A technological solution for Distributed Knowledge Management

M. Bonifacio[1,2], P. Bouquet[1,2], G. Mameli[2], and M. Nori[2]

[1] Dept. of Information and Communication Tech. – University of Trento (Italy)
[2] Istituto per la Ricerca Scientifica e Tecnologica, Trento (Italy)

Abstract. Distributed Knowledge Management is an approach to Knowledge Management based on the principle that the multiplicity (and heterogeneity) of perspectives within complex organizations is not be viewed as an obstacle to knowledge exploitation, but rather as an opportunity that can foster innovation and creativity. Despite a wide agreement on this principle, most of current KM systems are based on the idea that all perspectival aspects of knowledge should be eliminated in favor of an objective and general representation of knowledge. In this paper we propose a peer-to-peer architecture (called KEx), which embodies the principle above in a quite straightforward way: (i) each peer (called a K-peer) provides all the services needed to create and organize "local" knowledge from an individual's or a group's perspective, and (ii) social structures and protocols of meaning negotiation are defined to achieve semantic coordination among autonomous peers (e.g., when searching documents from other K-peers).

1 Introduction

Distributed Knowledge Management (DKM), as described in [4], is an approach to KM based on the principle that the multiplicity (and heterogeneity) of perspectives within complex organizations should not be viewed as an obstacle to knowledge exploitation, but rather as an opportunity that can foster innovation and creativity.

The fact that different individuals and communities may have very different perspectives, and that these perspectives affect their representation of the world (and therefore of their work) is widely discussed – and generally accepted – in theoretical research on the nature of knowledge. Knowledge representation in artifi cial intelligence and cognitive science have produced many theoretical and experimental evidences of the fact that what people know is not a mere collection of facts; indeed, knowledge always presupposes some (typically implicit) interpretation schema, which provide an essential component in sense-making (see, for example, the notions of context [13, 5, 10], mental space [9], partitioned representation [8]); studies on the social nature of knowledge stress the social nature of interpretation schemas, viewed as the outcome of a special kind of "agreement" within a community of knowing (see, for example, the notions of scientifi c paradigm [11], frame [**?**]), thought world [**?**], perspective [2]).

Despite this large convergence, it can be observed that the high level architecture of most current KM systems in fact does not reflect this vision of knowledge (see [3, 4, **?**] for a detailed discussion of this claim). The fact is that most KM systems embody the assumption that, to share and exploit knowledge, it is necessary to implement a process of ‘knowledge extraction and ‘refi nement”, whose aim is to eliminate all subjective and contextual aspects of knowledge, and create an objective and general representation that can then be reused by other people in a variety of situations. Very often, this process is fi nalized to build a central knowledge base, where knowledge can be accessed via a knowledge portal. This centralized approach – and its underlying objectivist epistemology – is one of the reasons why so many KM systems are deserted by users, who perceive such systems either as irrelevant or oppressive [7].

In this paper we propose a peer-to-peer (P2P) architecture (called KEx) which is coherent with the vision of DKM. Indeed, P2P systems seem particularly suitable to implement the two core principles of DKM, namely the principle of autonomy (communities of knowing should be granted the highest possible degree of semantic autonomy to manage their local knowledge), and the principle of coordination (the collaboration between autonomous communities must be achieved through a process of semantic coordination, rather than through a process of semantic homogenization) [4]. In KEx, each community of knowing (or Knowledge Nodes (KN), as they are called in [**?**]) is represented by a peer, and the two principles above are implemented in a quite straightforward way: (i) each peer provides all the services needed by a knowledge node to create and organize its own local knowledge (autonomy), and (ii) by defi ning social structures and protocols of meaning negotiation in order to achieve semantic coordination (e.g., when searching documents from other peers).

The paper goes as follows. In section 2, we describe the main features of KEx, and argue why they provide a useful support to DKM; in 3, we describe its implementation in a peer-to-peer platform called JXTA; fi nally, we draw some conclusions and future work.

2 KEx: a P2P architecture for DKM

KEx is a P2P system which allows a collection of KNs (individuals or groups) to search and provide documents on a semantic basis without presupposing a beforehand agreement on how documents should be categorized, or on a common language for representing semantic information within the system. In the following sections, we describe the high-level architecture of KEx, and explain what role each element plays in a DKM vision.

