



Speech acts meet tagging

Alexandre Monnin, Freddy Limpens, Fabien Gandon, David Laniado

► To cite this version:

Alexandre Monnin, Freddy Limpens, Fabien Gandon, David Laniado. Speech acts meet tagging. 5th AIS SigPrag International Pragmatic Web Conference Track (ICPW 2010), Conference Track of I-Semantics, 2010, Sep 2010, Graz, Austria. pp.1-10, 10.1145/1839707.1839746 . hal-01170885

HAL Id: hal-01170885

<https://inria.hal.science/hal-01170885>

Submitted on 2 Jul 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Speech acts meet tagging: NiceTag ontology

Alexandre Monnin
Philosophies contemporaines, ExeCO,
Univ. Paris I Panthéon-Sorbonne
17, rue de la Sorbonne
75005, Paris

Freddy Limpens,
Fabien Gandon
Edelweiss, INRIA Sophia-Antipolis
2004, route des Lucioles BP 93
06902 Sophia Antipolis Cedex

David Laniado
DEI, Politecnico di Milano
Piazza L. da Vinci, 32
20133 Milano

Alexandre.Monnin@malix.univ-paris1.fr

Freddy.Limpens@sophia.inria.fr
Fabien.Gandon@sophia.inria.fr

david.laniado@elet.polimi.it

ABSTRACT

Current tag models do not fully take into account the rich and diverse nature of tags. Each model makes different partial assumptions as to the definition and attributes a tag should receive. In this paper we propose an ontology, NiceTag, whose primitives are “tag actions” modeled with RDF named graphs. This mechanism allows us to type, describe and thus ensure the traceability of each single act of tagging. Our named graphs contain at least a resource linked to a “sign”, which can be any resource reachable on the Web (an ontology concept, an image, etc.). The *resource*, the *sign* and the *link* between them are the three components of the acts of tagging that we want to explicitly represent as social actions, akin to speech acts. The purpose of our model is threefold. First, to be able to describe acts of tagging in a very precise and general manner, consistent with the principles behind the architecture of the Web. To reconcile and bridge existing tag models (Newman ontology, Tagont, ES, SCOT, SIOC, CommonTag, MOAT, NAO). And finally, to propose a viable way to reify and represent the intention behind an act of tagging and leverage its semantics.

Categories and Subject Descriptors

H.5.3 [Information systems]: Group and organization interfaces – *Collaborative computing*. H.1.1. [Information systems]: Systems and information theory – *Information theory*. I.2.4 [Computing methodologies]: Knowledge representation Formalisms and Methods.

General Terms

Documentation, Human Factors, Standardization, Languages.

Keywords

Tagging, resources, social acts, speech acts, meaning, sense-making, pragmatic, identity crisis, annotation, named graphs.

1. INTRODUCTION

Tags in current tag ontologies are almost always modeled so as to associate a “tag”, a “user” and a “resource”. Unfortunately these primitives have barely been theorized in academic discussions

related to tagging. Though seemingly unchallenged and despite it being implemented in many systems, a close look at each component of the aforementioned tripartite tag/user/resource model reveals a lot of unnoticed difficulties thus calling for dire improvements.

For that reason, our main purpose will be to provide a model that answers one – apparently - simple question “What exactly is being tagged?”. On the Web, according to the guiding principles behind its architecture, the answer must always be: *a resource*. To get a clear idea of this rich notion it is necessary to go back to the specifications and debates where it was fleshed out, especially those debates that surrounded the *Identity Crisis*¹ of the Semantic Web as their outcome was a precise characterization of this notion.

In what follows we rely heavily on these works to answer our own question. Hence, a lot of space will be devoted to the task of underlining the relevance of the lessons learned from the Identity Crisis resolution to tagging. By doing so, we also voluntarily anchor tagging in the specific environment where it thrived: the Web. There is no doubt tagging is but one form of annotation among many. Yet, the need to be more specific is felt since the answer to our first question is determined by the requirements of a very specific architecture that will henceforth constitute our operational framework.

There are many graph-based knowledge representation formalisms from Conceptual graphs, which are historical descendants of semantic networks, to Entity-Relationship models, UML, Topic Maps or RDF that share the threefold structure of the tag/user/resource model. This similarity has led to many attempts to bridge the gap between social tagging and Semantic Web technologies. An effort we clearly take up by relying on named graphs, an extension of RDF. Yet, the contribution of this article is not to propose yet another graph-oriented formalism. Rather our intention is to devise a conceptual framework able to represent tags or rather *actions of tagging*. This is achieved by relying on an existing formalism (RDF) plus an extension (named graph model and syntax) and a schema (an OWL upper ontology) in order to make it possible to capture and type tags, their structure and uses.

