KNOWLEDGE REPRESENTATION THROUGH COHERENCE SPACES

A Theoretical Framework for the Integration of Knowledge Representations

V. Michele Abrusci, Marco Romano
Dipartimento di Filosofia, Università Roma Tre, 234 via Ostiense, Roma, Italy

Christophe Fouqueré
LIPN, Université Paris 13 et CNRS, 93430 Villetaneuse, France

Keywords: Coherence spaces, Folksonomy, Knowledge Representation, Logics, Ontology.

Abstract: This work is an ongoing research (sponsored with a PhD grant by Epistematica Srl) about the interpretation of ontologies and their operations on specific graphs called *Ontological Compatibility Spaces* (OCS). Such graphs are particular Coherence Spaces that are used to give a denotational semantics to Linear Logic, hence giving a solid logical basis to our work. Using a graph framework, we depart from traditional set representation. It provides us with tools to represent actual relations among resources (and data) over the WorldWideWeb. After defining such OCSs and their correspondence with ontologies, we show to which extent folksonomies may also be of use to discover ontologies by means of OCS. Then we briefly discuss what we may obtain thanks to such an interpretation, confident that it may benefit the Semantic Web initiative with a tool for dynamic data and resources exchange.

1 INTRODUCTION

Within the Semantic Web initiative it is assumed that expressive formal description of data sources will lead to their interconnection throughout the WorldWideWeb *via* logical inter-definition of concepts appearing in different descriptions. According to this core idea, the research area of Knowledge Representation (KR) has been co-opted; the merge of the most promising KR models (namely ontologies) with standard markup languages to be used in the Web has been followed by the adoption of inference engines to be plugged into knowledge bases (KB) in order to exploit the expressivity of languages that implement some Description Logic (DL), namely OWL. Thus, the interconnection of datasources depends on their specification in a formal way suitable for automatic reasoning. In the meantime, the social evolution of the Web (Web2.0) has showed the effects of collective intelligence and its potential to manage large amounts of information with user-friendly tools requiring no specific (or none at all) skill in KR. We think that the Knowledge Engineers (KE) community that is developing the technological layer of the Semantic Web should strictly cooperate with, and take advantage of, the Web2.0 communities that spontaneously provide the Web with collections of resources that are categorized (though roughly) according to some collectively developed knowledge framework.

Therefore, after a brief discussion of the logical assumptions claimed or implied with KR involvement – and of their aptness to account for low-level, untrained and spontaneous categorization of resources – we propose an alternative logical framework, that of Linear Logic, which, we argue, can trigger the dynamic exchange of resources and data through different datasources. In particular we expect that it can provide the tools to describe and realize the passage from a datasource to any other without the need for a given-in-advance and usually "hand-made" formal description of any single datasource involved and of the mapping between every pair of them. This way, we could also get Semantic Web closer to Social Web using directly the knowledge put into Web2.0 environments (e.g. tagging spaces) without the step of ontology extraction that requires the definition of a conceptual hierarchy to be validated somehow. In fact we should build knowledge representations out of Web2.0 environments using the minimum of the appropriate logic that allows also for the representation

Abrusci V., Romano M. and Fouqueré C. (2009).
KNOWLEDGE REPRESENTATION THROUGH COHERENCE SPACES - A Theoretical Framework for the Integration of Knowledge Representations.
In *Proceedings of the International Conference on Knowledge Engineering and Ontology Development*, pages 220-225
DOI: 10.5220/0002301002200225

of operations between knowledge representations.

2 KR AND WEB RESOURCES

About ten years after the adoption of KR to support the building of the Semantic Web it is apparent that there are some major problems that make it a long and hard way to walk. Indeed, to have knowledge representations suitable for use with current technologies and according to the theoretical background of KE, we need both given in advance formal descriptions of the datasources we would like to put together over the Web, and the provision of rules for the mapping between every couple of datasources to have them really exchanging data and resources. As regards the need for formal descriptions, it requires both specialists to define the domain ontologies (Gruber, 1995) and some kind of professionals to give them the formal dressing. As regards the mapping issue, it prevents to compare any two datasources (i.e. to use them in the sense of Semantic Web) unless someone has provided the mapping instructions. Put in this way, the Semantic Web looks like too ambitious a research initiative, since it expects enormous efforts to properly define suitable ontologies, while as a commercial initiative (as it is at least partially) it follows a completely top-down and anti-dynamic approach since ontology definition invariably freezes the knowledge about something according to its understanding on the part of a restricted group of experts, thus behaving in an opposite way to the dynamic and unpredictable character of Web communities.