2.1 K-peers

KEx is defi ned as a collection of peers, called a K-peers, each of which represents a KN, namely a semantic perspective on a given body of knowledge. Each K-peer can play the two main roles: *provider* and *seeker*. A K-peer acts as a provider when it “publishes” in the system a body of knowledge, together with an explicit perspective on it (called a *context*, e.g. a topic hierarchy used to categorized local documents); a K-peer acts as a seeker when it searches for information by making explicit part of its own perspective, and negotiates it with other K-peers.

Each K-peer has the structure shows in the Figure 1. Below we illustrate the main modules and functionalities.

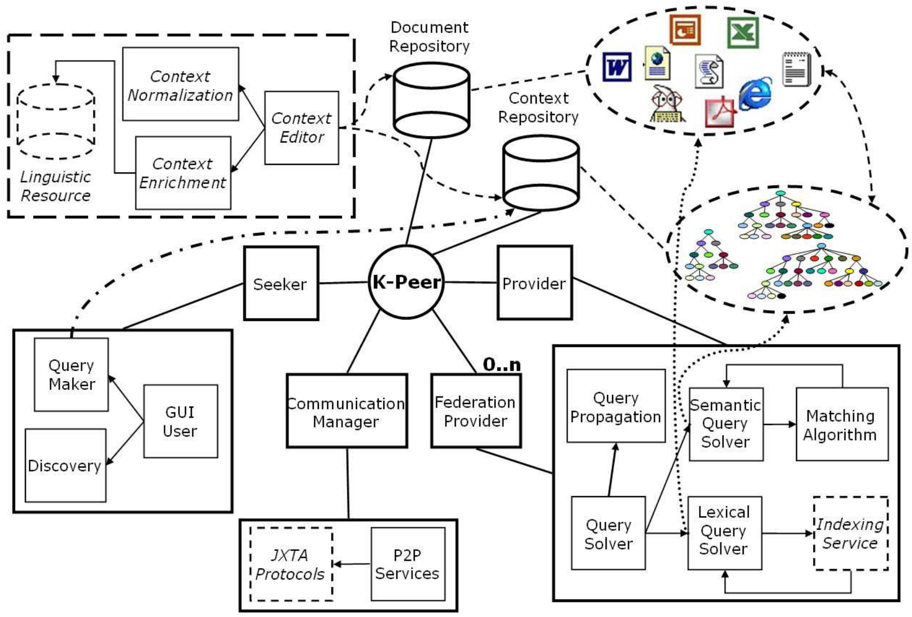

Fig. 1. The KEx's main components

Document Repository. A *Document Repository* is the place where each KN stores its own knowledge. We can image a private space in which the KN maintains its document and data, maybe using a local schema (e.g., a file-system structure, or a database schema) or a document management system in order to organize and access them.

Context Repository. Following [1], we define a context as a partial and approximate representation of the world from an individual's or a group's perspective. The reason why we adopt this notion of context is that it provides a robust formal framework (called Local Models Semantics [10]) for modeling both contexts and their relationships.

In order to use contexts in KEx, we adopted a web-oriented syntax for contexts, called CTXML. It provides an XML-Schema specification of context for document organization and classification[3].

In KEx, each context plays the role of a category system for organizing and classifying the documents (or any other kind of digital information) stored in a KN's document repository. Each peer can use more than one context to classify local documents; the contexts of a K-peer are stored in a *context repository*.

From the standpoint of DKM, contexts are relevant in two distinct senses:

[3] Currently, contexts are trees, whose nodes are labelled with words defined in some name space. Arcs are Is-A, Part-Of or generic relations between nodes. Details can be found in [6].

- on the one hand, they have an important role within each KN, as they provide a dynamic and incremental explicitation of its semantic perspective. Once contexts are reified, they become cognitive artifacts that contribute to the process of perspective making [2], namely the consolidation of a shared view in a KN, continuously subject to revision and internal negotiation among its members;
- on the other hand, contexts offer a simple and direct way for a KN to make public its perspective on the information that that KN can provide. Therefore, as we will see, contexts are an essential tool for semantic coordination among different KN.

It is important to observe that contexts provide only a common syntax for classification structures. Indeed, we could see them as a language for wrapping any classification structure (e.g., like directory systems, databases schemas, web directories). This means that in principle people can continue to work with their preferred document management system, provided it can be wrapped in CTXML.