Replacing the primitive of our ontology with “*tag actions*” instead of tags ensures that our model is a meta-model of tagging, able to bridge the gap between existing ontologies. More importantly maybe from a theoretic point of view, actions of tagging are

¹ This expression refers to the difficulties met by the Semantic Web community when the transition between a Web of “documents” and a Web of “non-information resources” (or things in the world), as advocated through the Linked Data initiative, became a pressing and concrete issue. See [1].

hereafter acknowledged and modeled as *social acts* or *speech acts*. Not unlike language, tagging should not be restricted to descriptive assertions if we are to account for existing usages in the hope to better capture their semantics.

This paper is organized as follows. In section two we propose a variety of scenarios to analyze the entire process of tagging and its many facets in an attempt to answer the question “what do we tag?” In the third section, we introduce the main feature of the NiceTag ontology, the use of named graphs to model tag actions, before we give, in section four, more details on the relations between a tag and a tagged resource. Section five discusses the current implementation of tagging in existing platforms such as delicious.com and the identity of tags. Section six works as a practical guide to NiceTag by presenting a variety of potential and existing uses-cases. Section seven concludes this article.

2. THE ANATOMY OF TAGGING

2.1 In the beginning were resources:

The advent of the Linked Open Data (LOD) paradigm was instrumental in revealing the intricacies behind Web addresses. For those who have no idea of how the Web is built, only URLs, links to documents conceived as HTML files, exist. This is roughly the layman’s point of view. However, the Web’s architecture wasn’t conceived that way. URIs do not only *give access* to Web pages, they primarily *identify* resources, information objects - that which a Web page represents.

Resources, as defined by the now outdated but still relevant for that matter RFC 2396, can be just about anything as long as they are identified by URIs. A resource is thus created by the person who publishes its corresponding URI. Each URI identifies no more than one resource. This is the *dictum* that best summarizes the architecture of the Web². Such a relation of identification is functional in the mathematical sense of the word. These resources are the responsibility of the person who published the URI that identifies them. Hence, according to many informal documents, some of them authored by Tim Berners-Lee himself, the authority of the publisher is what really determines the reference/meaning of a URI (i.e.: the resource).

Both relations of *access* and *identification* are governed by specifications and protocols alike. By contrast, the linguistic relation of *reference*, common in the philosophy of language, breaks away from the technical shackles imposed by technology³. For instance, nothing prevents two persons from referring to different things with a single URI. This conceptual framework, characteristic of the LOD initiative, opened up the possibility to handle things – resources – that are not Web accessible – not just “documents” or information objects. The question “How then are we to distinguish URIs from Linked Data URIs?” ensued as a consequence and led to what became known as the *Identity Crisis of the Semantic Web*.

While the Technical Architecture Group’s (TAG) own solution became known as the HTTP range14 decision, we will here focus

on the work of clarification accomplished by H.Halpin and V.Presutti [2]⁴ for it offers decisive insights which may eventually lead to an answer to our own question: “What exactly is being tagged?”.

Halpin and Presutti distinguish two main types of resources according to their various states:

- a) Information resources
 - a. Web resources
 - b. Web representations
- b) Non-information resources

Tags, contrary to URIs, are not necessarily what we call “*Web signs*”, signs whose functioning is largely determined by technical rules and specifications. That’s why they can easily violate one of the main axioms of the Web: that a URI should never be allowed to identify more than one resource.

The publisher of a URI, being granted with the power to state which resource it identifies, undoubtedly holds a privileged position according to the W3C specifications. This is precisely why these identifiers have been dubbed “Web proper names” [3]. Some philosophical accounts of logical proper names similarly insist on the original dubbing ceremony through which a proper name is directly assigned a reference, which is henceforth transmitted to other speakers of a language through causal chains [4]⁵.

“Direct reference” is a position shared by almost all the participants in the debate surrounding the Identity Crises. Yet, it fostered at least two conflicting positions. The first, that insists on the functional relation implied with direct reference and the authority of the publisher, is consistent with the idea that both a URI and a logical proper name (the name of a person, a common noun of kind like “water”, etc.), once the dubbing ceremony is accomplished, can refer to only one thing. This aspect was formalized in the IRW ontology proposed by H.Halpin and V.Presutti [2] with the property `irw:identifies`.

The other position was advocated by Patrick Hayes. It contrasts with the great emphasis Semantic Web specifications put on the owner of a URI considered as the sole person habilitated to specify its reference/meaning. Separating issues pertaining to language from technical ones, it holds that proper names can *refer* to whatever people like (a la Humpty Dumpty). Hence, if URIs are real proper names, linguistic signs (i.e. not Web signs⁶), they should work similarly. This was captured with the property `irw:refersTo`.