By contrary, within the Web2.0 movement have emerged some systems for quick and easy classification of resources by both the authors and the users that we can generally indicate with the term *tagging*, which is studied as the form of emerging categorization called *folksonomy* (Vander-Wal, 2007). Compared to ontologies, such bottom-up and popular systems are expressively very poor and usually generate no taxonomy but flat spaces where, even though there may be some hierarchy among the concepts that are identified by tag-terms, it is neither showed nor recorded when tagging. We will not hold that tagging spaces are first class knowledge representations since they completely lack any interest for the global consistency and awareness of the complete resulting terminology or concept-scheme. Nevertheless tagging is the form of "naive categorization", made by sticking labels to resources in the Web, where folksonomy comes from. Folksonomies are the aggregations of tag-terms, i.e. what can be used to find out the resources that match some tag and, with it, also the corresponding concept expressed by the term, precisely as it should be for Semantic Web ontologies. Secondly, flat tag spaces are classifying – yet poorly, but they are really doing it day by day – greater and greater amounts of web resources. So they seems able to do for Semantic Web that part of the work that ontologies cannot manage, i.e. large scale categorization. If we then consider that KR has entered the Web in order to help in making machine-readable the information in it, and not primarily to develop the finest descriptions of particular worlds, we should exploit any contribution able to cooperate in achieving Semantic Web.

2.1 A Logical Model to Rely On

Being a major result of KE, ontologies and especially DL ontologies have a precise logical meaning, i.e. a definite semantic interpretation (Baader et al., 2003).

Let's observe what is their semantic interpretation. It relies on Set theory and associates to the ideas expressed by concept names (and their logical definitions) the corresponding set of "all the objects that ..." within the specific domain of interest that the ontology is designed to describe.

More formally, let I be the interpretation function from an ontology to a non-empty set Δ, which is the domain of interpretation. An atomic concept A is interpreted as $A^I \subseteq \Delta$ while a role R as $R^I \subseteq \Delta \times \Delta$. Let C, D be concepts; R a role between C and D; a, b individuals; the semantics of an ontology results as: $I(\top) = \Delta$; $I(\bot) = \emptyset$; $I(C) = C^I (\subseteq \Delta)$; $I(D \sqsubseteq C) = D^I \subseteq C^I$; $I(C(a)) = a \in C^I$; $I(R) = R^I \subseteq \{\langle x, y\rangle \,|x \in C^I \wedge y \in D^I\}$; $I(R(a,b)) = \langle a, b\rangle \in R^I$. The inferences that a reasoner can draw from such KBs are generated by the axioms offered as concept definitions and possibly also presented as A-box assertions about individuals. Legal inferences are those which are valid for all the possible models based on the particular domain Δ and function I chosen. By the way we remark that generally the interpretation domain Δ is never given any concreteness, so that operations between ontologies affect primarily the intensional level and then the semantic interpretation is produced simply as a by-product of formalization, since it is always possible to declare which set should be the result of some operation. This would not be a problem if ontologies were not to assist interchange of data in a Web of concrete resources. But Semantic Web ontologies are precisely to describe what is within different datasources and to enable the jump from each other so that what exactly is Δ is not a minor problem. In order to achieve a working Semantic Web, we think we should talk about *applied ontologies*, i.e.

ontologies together with the collection of data or resources that are accounted for within the ontologies, so that also the operations between ontologies should be considered primarily as affecting the resources.