Context management module. The context management module allows users to create, manipulate, and use contexts in KEx. The module has two main components:

- **Context editor**: provides users with a simple interface to create and edit contexts, and to classify information with respect to a context. This happens by allowing users to create links from a resource (identified by a URI) to a node in a context. Examples of resources are: documents in local directories, the address of a database access services, addresses of other K-peers that provide information that a KN wants to explicitly classify in its own context.
- **Context browser**: is part of a GUI Seeker Admin component (Figure 1). It allows users to navigate contexts from the context repository. The main reason for navigating a context in KEx is to build queries. The intuitive idea is that users can make context dependent queries (namely, from their perspective) by selecting a category in one of the available contexts. Once a category is selected, the context browser builds a *focus* – namely a contextual interpretation of the user's query – by extracting a semantically relevant subset of the context to which the category belongs (see [12] for a formal definition of focus). The focus is then used as a basis for meaning coordination and negotiation with other K-peers during the search.

2.2 Roles of K-peers in KEx

Each K-peer can play two main roles: seeker and provider. Their interactions are represented in Figure 2, and described in detail in the following two the sections.

Seeker As a seeker, a K-peer allows users to search for documents (and other information) from other K-peers and federations (see Section 2.3). The seeker supports the user in the definition of context-dependent queries through the context browser. A query is composed by a query expression and a focus. A query expression is a list (possibly empty) of one or more keywords provided by a user; a focus is a portion of a context determined by the category that the user has selected. Moreover, the seeker provides the

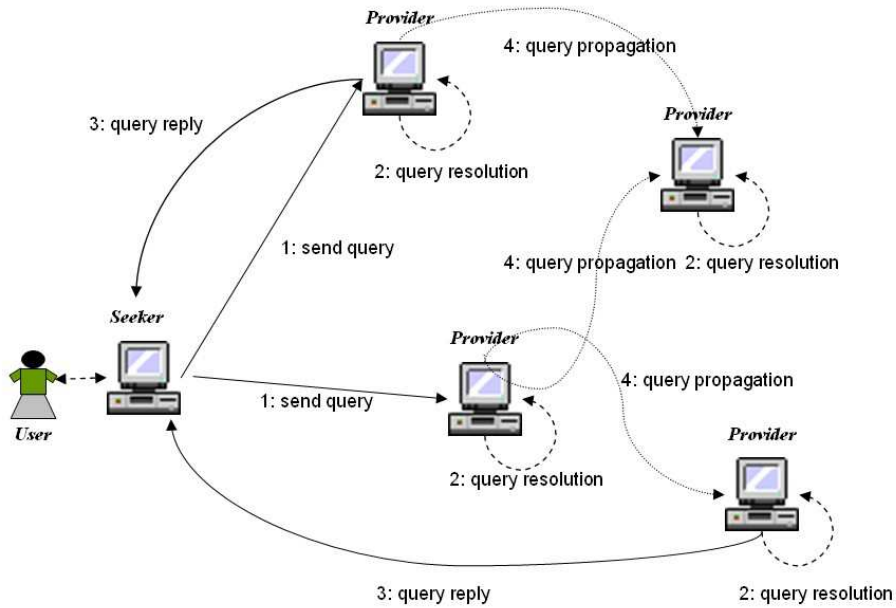

Fig. 2. The KEx system: interaction between Seeker and Provider roles

discovery mechanism, used to find resources to which the query is to be sent. The user decides to send the query to some of the available K-peers and federations. When the user submits the query, the seeker activates a session associated to that query (there can be only one active session for each seeker). In a session, a seeker can receive several asynchronous replies from the providers which resolved the query (through the meaning negotiation protocol, see below) and called back the seeker. The results returned to the user are composed by the aggregation of all the results received from the providers; each result is made up of a list of document descriptors (i.e., name of the document, short description, and so on). Each result is presented together with the part of context that the provider has matched against the current query. This relationship between contexts can be used as an opportunity for learning relationships across contexts of different KNs that the seeker can store and reuse for future queries (see 2.3). Finally the seeker allows the user to access the K-peer download service (Document Provider component in the figure 1). If the user finds one or more interesting documents, it can contact the peer that has the document and, if possible, download it.