Both positions assume that reference is best stated in private ceremonies, either by one specific person or by anyone. The

² “The relations between a URI and resource (...) is functional as the W3C states one should “assign distinct URIs to distinct resources [2]”.

³ A point that was made by P. Hayes: “the architecture of the Web determines access, but has no direct influence on reference” [1].

⁴ This was anticipated by previous works by Presutti and Gangemi drawing on the DOLCE+DnS Ultra Lite ontology, see [5], [6], [7] and [8]. The OKKAM project deals with the same issues, especially [10]. The W3C keeps working on this question inside the “Architecture of the World Wide Semantic Web” Task Force (AWWSW), <http://esw.w3.org/AwwswHome>.

⁵ See also [10] for an account of direct reference theories.

⁶ Judging from his keynote speech at ISWC 2009 entitled “Blogic”, such a position may longer be the one advocated by Hayes.

actual situation involved in tagging may well have more in common with the second account. After all, the person who is tagging is not bound by all the rules of the semantic Web since, by all odds, she is not the publisher of the URI. Hence, she will be essentially concerned with the URI understood as a linguistic sign or as an address since on the Web one is granted access to the *representation of the resource* and from this representation may *infer* what the resource is *from her own point of view* (with results that may well vary from person to person, the same URI will hence refer to a picture of a landscape, that landscape, an image, a digital object, the place where one was born, a picture of the place where one was born, etc. – a feature that goes against functional identification). Furthermore, the act of tagging itself will involve all kinds of signs, not just directly referring proper names.

As a consequence, the potential shift of referent from one act of tagging to another one, when both are prompted by the same accessed `irw:WebResource`, is a great cause of ambiguity, very much reminiscent of the symptoms surrounding the identity crisis. IRW ontology allows for a fine-grained rendering of these subtleties by accounting very precisely for the possibility that users would refer to something else than what a URI either identifies or grants access to (a possibility integral to tagging). That is why we make use of an equivalent of the class `irw:Resource` to model the tagged resource (`:TaggedResource`) and offer a different class to model annotated resources. The specific trait of `nt:annotatedResource(s)` is that they are *browsable resources*, in other words, *web-accessible representations of resources* (an equivalent of `irw:WebRepresentation`).


2.2 Tags and resources:

2.2.1 Annotation:

One may ask whether the very notion of tags is misleading. After all, annotation precedes tagging by far. In a sense, it might be seen as a way to simply annotate shared resources on the Web. There's no doubt annotation is a defining feature of tagging. Muxway, the first tagging Web service, resulted from an attempt to add annotations to personal bookmarks online. "Tags", as J. Schachter dubbed these annotations, were only "blank spaces" where various labels could be inscribed and attached to a bookmarked URL.


The architecture of the Web largely determines what can be tagged in this context: instead of either objects or documents, on the Web one is dealing (at least under current constraints) with *resources*. Various characterizations of resources have been given since the beginnings of the Web. The first HTTP RFC stated that resources were "a network data object or service identified by a URI".

Just like material tags are attached to material objects, so are annotations attached to resources on an information network. In both cases, however, the inscription (what we call the "label" of a tag or, more generally, of an annotation) while attached to an object on the network, may apply to just about *anything*: the constraints of meaning aren't those of annotation – only resources *accessible* on the Web via the HTTP protocol, browsable resources in other words, are available for annotation. Consequently, all that is required for annotation on the Web is to attach "tags" to an HTTP resource (one that is identified by a single URI and accessed through the set of its representations). Tags, on the other hand, need not refer or apply to resources identified by URIs nor take their meaning solely from them.



2.2.2 Lexical meaning and direct reference:

But this is only one aspect of tagging. A tag may be attached to a resource without either the least *referring* to this resource. If one decides to tag an image of one's parents with the label "parents", the word "parents" has some specific lexical meaning. By contrast, "my parents" should refer to a pair of human beings while proper names like "John" and "Johann" may each have a FOAF profile, identified by a Linked Data URI that could be used to specifically identify and designate both of them.



Compared to our previous schema, this relation starts from the label. The reason here is quite simple: if the label happens to be a word, it may very well already have a meaning of its own, or if it is a proper name for example, it may refer to some precise entity.

Existing tag ontologies have almost exclusively focused their efforts on harnessing the meaning of tags - essentially conceived as signs. This explains why issues of synonymy, spelling variations and plain misspellings are regularly brought back against tagging.

A seemingly indisputable proof of tagging's serious lack of accuracy is that they were among the first issues to be tackled in Newman's ontology⁷ and SCOT⁸. Likewise, by allowing anyone to link a tag to a well defined meaning, as do MOAT or Common Tag, a relationship is created which helps to face the problem of the different acceptations a term can receive in different contexts and for different communities [11].