Now let's observe the semantic interpretation of folksonomies. To be honest, they have no definite semantic interpretation, but it seems quite easy to adopt once again Set theory and consider them as very poor ontologies without concept definitions. It would be straightforward to take some domain Δ and interpret all the resources that share a specific tag as a particular subset of Δ identified with the concept lying behind the tag-term. As regards the calculus they support, no inference is allowed, but mere resource retrieval by using tag-terms as search keys. Operations between folksonomies have not even been conceived.

2.2 One Step Further with More Structure

We propose a theoretical framework for KR specially conceived for use within the Semantic Web scenario. Behind this proposal there is a doubt about the adequacy of the current approach based on Set theory (for KR) and linguistic analysis of KBs' intensional level (for the discovery of compatible resources) with respect to the full achievement of the Semantic Web. Thus, instead of focusing on conceptual schemes of ontologies, we propose to focus on the extensional level of KBs, i.e. on "real" objects, by adopting a logical framework capable to geometrically represent relations among resources, possibly discovering the concepts from resource aggregations that are actually in the Web. In particular we suggest to consider another kind of semantic interpretation that relies on structures richer than sets, the *coherence spaces*, where the interpretation of a concept (or tag-term) produces graph theoretical objects along with the determination of the extensional counterpart within the collection of resources that is the domain of interpretation.

Coherence spaces (Girard et al., 1989) are webs whose points may or not be linked to each other according to a binary reflexive relation called *coherence*. They allow for the definition of a denotational semantic, so that we can get one also for data exchange within the WorldWideWeb and for a definition of operations between ontologies that is primarily focused on resources. Coherence spaces come from Linear Logic (LL) and have been the first semantic interpretation of that logical system. What is more is that it is not truth-valued semantics, useful only for talking about formulas – as it always happens with Classical Logic (LK) – but precisely denotational semantics, useful for talking also about proofs. Indeed, they offer the domain of interpretation of the objects manipulated by the logical calculus, i.e. proofs.

LL (Girard, 1987) is a logical system developed by imposing some restrictions on the use of structural rules for the construction of deductions within Gentzen's proof calculus for first order logic (FOL) known as sequent calculus. The affected rules are Contraction and Weakening which deal with the number of times formulas may be used within the same proof. In LL they are re-defined in form of logical rules (instead of structural) so that their usage has to be marked with specific connectives, called *exponentials*. This way everything that holds in LK also holds in LL, although LL is able to better describe what is happening in a proof: the ability to mark for which formulas it is licit to have weakening and contraction means that one can control the times resources are used. We note, by the way, that this one looks like an interesting property to have at hand while working for Semantic Web. As a consequence of the control on contraction and weakening, LL deconstructs the connectives $\wedge$ and $\vee$ doubling them in the multiplicative and additive variants, since their behaviour is different according to the possible uses of the context (i.e. the other formulas) where the formulas in which they occur are interacting with. To put it in a nutshell, the multiplicative connectives operate on the coherence space resulting as the product of the coherence spaces corresponding to the proofs of the connected formulas, while the additives on their disjoint union. Moreover, LL has developed a geometrical representation of proofs by means of graphs called proof-nets. It exploits graph structures to compose partial proofs and provides graph properties to determine when a proof structure is correct. Such graph structures have a model in coherence spaces, where the denotation of the proof of a formula is a set of pairwise coherent points, called *clique*. The operations between coherence spaces interpret the composition of formulas and their proofs according to LL connectives. In order to account for the interpretation of ontologies through coherence spaces we need a particular class of them, with a typical support set (see below). In addition, because of the context of Semantic Web and its need for datasource integration, we propose to call the structuring relation of such coherence spaces compatibility rather than coherence, and to call this special class of coherence spaces "Ontological Compatibility Spaces" (OCS).

3 ONTOLOGICAL COMPATIBILITY SPACES

Let A be a semantic web ontology in a language like OWL. Let $\mathcal{L}(A)$ be the set of all the symbols for individuals, concepts and roles of A (i.e. individual symbols, unary predicate symbols, binary predicate symbols), M a set of data, ϕ a valuation through suitable mapping from $\mathcal{L}(A)$ to M, i.e. if c is an individual symbol of $\mathcal{L}(A)$ then $\phi(c)$ is an element of M, if P is a unary predicate symbol of $\mathcal{L}(A)$ then $\phi(P) \subseteq M$, if P is a binary predicate symbol of $\mathcal{L}(A)$ then $\phi(P) \subseteq M \times M$.