Provider. The provider is the second main role in the KEx system. It contains the functionalities required to take and resolve a query, and to identify the results that must to be returned to the seeker. When a K-peer receives a context-dependent query (keywords + focus), it instantiates a provider, configured to use a set of contexts and some documents (a particular directory), to resolve the query in two ways:

– **Semantic search**: using for a context matching algorithm [12], which finds relations between a providers context and the querys focus. More specifically, the matching algorithm searches categories whose associated contextual information in the providers contexts (called a provider's focus) match (in a sense defined in [12]) with the querys focus. Provider's focuses, together with the URIs of the associated resources, are returned to the seeker. If the provider's focus contains also links to resources in other K-peers, the provider propagates the query to the K-peers that own the information.
– **Lexical search**: using an indexer to search for the occurrence of specific keywords into the set of documents owned by the provider.

If the query is composed only from keywords, the provider will use only the lexical search; if it is composed only from focus, the provider will use only the semantic search. If both are available, the final result will be the merge of the semantic and lexical results that the provider will send to the requesting seeker.

2.3 K-Services

KEx provides a collection of services which have an important role in supporting knowledge exchange (that's why we call them K-services). The more important among them are described in the following sections.

Context Enrichment. The *Context Normalization and Enrichment* allows to perform a linguistic normalization on user defined contexts (e.g., deleting stop words, tokenizing, tag part-of-speech, etc.) and to use external linguistic resource (like WordNet) to add semantic information to the categories in a context.

Normalization uses pretty standard NLP techniques, so we do not discuss it here. As to enrichment, this component – described in detail in [12] – is applied offline to a context defined by a user. It takes a context (in the sense of [6]) as input and returns an enriched context as output. It is important to say why enrichment is not equivalent to introduce a shared semantics in KEx. Indeed, the intuition is that the meaning of a category in a context has two components:

– the first is the linguistic component, which means that the words used as category labels have a standard meaning (or, better, a set of meanings) in a "dictionary". This helps, for example, to distinguish between "apple" as a fruit and "apple" as a tree;
– the second is a sort of pragmatic component, which is given by its position in a context (e.g., its path in a category tree). This helps in understanding what the user means on a particular occasion with a word (e.g., "apple" in a path like "computer/software/apple" is different from "apple" in a path like "computer/hardware/printers/apple", even though "apple" has the same dictionary meaning').

The first component of meaning is dealt with by the normalization and enrichment phase; the second is dealt with by the meaning negotiation protocol, and cannot be computed beforehand, as it expresses a user's perspective in making a query (so this is the more "perspectival" aspect of meaning).

It is extremely important to notice that different linguistic resources can be used to enrich a context. So far, we've been using WordNet, but there's no reason why other resources (like CYC or any other ontology) can't be used to replace WordNet. Of course, this means that one gets different context's versions (i.e., context translation in different languages) and a emerging of spontaneous community of K-peer that "speak" the same language.

K-federation. A *K-federation* is a group of K-Peers that agree to appear as a unique entity in respect to K-peers that perform a search. In other words, a Seeker can send a query to a K-federation, and the query will be forwarded to each federation member. As a consequence, the response of the K-federation to a query is the same as if the query was sent directly to all the members of the Federation (even if in the returned result set is specified if the Provider answers as a member of a Federation).

Each Federation can be though as a "social" aggregation of K-peers that display some synergy in terms of content (e.g., as they provide topic-related content or decided to use the same linguistic resource to create a common "vocabulary", thus providing more homogeneous and specific answers), quality (certify content) or access policies (certify members).

To become a member of a K-federation, a K-Peer must provide a K-federation Service (quite similar to that required by the Provider role) that implements the required federation protocol (reply to queries sent to the K-federation) and observes the federation membership policy.

Discovery. Discovery is a mechanism that allows the user to discover resources in the P2P network. The user discovers K-peers or K-federations available in the network in order to contact them and send them the queries. A peer advertises the existence of resources publishing a XML document (advertisement). In the KEx system we advertise two type resources:

– **K-peers** that have a provider service to solve queries. The main elements of the advertisement are a description of the peers contexts and an address to contact the peer in order to send it a query or retrieve documents;
– **K-federations**, namely set of peers that have a federation service to solve similar queries. The federation assures that a query sent to a federation is propagated to all active peers that are member of the federation. In this case the main elements of the advertisement are the federation themes, how to contact, how to join.

In order to discover resources in a P2P network, a peer sends a discovery request to another known peer, or sends a multi-cast request on the network, and receives responses (list of advertisements) that describe the available services and resources. It is possible to specify a searching criteria (such as a keyword or textual expression) that is matched against the contents provided by the advertisement related to each peer or federation description.