Still, where a tool like MOAT best achieves its purpose is when it is applied to disambiguate the use of a directly referring proper name, be it that of a celebrity "Paris" (Hilton), or of a town "Paris" (France) (as bespeaks the fact that the examples given by Passant [3] are almost all of that kind).

However, an application that facilitates the identification of distinct referents of two otherwise identical proper names (homonymy due to homographs) does not provide a solution to the problem of polysemy (the multiple meanings of a sign). This is evidenced by the choice of Wikipedia, an encyclopedia, as the knowledge base from which the "meanings" attached to labels in MOAT are extracted in [3]. Contrary to dictionaries or thesauri, encyclopedias would simply be unsuited to this task, being catalogs of entities rather than words senses⁹.


2.2.3 Ostensive definitions and demonstratives:

This is especially useful with ostensive definitions and demonstratives. On a photo-sharing service for instance, someone may choose to refer directly to the person depicted on a picture rather than to the picture itself.

⁷ <http://www.holygoat.co.uk/projects/tags/>

⁸ <http://scot-project.org/>

⁹ Of course, one possible answer would be to treat words senses as things and to refer to these entities with proper names (a serious stretch to common sense).



From the TAG's (Technical Architecture Group) definition (anything which might be identified by a URI) to Tim Berners-Lee's, a slight shift can be observed. For Berners-Lee, indeed, a resource is anything that has an identity. Conversely, for anything to have identity, to exist in a tractable fashion on the Web means that it is *ipso facto* a resource identified by a URI. Hence, W. Quine's famous *dictum* "no entity without identity", is still true on the Web provided it is clear that, as a technical system and mediator, the Web only knows resources as entities. To put it another way, *the ontology of the Web is an ontology of resources; objects once mediated by the Web are qualified as resources*.

People (as modeled by the FOAF vocabulary), countries (as modeled by geonames), genes (by the gene ontology), species (by geospecies), concepts and abstract entities, in other words, Web-inaccessible things in general, when subjected to identification on the Web according to the very principles behind the Linked data initiative all fall under the *non-information resources* category¹⁰. Therefore, on the Web, the space of qualities and the space of objects is altogether replaced by the hierarchy of resources described in the IRW ontology.

Thus, when tags are used in ostensive definitions or as demonstratives, the resources they're referring to are no longer restricted to Web resources and may be outside the scope of *annotation* as we've defined it. Yet tags are not sufficient to functionally *identify* these resources since the link between a tag and a resource is much less formal than the link between a URI and a resource. Only URIs will make it possible to state the identity of resources on the Web. This is precisely what an ontology like MOAT does. If the label of a tag is used to refer to a non-information resource already identified by a linked data URI, a relation is available in MOAT which will let users link that URI to the label of the said tag.


2.2.4 Relations and reference shifting:

Of course, what's equally important is the way labels are related to resources and how we are manipulating them, whence the paramount importance of the meaning of relations established between the resource and the tag. Still, no known effort was undertaken to type the relations between those signs and resources. This task is left entirely up to triple and/or machine tags were predicates are stated, though it is not accomplished according to any Semantic Web standards (explaining why they were dubbed "the poor man's RDF" from the inception). The situation is officially acknowledged as the main source of ambiguity of tagging by the W3C itself¹¹.

¹⁰ This is a moot point in the case of books, either treated as texts, information resource, or as physical objects, - their material realizations. This example (in particular *Moby Dick*) played a paradigmatic role in discussions belonging to the Identity Crisis. It is also central in debates at the crossroads of philosophy, ontology and linguistic on systematic polysemy.


¹¹ See: <http://weblog.scifihi.com/2005/08/05/meta-tags-the-poor-mans-rdf/>

As a matter of fact, homonymy, polysemy, synonymy, variations of spelling and utter misspellings are not the only sources of ambiguity encountered with tagging. A user may either want to say, somehow lacking precision as befits human communication (making it possible), that this *is* T. B.-L., or that this picture *represents* T. B.-L. Similarly, other tags may either describe the qualities of the man or the picture.



2.2.5 Same resource, different relation and different assertions:

Other uses-cases, outside of reference-shifting, call for an explicit integration of relations in tagging ontologies. Far from being benign, the simple fact that a document *is about* a thesis or is *relevant for* a thesis, and that we are left with the choice to assert the first relation, the second or both, is of utmost importance. Current bookmarking systems do not allow for a sign to be used more than once to *annotate* a given Web resource. Such limitations stem from leaving implicit the relation between labels and resources and consequently assuming that it is always a descriptive one. Tags, accordingly, are often believed to be tantamount to "topics". Yet, should relations in their diversity be made explicit, nothing would no longer prevent considering the following example as valid:



