*.

[Santini & Jain, 1999]Santini, S. & Jain, R. (1999). The "el niño" image database system. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, (pp. 524–9).

[Scherp & Jain, 2007]Scherp, A. & Jain, R. (2007). Towards an ecosystem for semantics. In *Proceedings of ACM Multimedia systems*.

[Stich, 1983]Stich, S. (1983). *From folk psychology to cognitive science*. Cambridge, MA: MIT Press.

[Uschold & Gruninger, 1996]Uschold, M. & Gruninger, M. (1996). Ontologies: principles, methods and applications. *Knowledge engineering review*, *11*(2).

[Vermeer, 1983]Vermeer, H. J. (1983). Ein rahmen eür eine allgemeine translatontheorie. *Aufsätze zur Translationtheorie*, 48–61.