An oddly-positioned position paper on context and ontology

Simone Santini

Universidad Autónoma de Madrid, Madrid, Spain

Abstract

This paper is a theoretical analysis of formal annotation and ontology for the expression of the semantics of document. They are found wanting in this respect, not only for technical reasons, but because they embody a fundamentally misunderstood model of the process of signification. I propose an alternative model in which the interpretation context plays a fundamental rôle, and briefly discuss it and its current technical embodiment.

1 Introduction

This is a position paper on meaning, on the process of signification, of its importance for data processing problems, and on the (in my view) misguided ways in which the problem has been heretofore regarded within the computing community. Quite a lot of stuff to fit into a conference communication. I called the paper "oddly-positioned" in part, I admit it, for the *divertissement* of the word play in the title but, a bit more substantially, because the position that I am defending here is not terribly popular. Computing scientists and professionals are oriented towards a rather specific way of considering meaning and, especially under the pressure of lavishly funded eforts such as the "semantic web", one has the impression that views that dare to question the received wisdom are not entirely welcome. I do, of course, hope of being mistaken because, as it happens, I regard the common way of considering semantics a rather naïve one, one that ignores what we have come to know in the last century about semantics.

Some people might find my arguments a bit on the philosophical side, but I shall offer no apologies for that. Semantics is essentially a philosophical issue. The computer scientist who design programs for particle physics can't help but coming in touch with a bit of quantum physics, because that is what particle physics is about. Analogously, computing scientists who want to work on semantics can't help but deal with (as Sam Spade would have put it) the stuff semantics is made of: philosophy.

2 Ontology for the representation of semantics

One solution to the problem of semantic data processing, quite popular in the computing milieu these days, entails the *formal annotation* of the semantically problematic data. Annotating the data with a formal language serves, in this vision, two purposes. On the one hand, it is supposed to represent a "semantic" annotation: an annotation that records the *pure* meaning of the data, distilling it from the super-structures and the uncertainties of natural language. On the other hand, being formal, the language of the annotation allows one to make the same semantic assumptions that one does in standard data bases, namely that semantics can be formalized using method similar to the formal semantics of programming languages. This semantic programme is based on the assumption that the semantic problems that we face *vis à vis* our data are not due to an inherent characteristic of the data themselves, but to the defective way in which the meaning of the data is carried by the language in which they are expressed. It is assumed that, at the core, web data are not that different from the ones in data bases, in that they do have an inherent meaning that can be made sense of by an algorithm. The problem is that these data are expressed in semiotic systems (from natural language to images and video) that make it difficult to *extract* this meaning. (The word "extract" does a lot of work here, since it underlies one foundational assumption of this approach, namely that meaning is an inherent quality of the data.) It is assumed, however, that meaning pre-exists the text (logically, not necessarily chronologically), that can be expressed in a suitable formal system and attached to the text in a form that can be understood by an algorithm.

The "hubs" of systems organized along these lines are the so-called *ontologies*, collections of axioms that supposedly capture the semantics of the terms used in a certain text and make the text amenable to treatment using standard data base techniques. The most diffused and celebrated (and lavishly funded) form of this approach can be found in the so-called *semantic web*.

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.

One of the problems of this point of view is that if you hold it as a theory of meaning is very hard to get out of radical nativism. That is, this model leads you almost automatically to admit that all concepts are innate, and almost none of them is acquired. This is quite absurd, of course: Kant would have gladly admitted that we have the concepts of space and time without learning them but, as Fodor says

> [...] how could DOORKNOB be innate? *DOORKNOB*, of all things![1]

Fodor escapes this trap somehow, but this theory of meaning has an even harder time explaining the deep different in the creation of concepts between different languages.

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). 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 [14] and structural semantics [4], 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[2]. 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 ad-

---

[1] ibid. p. 123, emphasis in the original.

[2] This example is an extension of a similar one in [3].

jectives, 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*.

Italian	Spanish	French
	añejo	
vecchio	viejo	vieux
anziano	anciano	âgé
	mayor	
	antiguo	ancien
antico		antique

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.

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[3]. 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.

---

[3]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 [9]. Rather, different translations are possible depending on the rôle that the translation will play in the receiving culture.

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 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 [10]. 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 [15] 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.