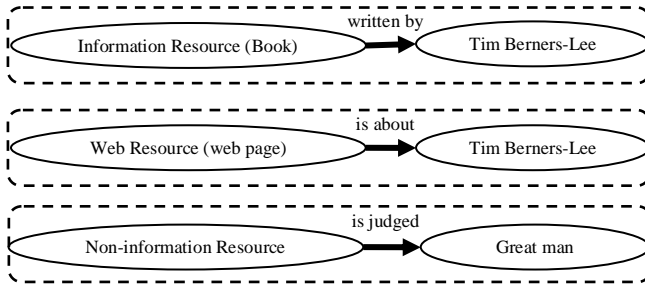
Stating these relations also presents one more advantage since some of them, to hold meaningfully, necessarily *coerce* their *relata* (this is accomplished in RDF through `rdfs:range` and `domain`). In other words, when something is declared to be "written by" someone, this statement embodies information about both sides of the relation (namely, *that which can be written* and *that whom can write*). Hence, it is easy to see that *tagged resources* vary, as well as the meaning of tags, according to the relation linking both.

Of course, what *triggers* the act of tagging (the *annotated resource*) still remains a browsable Web resource, identified by a dereferenceable URI. Yet, what users are *referring to* depends on how they envision the nature of the relations between tags and their resources. The complexity of resources and the ambiguity it creates lie in their being conceivable from different angles. The inherent "aboutness" they share with documents ensures that either themselves or what they represent can be tagged.

2.2.6 Different resources, different relations, different assertions

Sometimes, depending on the chosen relations, it is possible to distinguish between the two aspects. Let us take a bookseller's webpage about a book, *Weaving the Web* by Tim Berners-Lee for example. It could equally be tagged with a relation expressing what the *webpage* is about (and Berners-Lee certainly qualifies as

a good candidate in this context) or a relation expressing the fact that the *book Weaving the Web* was written by Berners-Lee, or even with a relation expressing feelings towards him ("great man").



2.2.7 Assertions and other Tag Actions :

In spite of all their differences, the previous use cases share one trait in common: the meaning of each act of tagging is function of the assertion performed. Of course, this treatment of tags is strongly reminiscent of "speech acts". Just as language can be used to perform actions and not simply in a purely descriptive fashion despite the variety of possible assertions, tags also acquire a meaning in context once it becomes obvious which action they help to perform. More details on Tag Actions are offered in section 4. Let us now turn our attention to the mechanism that is used to model the previous use cases in NiceTag.

3. NICETAG¹² AND RDF NAMED GRAPHS

3.1 Modeling and Typing Tag Actions

Carroll et al. [12] remarked that RDF does not provide any operational means, apart from reification, for making statements about graphs and relations between graphs. As a solution to this problem they proposed Named Graphs in RDF to allow publishers to communicate assertional intent and sign their assertions. The need they expressed to embody "social acts" with some record clearly resonates with scenarios of social tagging.

To model tag actions we defined a subclass of named graphs (modeled as `rdfg:Graph`, see [12]) called `nt:TagAction` which embodies one single act of tagging (see fig. 1 below). The triples contained in the named graph represent the link, modeled with the property `nt:isRelatedTo`, between an instance of the class `irw:Resource` and a sign (modeled as an instance of `rdfs:Resource`). Starting from this point, NiceTag is able to serve as a pivot-model as the signs used to tag can be modeled with all the other currently available models of tags.




Fig. 1. A Tag Action as a named graph ("nt" stands for nicetag namespace)

More importantly, our paradigm opens up new perspectives on modeling tags by providing for three degrees of freedom. (1) The model of the tagged resource can be extended to contribute to a solution to the identity crisis of the Semantic Web akin to the one proposed by [2]. (2) The modeling choice of the sign used to tag is let free (thanks to the use of `rdf:value`, cf. *infra* 5.1.1). (3) The

relation between the tagged resource and the sign allows for a fine-grained account of the semiotics of tagging. The possibilities to capture the nature of this relation are twofold: (a) the `nt:isRelatedTo` relation can be declined to *faithfully model all the possible uses of tags* already described in academic literature and their possible evolution and (b) *tag actions can be typed and specified with extra subclasses* to capture other dimensions of the tag including context and other pragmatic dimensions.