This defines an applied ontology A_M and we can represent every concept, role and individual of A_M by means of a coherence space $[A,M,\phi]$ with support $|[A,M,\phi]| = M \cup (M \times M)$ provided with a coherence relation noted $\frown_{[A,M,\phi]}$ between its points: $x \frown_{[A,M,\phi]} y$ iff $\exists P \in \mathcal{L}(A)$ such that $\{x,y\} \subseteq \phi(P)$. For sake of clarity, we may note simply $\frown$ when the coherence space is obvious. This coherence relation formalizes the notion of compatibility emerging whenever an ontology is applied to a set of data. This way we have defined an OCS.

Indeed, from an abstract point of view, the retrieved values for any predicate symbol P of A form some subset of $M \cup (M \times M)$ whose elements share with each other something more than all the other elements of $M \cup (M \times M)$. Such a property, instead of being named according to any specific symbol P occurring in the ontology, may be rewarded as the compatibility between all the points of that subset of $|[A,M,\phi]|$. Such a group a of pairwise compatible points of $[A,M,\phi]$ is called a *clique* and is noted $a \sqsubset [A,M,\phi]$.

Defined as a coherence relation, the compatibility relation is reflexive, symmetric and not transitive. Reflexivity ($\forall x \in |[A,M,\phi]|, x \frown x$) assures that every point of the coherence space is compatible with itself. Since we deal with a coherence space whose web is provided by the set $M \cup (M \times M)$ we have also pairs $\langle y,z \rangle$ as points of $|[A,M,\phi]|$. Thanks to reflexivity, every point is a clique and may be considered as a minimal class of compatibility. As an example, singletons that are retrieved from M through ϕ (for individuals of A) are interpreted in such cliques. Symmetry ($\forall x,y \in |[A,M,\phi]|, x \frown y \iff y \frown x$) expresses the core of the idea of compatibility, since for compatibility we mean the possibility to put two "objects" together based on their sharing of some common property – not necessarily an expressed one – and such a commonality has to be inevitably a reciprocal fact. Non-transitivity prevents the overwhelming distribution of compatibility that would mix up different Ps of A whenever they have some common points. However transitivity may appear under certain conditions, e.g. within a clique, where the points of a clique are all pairwise compatible.

A coherence space can be represented by a graph whose set of nodes is given with the web and the edges reflect the relations of coherence of every point with the others. Also an OCS can be represented through a graph. The graph $G(V,E)$ of our OCS $[A,M,\phi]$ may be defined by the set of nodes $V = |[A,M,\phi]|$ and that of edges $E \subseteq |[A,M,\phi]|^2$ with the constraint that for every two points x,y of $|[A,M,\phi]|$ we have $\{x,y\} \in E \iff x \frown y$ that is to say $\{x,y\} \in E \iff \exists P \in \mathcal{L}(A)$ s.t. $\{x,y\} \subseteq \phi(P)$. Based on the definition of coherence, when looking at the graph of an OCS, a clique $a \sqsubset [A,M,\phi]$ turns out to be a completely connected portion of the graph (a subgraph).

We observe that the interpretation of ontologies as OCSs successes: i) every individual of A_M is represented within the OCS as a clique of a single point (the smallest class of compatibility); ii) every concept or role of A_M is represented as a class of compatible points (a general clique). Following the inverse direction, we observe that every clique of the OCS is the denotation of: i) some concept or role of A_M; ii) or some subconcept (or subrole) not identified by any predicate of $\mathcal{L}(A)$ yet recognizable in A_M as a subset of some concept (or role) identified by some predicate in $\mathcal{L}(A)$; iii) or an individual of A_M.

4 OCS VS ONTOLOGY

We have designed OCSs specially for representing ontologies. We now discuss the advantages that our proposal may bring to the development and usage of folksonomies and suggest to which extent all this may benefit Semantic Web.