2.1 Ontology as non-contextual 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[4]. Where does it come from? For such an hypothesis to work, meaning must exist prior to text and independently of the language in which it is expressed. The scheme is pretty much that of a communication channel.

meaning → encode — language — decode → meaning
noise ↑ (encode)   noise ↑ (decode)
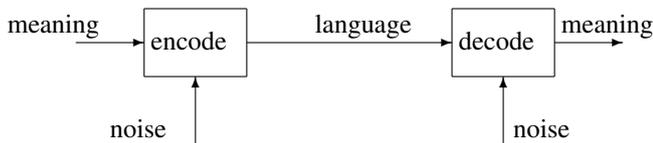
The origin of the communicative act is a *meaning* that resides with the author, and that the author wishes to express in a permanent text. This meaning is a-historical, immutable, and pre-linguistic. In order to communicate meaning, the author translates it into the shared code of language, and sends 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):

[4]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.

formula — formal system — algor. → meaning
meaning → encode — language — decode → meaning
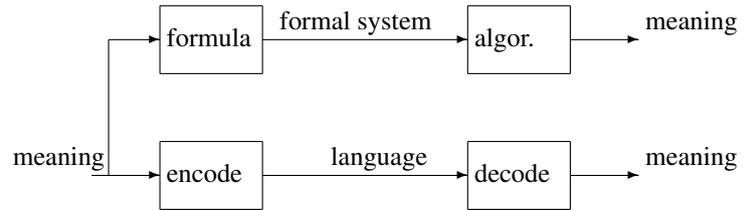
The conclusions of much of the linguistics and philosophy of language of the XX century, however, point in a different direction. There can be no meaning before language and independent of it: meaning can only exist within the categories and the strictures of language [5]. Not only meaning, but the signifying subject as well are a product of language [8]. 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* [2, 1].

Lest should you think that all this applies only to literature and not to the prosaic world in which ontology operates, let me take the most prosaic example I can think of: a sign on a door that says "trespassers will be prosecuted." The hermeneutical activity necessary to understand this sign is considerable. I must understand, for instance, that this sign is not informative in the sense that a newspaper headline is: I am not being informed that there have been trespassers somewhere and that they will be prosecuted sometime in the future: in western societies at least, information of this kind is not written on signs hanging from doors, especially if the sign is made of plastic or wood (and therefore is durable) and the writing is not dated. Such a sign typically is a prohibition, a threat; the word "trespasser" refers to me (the reader) in case I decide to walk through the door, and it threatens me of prosecution if I do so. The threat also implies that prosecution is likely to result in punishment. I must understand that trespassing in this context means to cross *this* door, not some door in the palace of the king of Siam. I have to have a general knowledge of private property to understand that preventing people from entering a building is one of the rights that society grants to propri-

etors (while, for instance, preventing people from looking at the building is in general not such a right), that there are authorities that will guarantee the respect of these rights, that they will punish people who infringe these rights, that the sign has been placed there with their tacit approval, and so on...

None of these elements, necessary for the interpretation of the sign, is *in* the text: they must be supplied by a specific situation. The text here takes meaning by being situated (viz. placed in a situation: a door on a building rather than, say, a shelf on a store that sells signs, a situation in which the text would have a completely different meaning) and in a certain relation with other texts that are not present, namely the political discourse that regulate private property, the speech through which certain customs have been implanted in the reader, and so on. Finally, all this linguistic discourse and all this hermeneutuc activity rest on a substratum of human practices and action: the political relation of power between authority and citizens, and the fact that in order to understand punishment one must understand pain (psychological pain, at least).

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.

3   Context-based retrieval

In the light of the previous observations, it seems clear that one can't hope to simply encode the semantics of a document in manner independent of the act of access: meaning is created anew with each data access, and it is a result of that operation. Just like the tree falling in a deserted forest that makes no noise (although it does provoke acoustic waves), so a text, when it is not accessed, has no meaning (although it has the potential for signification). Our problems, then, are basically three: given a data access situation, we must (i) find a suitable context in which the data access is situated, (ii) find ways to formalize this context, at least to a certain degree (we are, after all, computing scientist, and we can only work with what we can formalize), and (iii) find ways in which the context can interact with the data to generate meaning.

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:

$$\mu_C(d) = \Delta(\mu(d)(C), C) \tag{1}$$

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:

$$\perp \xrightarrow{\mu(d)} C \tag{2}$$

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* [13, 12]: 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:

$$C \underset{u}{\overset{\mu(D)}{\rightleftarrows}} C' \tag{3}$$

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 $\mu(D) \circ u$.

Our next problem is how to capture ongoing activities, how to represent them and, for what is possible, formalize them, in such a way that they can be used as a basis for data access. In general, of course, this is impossible. If a person is, say, shopping for detergent and wants to search the internet for brands with certain characteristics, there is very little hope that we can represent the activity "shopping for detergent" in a computer system: we are in this case in the presence of a physical activity that leaves no *digital trace*, so to speak.