These subclasses can help distinguish, for instance, automatic tagging (`nt:AutoTagAction`) from manual tagging (`nt:ManualTagAction`). They can also help in accounting for the way in which tags are expressed. The `nt:WebConceptTagAction` would be used when signs are computer processable by design, like URIs in MOAT and CommonTag. We intentionally add "by design" because a URI acting as a MOAT "meaning" would be a "Web concept", whose meaning is sometimes construable by a human (a DBpedia URI) sometimes not (a geoname one). `nt:SyntacticTagAction` suits tagging involving complex signs like machine tags, that is, tags decomposed in a plurality of elements (machine tags can thus be seen as a kind of triple tags), that sometimes have the particularity of following a specific syntax in order to be processable by APIs. The name "machine tags" was chosen by Flickr engineers to mimic in an unconstrained fashion the structure of RDF triples. Hence, the "namespace:predicate="value"" (a very simple and unspecified way to represent things thanks to strings of characters whose components are ordered like classical EAV models) became officially supported as a means to browse in a facet-like fashion Flickr's collections

Finally, the `nt:TagAction` class is declared as a subclass of `sioc:Item` in order to account for the shareable nature of tags, which can be seen as some sort of post. This, in turn, makes it possible to describe the place where tag actions are stored with the property `sioc:has_container`, and the account (`sioc:UserAccount`) of the user (`foaf:Person`) of the tag with `sioc:has_creator`. All these elements could easily be added automatically, unbeknownst to users.

3.2 RDF/XML Source declaration

In SPARQL when querying a collection of graphs, the GRAPH keyword is used to match patterns against named graphs. However, the RDF data model focuses on expressing triples with a subject, predicate and object and neither it nor its RDF/XML syntax do provide a mechanism to specify the source of each triple.

A means to palliate such a lack was proposed in the W3C Member Submission *RDF/XML Source Declaration* [13] as an XML syntax associating to the triples encoded in RDF/XML a URI specifying their origin. It makes use of a single attribute to specify, for triples represented in RDF/XML, the source they should be attached to. The URI of the source of a triple is then: (1) the source URI specified by a `cos:graph` attribute on the XML element encoding this triple, if one exists, otherwise, (2) the source URI of the element's parent element (obtained following recursively the same rules), otherwise, (3) the base URI of the document.

The scope of a source declaration extends from the beginning of the start-element in which it appears to the end of the corresponding end-element - excluding the scope of any inner source declarations. Such a declaration applies to all elements and

¹² The complete schema of the NiceTag ontology is available here: <http://ns.inria.fr/nicetag/2009/09/25/voc#>

attributes within its scope. If no source is specified, the URL of the RDF/XML document is treated as the default source. Only one source can be declared as attribute of a single element.

```
1.<nt:TaggedResource rdf:about="twitpic.com/l4boe"
2.  cos:graph="http://delicious.com/url/7f3e344f4abbbcc3a35c243d2ad5bec#fabien_gandon_:-)">
3.  <nt:emotionalReaction>:</nt:emotionalReaction>
4.</nt:TaggedResource>
5.<nt:ManualTagAction
  rdf:about="http://delicious.com/url/7f3e344f4abbbcc3a35c243d2ad5bec#fabien_gandon_:-)">
6.  <sioc:has_creator rdf:resource="http://delicious.com/fabien_gandon"/>
7.  <sioc:has_container rdf:resource="http://delicious.com/">
8. </nt:ManualTagAction>
```

Listing 1. A tag as a named graph using RDF/XML

```
1. SELECT ?t ?a ?g WHERE {
2.   GRAPH ?tag { ?t ?a ?g }
3.   ?tag rdf:type nt:ManualTagAction }
```

Listing 2. SPARQL query to retrieve tags as named graphs.

The example in listing 1 shows how this applies to declare a tag as a named graph. Lines 1 to 4 declare the tag as a graph named http://delicious.com/url/7f3e344f4abbbcc3a35c243d2ad5bec#fabien_gandon_:-). Lines 5 to 8 reuse the name of the graph to qualify the tag as a tag created manually by user "fabien_gandon". Loading this RDF in a compliant triple store one can then run SPARQL queries like the one in listing 2 where line 2 searches for named graphs and the triples they contain, and line 3 enforces these graphs to be manually generated tags.

Note that the URI of the `nt:AnnotatedResource` in this schema is the one that identifies and gives access to the `irw:WebRepresentation` (the representation of a resource after content negotiation) which prompts the act of tagging. Tagging is made possible on the Web in this context. *Neither* shall it be conflated with the URI of the tag action, which is another URI, *nor* with the `nt:TaggedResource` being referred to, which is in no way bound by the constraints found in technical specifications.

While RDF does provide constructs to write reification quads, asserting the *reification in RDF* is not the same as asserting the original statement – and neither implies the other. Moreover, reification expands the initial triple into a total of five triples (a triple plus a reification quad) and the link between the initial triple and its reification quad is not maintained.

The attribute `rdf:ID` can be used in a property element to produce a reification of the triple that the property element generates and assert it at the same time. However, this mechanism remains at the level of triples and there is nothing in the resulting triples that explicitly identifies the original triple and links it to the reification quad. RDF provides no way to associate the subject of the reification triples with an individual triple.