Query Propagation. This functionality allows the KEx system to satisfy the requirements of a better and more specific result set and, at the same time, showing a adaptability to a highly dynamic environment.

When a Provider receives a query, it can decide to forward it to another Provider that is considered "expert" about the query's topic. To decide to which peers the query is to be forwarded, a peer has two possibilities:

– physical "proximity": the query will be sent to peers known through the discovery functionality. This way, peers or providers that are non directly reachable from the Seeker, or have just joined the system, can advertise their presence and contribute to the resolution of queries;
– semantic "proximity": if the Provider computes some matching between a query and concepts in its own context, the system will look for addresses of other peers that the K-Owner has associated to that concept as expert in the query's topic. If one or more address are found, the Provider forwards the query to those peer, so the query will be propagated only to relevant Provider and the answers will be more accurate.

The propagation algorithm is based upon a cost function which allows choosing peers that are regarded to provide more relevant information (assigning a higher value to peer discovered through semantic criteria than to peers reached through physical proximity criteria).

Obviously, there are several parameters and mechanism controlling the scope of the search and prevent a message "flooding": set a time to live (TTL), limit the number of hop, store in the query the name of peers that already received the query, and so on.

Learning. When the matching algorithm finds a semantic correspondence between concepts of different context, the Provider can store this information for future reuse. This information is represented as a semantic "mapping" between concepts (see [6]), and can be used in three ways:

– when the K-Peer receives a query from a seeker, it can reuse the corresponding stored mapping to facilitate (or eventually don't perform) the matching algorithm;
– a Provider can use the existing mapping to forward the query to other peers that present a semantic relation with the topic of the query (see above);
– the Seeker can search into the mapping relations in order to suggest to the user a set of Provider with which it already had previous interactions and are considered qualified with respect to the semantic meaning of the concept selected in a query.

Using this mechanism, the K-Peer network will define and increase the number and quality of the semantic relations between its members, so that it becomes a dynamic web of knowledge links.

3 Development Framework

In this section we briefly show how the non-functional requirements of a DKM system drive the choice of a particular architectural pattern design (a *peer-to-peer* system) and an underlying technology framework (the *Jxta* Project).

In particular this knowledge exchange system is under development within the business setting of an Italian National Bank[4]. For more details see [**?**].

In the DKM approach, a great emphasis is given to autonomy and coordination aspects, so that every KN can access external information and can allow other KNs to access its knowledge. As we already said, this meaning negotiation is mediated through the use of contexts [6] that are a personal explicit representation, typically in form of classifi cation, of the KN's system of meanings. With respect to these requirements, a peer-to-peer system (see [**?**]) can be depicted with the following capabilities:

– being autonomous: every peer can decide how to provide a service;
– being dynamic: peers and resources can be added or removed at any time;
– being decentralized: the community of peers is able to achieve its goal independently from any specifi c member or component;
– being composed by entities that have equal capabilities: every member must provide resources or services, as well as has a right to access the others'resources and services.

These features of a peer-to-peer system seem to match the spirit and the main non-functional aspects of a KN in a DKM application, and suggest this architectural solution as a logical choice. In particular autonomy is guaranteed by the fact that each peer allows its users to manage knowledge in a personal way, using contexts that better suit their document classifi cation. Coordination is guaranteed by enabling peers to collaborate with each other, and exchange information, without imposing a common interpretation schema, but through a meaning negotiation service that automatically maps concepts among different systems of meanings.

From an implementation point of view, we focus on *JXTA* [5], a set of open, generalized peer-to-peer protocols that allow devices to communicate and collaborate through a connecting network. This P2P framework provides also a set of protocols and functionality as a decentralized discovery system, an asynchronous point-to-point messaging system, and a group membership protocol. A *peer* is a software component that runs some or all the JXTA protocols; every peer has to agree upon a common set of rules to publish, share and access "resources" (like services, data or applications), and communicate among each others. Thus, a JXTA peer is used to support higher level processes (based, for example, on organizational considerations) that are built on top of the basic peer-to-peer network infrastructure; they may include the enhancement of basic JXTA protocols (e.g. discovery) as well as user-written applications. JXTA tackles these requirements with a number of mechanisms and protocols: for instance the publishing and discovery mechanisms, together with a message-based communication infrastructure (called "pipe") and peer monitoring services, supports decentralization and dynamism. Security is supported by a membership service (which authenticates any peer applying to a peer group) and an access protocol (for authorization control). The flexibility of this framework allows to design distributed systems that cover all the requirements of a DKM application, using the JXTA P2P capabilities, completed and enhanced through

[4] This architecture is under development as part of EDAMOK, a joint project of the Institute for Scientific and Technological Research (IRST, Trento) and of the University of Trento.