On the other hand, a significant number of daily activities are, for many of us, executed on or with the aid of a computer. In this case, they do leave a digital trace, one that can be recorded and used as a context for a search carried out as part of that activity. Suppose I am preparing a presentation for a conference to which I had submitted a paper and that, during this process, I need to clarify a point or to look for an illustration for my presentation. In order to prepare my presentation, I have created a document in a directory (let us say the directory *presentation*) where I have possibly copied some documents that I thought might be useful. This directory is likely to be placed in a hierarchy, something like this:
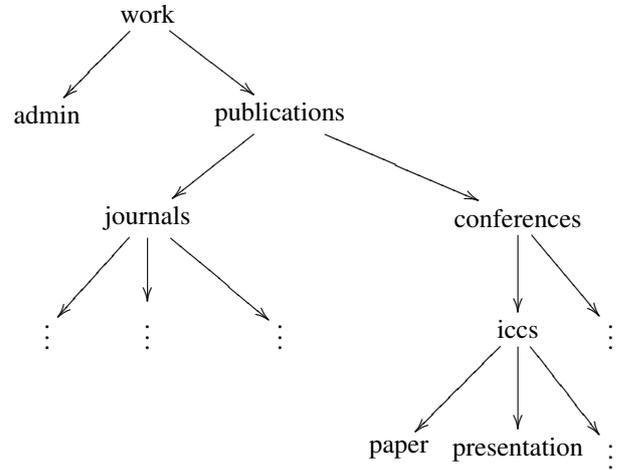
Its sibling directories will contain documents somehow related to the topic at hand although, probably, not so directly as those that can be found in the work directory. The siblings of the conference directory (and their descendants) will contain documents related to my general area of activity, although not necessarily directly related to the topic of the presentation. This project, in its context search component, will look for ways to use this information in order to direct and focus the search. This information will constitute the context of the search. One consequence of this point of view is important enough to be noted from the outset: data access is no longer an independent activity, but can take place only in the context of a certain activity.

3.1 Context representation

The problem of representing context and, most importantly, how to make it interact with the documents, are still

largely unexplored, at least as far as computing science goes, and it is not clear in which direction one should look for a proper representation of context.

As a first step, one might consider the use of techniques from information retrieval. Here, I will give some pointers on the possible use of a model based on a vector space representation of word contexts [6], and a self-organizing map to give a non-linear form of latent semantics. This is the model that we are currently using in our activity. Word context are groups of sequential words that occur in a text. In this case, they are more representative than single words because they capture, statistically, word co-occurrence, which is a strong indicator of the semantics in which a word is used in a given context (if the word "bank" co-occurs with "investment", it is likely to mean a financial institution, if it co-occurs with the word "river", it is likely to indicate the border of a body of running water, and so on). These co-occurrences will be represented in a suitable feature space and a Self-Organizing Map [7] will be used to cluster and represent its contents, using again fairly standard techniques. As a starting point for the representation of context, a self-organizing map appears to be a suitable choice for a number of reasons.

i) The map represents a sort of non-linear latent semantic subspace: it capture the statistically significant relations between terms in a given context.

ii) The learning algorithm gives us an obvious way to start including the structure of the directories into the context representation: learning may not be limited to the documents contained in the working directory, but can include those contained in the children/siblings/parents. Moreover, by varying the fraction of times the documents in each one of these directories are presented we can give more or less importance to certain structural relations.

iii) The map can constitute a query formalism. Given that the map represents the context, we can create a query by modifying it (i.e. distorting it) towards certain terms using Kohonen's training algorithm. In the previous case, the word "bank" would *pull* the part of the map closer to it towards it, and that part of the map would contain the context in which the word appears. The difference between the initial map and the deformed one is the meaning of the text that we are after. We are studying efficient ways to express this difference as a query.

iv) The map, being geometric in nature, suggests a way to extend the context representation to multi-media document or, at least, to documents containing images. We can extract image features that can be represented in geometric spaces [11] and derive the direct sum of the space of words and the space of features. This should allow the map to capture any statistical regularity, in the document corpus, between certain word combinations and certain image features. Note that if the features are extracted from regions of the image, one can use feature context techniques similar to the word context used for text, thus seeking statistical regularities between co-occurrences of words, and co-occurrences of localized features.

While the techniques used in this approach are fairly standard, its novelty is that, in this case, we are not using them in order to represent the data base in which the search is to be done, but to represent the environment from which the query originates. Note however, that point iii) and iv) above (the algorithm for the context distortion and that to translate the context difference into a query, and the deployment of the map in a multi-feature space, to seek statistical regularities between co-occurrences of words and co-occurrences of features) represent significant new technical challenges.