Associating URIs with specific statements has to be done using mechanisms outside RDF and is one of the motivations behind RDF 2.0¹³. Likewise, statements can be made using the URI of a document as commonly done by *annotations in OWL*. In an *ad hoc* application-dependent understanding, those statements could be interpreted as if they were to be distributed over all the statements in the document. But there again we are outside RDF

and OWL and relying on likening the document to its asserted content does not sound like a good practice.

4. TAG ACTIONS AND RELATIONS

4.1 TagActions as Social Acts

Tag actions are irresistibly reminiscent of speech acts. In fact our main inspiration here was drawn from Adolf Reinach's work on Civil right, *Die apriorischen Grundlagen des bürgerlichen Rechtes* (Reinach [14]). Reinach is best known as one of the brightest students of Edmund Husserl, the founding figure of phenomenology. Reinach himself is often described as the modern father of speech act theory, long before J.L. Austin or J. Searle. Whereas Reinach found posterity as a realist phenomenologist, what is most remarkable for us is the subtlety with which he analyses the mediation between the ontological and the normative levels of civil law. Granted, in the end, he defends a view that tightly articulates what he calls "social actions" (the equivalent of speech acts) and legal norms - law being incapable of producing any concepts a philosophical foundation is called for. Reinach's theory may be fruitfully construed as a theory of mediation between social acts (and their *a priori*, ontological foundation) and civil right. A norm, says Reinach, gives existence to what is not and does as if what exists did not. In other words, it is not restricted to reality as we know it. Hence, social acts are modified once embedded in legal norms (a legal act of promising for example can satisfy different criteria to hold compared to a simple act of promising).

Similarly, some tag actions are social acts mediated through a technological system that is no less than the Web itself. While a number of tag actions are fairly equivalent to existing speech acts (categories like asserting or expressing emotions), others are modified once mediated in a digital environment like the Web. Sharing, for example, is a social action that must meet precise criteria (that that which is shared with someone be publicly known from the two parties, etc.). The `nt:Share Tag Action`, by contrast, is accomplished by sending a resource to "someone" (this is done on delicious with "for:username" tags and with "@username" tags on Twitter). A username does not guarantee that a person, a human being, "someone", is behind that account. This is the very reason behind the choice that was recently made on the SIOC discussion list¹⁴ to replace the `sioc:User` class by the more agnostic `sioc:UserAccount` class (the owner of a user account might indeed be a robot, a person, an institution, etc). Hence, while sharing normally involves at least two conditions: **(a)** something is shared between at least two persons, and **(b)** both persons have to acknowledge that something is being shared, on social tagging Web services, both of these conditions are violated, sharing is thus closer to indicating and sending information to contacts.

The last category of tag action concerns those conceivable only in this context, depending as they do on the technical possibilities opened by the Web. On YouTube users now have the possibility to isolate media fragments of videos at will, in order to contextualize their comments by *pointing* (`nt:PointsAt`) at a specific part of a resource. This mechanism could be reused along the same lines with tagging so that it would become possible to tag specific parts of a media. This is especially useful when

¹³ <http://www.w3.org/2009/12/rdf-ws/>

¹⁴ The final decision was made the day this paper was submitted: <http://sioc-project.org/node/341>

dealing with temporal objects which, not unlike scrolls in contrast with codices, impose their unfolding order and raise the need for new forms of “bookmarking”. The need to fragment resources is also felt with traditional Web pages. One may for example want to tag a specific comment of an article. XPointer was thought to be a good candidate to fulfill this scenario. Unfortunately, arbitrary technical limitations¹⁵ were responsible for the lack of advancement and the project was abandoned. Other proposals include defining access roles through tagging (allowing one’s parents to access a picture simply by using the tag “parents” and the Tag Action `nt:GrantAccessRightsTo` – an idea currently examined by various research teams¹⁶).

Status	nt:TagAction	nt:isRelatedTo
New	nt:PointAt	nt:hasPart
New	nt:GrantAccessRightsTo	nt:canBeReadBy (etc.)
New	nt:Aggregate	nt:hasCommunityTag
Existing (modified)	nt:Share	nt:sentTo
Existing (modified)	nt:Assert	nt:isAbout nt:isRelevant nt:isRelevantToSo, nt:isRelevantToSt,
Existing	nt:Evaluate	nt:isWorth
Existing	nt:SetTask	nt:elicitsAction
Existing	nt:ExpressEmotion	nt:makesMeFeel
Existing	nt:Ask	nt:raisesQuestionAbout

Fig. 2. `nt:TagAction` subclasses and `nt:isRelatedTo` sub properties in NiceTag (“existing”, “modified” and “new” refer to the status of speech acts : those that existed before being translated on the Web, those that were modified and those that essentially depend on the Web to be performed).