[5] A P2P open source project started in 2001 and supported by Sun. `http://www.jxta.org/`

the implementation of user-defined services. As shows in the previous sections, in the Kex system we combine the P2P paradigm (characterizing a KN network as a network of distributed peers) and JXTA as an implementation infrastructure.

4 Conclusions and Research Issues

In this paper, we argued that technological architectures, when dealing with processes in which human communication is strongly involved, must be consistent with the social architecture of the process itself. In particular, in the domain of KM, technology must embody a principle of distribution that is intrinsic to the nature of organizational cognition. Here, we suggest that P2P infrastructures are especially suitable for KM applications, as they naturally implement meaning distribution and autonomy. It is perhaps worth noting at this point that other research areas are moving toward P2P architectures. In particular, we can mention the work on P2P approaches to the semantic web [**?**], to databases [**?**], to web services [**?**]. We believe this is a general trend, and that in the near future P2P infrastructure will become more and more interesting for all areas where we can't assume a centralized control.

A number of research issues need to be addressed to map aspects of distributed cognition into technological requirements. Here we propose two of them:

– **social discovery and propagation**: in order to find knowledge, people need to discover who is reachable and available to answer a request. On the one hand, broadcasting messages generates communication overflow, on the other hand talking just to physically available neighbors reduces the potential of a distributed network. A third option could be for a seeker to ask his neighbors who they trust on a topic and, among them, who is currently available. Here the question is about social mechanisms through which people find – based on trust and recommendation – other people to involve in a conversation. A similar approach could be used in order to support the propagation of information requests;
– **building communities**: if we consider communities as networks of people that, to some extent, tend to share a common perspective [2], mechanisms are needed to support the bottom-up emergence of semantic similarities across interacting KNs. Through this process, which are based on meaning negotiation protocols, people can discover and form virtual communities, and within organizations, managers might monitor the evolving trajectories of informal cognitive networks. Then, such networks, can be viewed as potential neighborhoods to support social discovery and propagation.

References

1. M. Benerecetti, P. Bouquet, and C. Ghidini. Contextual Reasoning Distilled. *Journal of Theoretical and Experimental Artificial Intelligence*, 12(3):279–305, July 2000.
2. J.R. Boland and R.V.Tenkasi. Perspective making and perspective taking in communities of knowing. *Organizational Science*, 6(4):350–372, 1995.

3. M. Bonifacio, P. Bouquet, and A. Manzardo. A distributed intelligence paradigm for knowledge management. In *AAAI Spring Symposium Series 2000 on Bringing Knowledge to Business Processes*. AAAI, 2000. Submitted.
4. M. Bonifacio, P. Bouquet, and P. Traverso. Enabling distributed knowledge management. managerial and technological implications. *Novatica and Informatik/Informatique*, III(1), 2002.
5. P. Bouquet. *Contesti e ragionamento contestuale. Il ruolo del ntesto in una teoria della rappresentazione della conoscenza*. Pantograph, Genova, Italy, 1998.
6. P. Bouquet, A. Donà, and L. Serafini. ConTeXtualizedlocal ontology specification via ctxml. In *MeaN-02 – AAAI workshop on Meaning Negotiation*, Edmonton, Alberta, Canada, 2002.
7. G. C. Bowker and S. L. Star. *Sorting things out: classification and its consequences*. MIT Press., 1999.
8. J. Dinsmore. *Partitioned Representations*. Kluwer Academic Publishers, 1991.
9. G. Fauconnier. *Mental Spaces: aspects of meaning construction in natural language*. MIT Press, 1985.
10. C. Ghidini and F. Giunchiglia. Local models semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence*, 127(2):221–259, April 2001.
11. T. Kuhn. *The structure of Scientific Revolutions*. University of Chicago Press, 1979.
12. B. Magnini, L. Serafini, and M. Speranza. Linguistic based matching of local ontologies. In *MeaN-02 – AAAI workshop on Meaning Negotiation*, Edmonton, Alberta, Canada, 2002.
13. J. McCarthy. Notes on Formalizing Context. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 555–560, Chambery, France, 1993.