4.1 Tagging and Ontology Extraction

Nowadays when looking at the set of tag-terms adopted within a community it is expected to reconstruct a formal ontology out of that, establishing a neat and formal hierarchy among concepts, useful for resource retrieval through progressive specification. The major side effect of such a reduction is the loss of the dynamic aspect of Web2.0.

We may observe that building the Semantic Web by means of ontologies requires predefined sets of metadata (the schemes) to be adopted and respectfully obeyed. Their usefulness – and the wealth of Semantic Web itself as the workplace of autonomous agents (Berners-Lee et al., 2001) – will depend in fact on the number of resources whose set of metadata matches one or more of the predefined schemes so that programs specially written according to the same scheme(s) will be able to use those resources. So the Web of data that W3C indicates turns out to be something like a huge database where a neat definition of the logical scheme (even composed of many different ontologies) can be achieved only thanks to the standardization of the metadata tags to be used to describe resources. In the opposite direction goes the practice of free tagging, so that folksonomies emerge as everlasting works in progress where concept-terms institution and resource description and classification always happen in the same time, with no hope for standardization. In fact, when people tag they freely choose and establish their own categories in an unending process of ontology elaboration. Moreover, while using their very personal categories people also express their own "world's understanding" so that tagging spaces are not only useful for classification, but also convenient for collective intelligence to share knowledge. We remark that tagging spaces publish enormous amounts of resources with some kind of classification while providing a cognitive framework that has not the claims of ontology but that is powerful enough to let one recognize and find classes of resources that are compatible, i.e. similar to some extent. Maybe such a cognitive framework is a lower quality contribution, in comparison to formal DL ontologies, to have the content of the Web surely recognizable, but it seems to show a more feasible way to do that. In fact, since it relies on no pre-emptive requirements, no standardized label to tag resources, it preserves the dynamical behaviour of Web2.0 and lets Semantic Web to be a "common people affair" as it has been for last years for WorldWideWeb, and for its big boom. However, it still lacks an appropriate theoretical framework to be successfully exploited to the benefit of Semantic Web. Therefore we propose one: to consider tagging spaces without any attempt to reduce them to formal ontologies. Instead of the usual techniques for tags clustering and concept extraction, we may exploit OCSs in order to recover from tagging spaces a description of the resources in a datasource that is formal enough to be useful for data exchange but that does not need *ad hoc* specification of a conceptual hierarchy. We look only for compatibility between resources observing the connections given within an OCS. Since operations between coherence spaces have already been largely studied in LL, we have almost ready-to-use precision tools to talk about usual operations (union of ontologies, module extraction, ... up to ontology merging) *via* their interpretation as operations between OCSs using the LL connectives – and thus opening to a new role to play in Computer Science for LL (Ehrhard et al., 2004). Our proposal is then to give more logical dignity to flat tagging spaces without rising too high in formal complexity so as to prevent large contribution from common web-users. We simply aim to describe flat spaces in such a way that it makes sense to talk about operations between them.