4 Words of parting

I have argued that formal annotation, and the general ontological programme that comes with it, might not be the proper way to consider the problem of the meaning of the data and, in general, to frame the issues related to semantics. This is not, I suspect, a popular position, and there are a few reasons that contribute to its unpopularity and to the exclusivity of the attention given to annotation and ontology.

First, there is the pull of certain common sense philosophy. We can look at texts, read them, and make sense of them, and it seems natural to interpret this act as unlocking the meaning that is *in* the text. After all, if I don't know from what gate does flight 354 to New York leave, and I read the announcement board of the airport, I end up knowing it. It is easy to model this situation as a transfer of a specific information (viz. that the flight leaves from gate C34) from the announcement board to me. The error is the failure to recognize that this is a limit case, namely a case in which the external context is so constraining that the readong of the symbol "C34" can basically have only an interpretation, and to extend the same model to the common situation, the one in which the interpretation context plays a much more important rôle. I have given my arguments (convincing, I hope) why I believe that this position represents a gross philosophical simplification, and I believe that it will ultimately result in the sterility of semantic computing.

Second, there is the normal inertia of the established position upon which truckloads of money have already been invested. The agencies and companies that have invested in annotation and ontology are understandably nervous when

they hear that they might have bet on the wrong horse. But, of course, if they have indeed bet on the wrong horse (and I believe they have), the sooner they know, the better.

Finally, there is a point related to the economy of the commercial web (which, unlike ten years ago, regrettably represents the vast majority of the web today). The model of meaning assumed by the semantic web is very appealing to web companies because, if meaning is inherent in a text, it can be owned, bought, and sold. In the ontology view, meaning is a property of the author (or of the organizations that bought it from the author), a property that can be exchanged with the reader using the currency of language. This "market" view of meaning open disturbing (or interesting, depending on what side you are looking at them) possibilities of copyrighting meaning, patenting meaning, and in general posing commercial restrictions to the free exchange of meaning. We have seen in the last few years the desperate (and, so far, successful) attempts of the industry to restrict all available intellectual property laws. The ontological model of meaning gives them a whole new area in which copyright can be enforced. For those of us who believe that the web should be a common good, in which commercial interests should never replace the free exchange of ideas, this is not an appealing perspective.

Technically, this paper has presented the outline of a different model of meaning, one in which the reader's context plays a preponderant rôle. I have presented a simple framework in which we are currently experimenting with this model, a framework that in the future will be extended in different directions: on the one hand, the integration in this framework of more formal representations, at least for those parts of the context that can be formalized; on the other hand, the development of suitable data base techniques to make this kind of query efficient. Our purpose will be, on one hand, to build a context-based data access client (configured as a plug-in to some word processing or presentation program, if possible) to make context based retrieval on general web sites and repositories and, on the other hand, to build a context-based access server. The latter will be akin to the servers built for search engines such as yahoo or google but, while these servers do not coöperate with the user's computer (apart from the elementary communication necessary to retrieve the query and return the results), the server that we consider here will be integrated with the user's computer from which it will derive the current context, and with which it will coöperate to support interaction.

References

[1] Umberto Eco. *The role of the reader*. Bloomington: University of Indiana Press, 1979.

[2] Hans-Georg Gadamer. *Truth and method*. London, 1975.

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

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

[5] Frederick Jameson. *The prison-house of language*. Princeton, NJ, 1972.

[6] S. Kaski. Computationally efficient approximation of a probabilistic model for document representation in the WEBSOM full-text analysis method. *Neural Processing letters*, 5(2), 1997.

[7] T. Kohonen. *Self-organizing maps*. Heidelberg, Berlin, New York:Springer-Verlag, 2001.

[8] Jaques Lacan. *Escrits: a selection*. New York: Norton, 1982.

[9] Virgilio Moya. *La selva de la traducción*. Madrid:Cátedra, 2004.

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

[11] S. Santini. *Exploratory image databases: context-based retrieval*. San Diego:Academic Press, 2001.

[12] Simone Santini, Amarnath Gupta, and Ramesh Jain. Emergent semantics through interaction in image databases. *IEEE Transaction on Knowledge and Data Engineering*, 2001.

[13] Simone Santini and Ramesh Jain. The "el niño" image database system. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, pages 524–9, 1999.

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

[15] Hans J. Vermeer. Ein rahmen eür eine allgemeine translatontheorie. *Aufsätze zur Translationtheorie*, pages 48–61, 1983.

Madrid, March 2008