4.2 Folksonomies out of OCS

In order to have an OCS out of a flat tagging space we need nothing more than what has been stated for ontologies, since we will consider it as a very simple ontology, with just some niceties: the language $\mathcal{L}(F)$ of such a folksonomy F will be the pair (T,I) where T is the set of tag-terms and I the set of URIs. The web of the OCS $[F,I,\phi]$ will be $|[F,I,\phi]| = I$. ϕ is the pointer from a tag-term $\mathbf{t}$ to the set $\phi(\mathbf{t}) \subseteq I$ of the resources tagged with $\mathbf{t}$ – i.e. it is a query. The compatibility relation is slightly revised as $x \frown_{[F,I,\phi]} y \iff \exists \mathbf{t} \in \mathcal{L}(F)$ s.t. $\{x,y\} \subseteq \phi(\mathbf{t})$. We observe some characteristics of such an OCS: i) for every tag-term $\mathbf{t}$: $\mathbf{t} \in T \rightarrow \exists x \in I\ x \in \phi(\mathbf{t})$; ii) if we assume that in a tagging space there cannot be a resource that has no tag, we have also the inverse, for every URI i: $i \in I \rightarrow \exists \mathbf{v} \in T\ i \in \phi(\mathbf{v})$. These mean that we have no empty concepts and that the graph associated to a folksonomy may be a fully labelled graph. Moreover, we remark that thanks to the compatibility relation we are now able to read and build concepts out of the tagging space in an original way. In fact we do not search for interesting concepts looking at recurring couplings of certain tag-terms stuck to different resources (looking for synonymy or some subsumption between terms). Quite the contrary, while searching for cliques we look at recurring couplings of certain resources under different tags, i.e. among the URIs retrieved for the tags. We look for compatibility between resources and try to collect all the classes of compatibility, what turns out to be a new tool for concept discovery. Based on the definition of compatibility, we may have a class of compatible resources which are not all together within the retrieved URIs for one single tag, yet they are all pairwise compatible. It is such a case that of a possible new concept not explicitly recognized on the part of the tagging community but implicitly present as an underlying idea of compatibility. Whether such a concept can be identified with some linguistic term or not, it is not important: we are abstracting from linguistic determination of compatibility classes and we can grasp some new concept which one can look for other elements compatible with.

5 CONCLUSIONS

In conclusion, we recall some of the most challenging aspects that are to be further investigated in our research along with the results that we can see at present. Recurring to coherence spaces to describe operations between ontologies (and folksonomies), we have to face some deep questions that emerge from aspects inner to LL and coherence spaces theory. Indeed, coherence spaces and their cliques represent proofs of (multisets of) formulas (from here on we use formula while meaning multiset). Each coherence space offers the place for the denotation of one formula and shows a correspondence regarding a clique, a formula and its proof, while operations between coherence spaces provide the denotation for the calculus when composing formulas. When adopting coherence spaces for ontologies, we have to find the right place to put other elements in this correspondence, i.e. concepts and resources, together with ontology. We have to clearly set the circle of correspondences and a first attempt is to compare the formulas appearing in the right-hand side of a sequent in LL with an ontology seen as the union of its concepts, so that the proof satisfying one of the concepts satisfies also the whole ontology.

For the time being, we are proposing something like an alternative descriptive language to represent data sources in the WorldWideWeb. With such a language we are given the easy of use of tagging space, the ontology/folksonomy extraction procedure described above and a calculus based on LL connectives, that are a little more and a little finer than those usually adopted working with ontologies. We have already mentioned the doubling of some connectives ($\wedge$ and $\vee$, in their multiplicative and additive formulation) in LL, together with the appearing of two new connectives, the exponentials, that control the use of contraction and weakening, i.e. the use of resources. Without doubt it is worth to assess the usefulness of such "logical controllers" in order to account for operations in the Web that involve resources. If we consider that the multiplicative fragment of LL shows interesting computational properties and that reasoners for LL are already available, we are further encouraged in developing such a language in order to assess its capability to satisfy Semantic Web needs.

REFERENCES

Baader, F., Calvanese, D., McGuinness, D., Nardi, D., and Patel-Schneider, P., editors (2003). *The Description Logics Handbook: Theory, Implementation and Applications*. Cambridge University Press.

Berners-Lee, T., Hendler, J., and Lassila, O. (2001). The semantic web. *Scientific American*.

Ehrhard, T., Girard, J.-Y., Ruet, P., and Scot, P., editors (2004). *Linear Logic in Computer Science*, volume 316 of *London Mathematical Society Lecture Note*. Cambridge University Press.

Girard, J.-Y. (1987). Linear logic. *Theoretical Computer Science*.

Girard, J.-Y., Lafont, Y., and Taylor, P. (1989). *Proofs and Types*. Cambridge University Press.

Gruber, T. (1995). Towards principles for the design of ontologies used for knowledge sharing. *Int'l Journal of Human and Computer Studies*.

Vander-Wal, T. (2007). Folksonomy. http://vanderwal.net/folksonomy.html.