4.2 isRelatedTo

Inspired by previous studies, in particular the seminal work of Golder & Huberman [15], we modeled the different possible uses of tags with sub-properties of the `nt:isRelatedTo` property. Two of our broadest classes, factual and personal, were already proposed in Sen et al. [16]. `nt:Assert` (used to express factual knowledge about the resource being tagged) is associated to the relation that appears to be the most widespread of all properties, employed when the sign used to tag describes the topic of a resource: `nt:isAbout`. Lots of models of tagging take for granted

that this is the relation by default. There seems, consequently, to be no pressing need whatsoever to implement it. It is not true as evidenced by two properties that are on the same level: `nt:isAbout` and `nt:isAKindOf`. `nt:isAKindOf` is intended for all cases in which a tag is used to distinguish between types of `nt:WebResource` (e.g.: forum, video, blog, picture, etc.). While `nt:isAbout` is perceived to be broad enough to justify its special status, it is certainly superseded in this regard by the `nt:Assert` Tag Action, a proxy for all the state-of-affairs RDF relations available in existing vocabularies that may be asserted between a given resource and the things a sign refers to (written by, bought by, composed of, caused by, theorem of, etc.). Another wholesome property whose virtue is to limit the inadequate use of `nt:isAbout` is `nt:relevant` (to someone or something). A resource might indeed be said to be relevant to a thesis, some studies, etc. without being even remotely about any of these elements. Accuracy dictates that aboutness and relevance be thus unmistakably distinguished from one another. `nt:express` Emotion associates a resource with an adjective or any kind of sign expressing a quality or emotion stirred up by a resource; typical examples are interjections and smileys (e.g.: “wow!”, “<:o)”). `nt:isWorth` is meant whenever a resource is evaluated, ranked, etc. (e.g.: “nice”, “*****” – Tag Action: Evaluate). `nt:makesMeFeel` (Tag Action :ExpressEmotion) is fit for use with signs expressing emotions. Another sub property of `nt:isRelatedTo` covers uses of tags intended to make sense first and foremost to the applier. This includes Golder & Huberman’s class “task organizing” (tag action `nt:SetTask` and relation `nt:elicitsAction` used with “todo” tags). Similarly, we introduce the property `nt:hasCommunitySign` for collectively approved signs designed to (Tag Action) `nt:Aggregate` resources revolving around a shared event, goal or entity known by all the members of a community/audience. For example, we used the tag “#vocampnice2009” to gather and share resources across multiple social Web applications about the VoCamp where the NiceTag framework was conceived.

5. TAG IDENTITY

5.1 From tagging to folksonomies

5.1.1 The identity of tags in NiceTag

We would like in this paragraph to examine the relationships between tagging models, uses and implementations by contrasting one model, at the core of the NiceTag ontology (proposed by [19] and [20]), and the results of the analysis of the data extracted from delicious.com (the choice of this website is motivated in recognition of the fact that it has set so many standards that the way it handles tagging has become almost proverbial). The former provides a vocabulary to harness the semantics of individual “tag actions” between one resource and one “sign”. Websites like delicious.com, in turn, offer no faithful equivalent of such a model because they represent tags according to much looser criteria. By distinguishing representations of tags on the Web from their models, it becomes clear that every single act of tagging can be decomposed in a set of basic elements, some of which are lost in most, if not all, current implementations. This is especially true of relations - sub properties of `nt:isRelatedTo` - generally left implicit. Through source declarations and named graphs, each tag action is given a LOD URI corresponding to the various dimensions implied every time a user (or a machine, in the eventuality of automatic tagging) tags a resource. Hence, tags,

¹⁵ Limitations exceeded by the W3C Media Fragment Working Group (MFWG) who discussed the need for various new HTTP headers in order to process URI + fragment requests (an issue raised by fragment parameters being stripped off from URIs by Web browsers and the issue of caching Web proxies).

¹⁶ [17]. We express our sincere thanks to Michel Buffa and Catherine Faron for calling our attention to this matter. See also [18] where similar ideas are presented.

instead of "information resources" [2], traditionally a word or group of words embedded in a hyperlink, might be better understood as non-information resources. Tags are sometimes taken to be no different from any HTML `<a>` element except for the fact that a tag may receive additional markup (the `rel="tag"` microformat in particular¹⁷, to indicate that the link is a tag).¹⁸

While the tagged resource in NiceTag basic schema (fig.1) espouses the many distinctions drawn in IRW ontology, it remains clear, however, that this is in no way tantamount to loosening the identity criteria of tag actions. Tag Actions in NiceTag, generally called "tags" elsewhere, consist in at least one resource, one relation and one sign. A set containing all these elements is being typed with additional information thanks to the mechanisms described in sections 3 and 4. This definition rests on very strict identity criteria for Tag Actions. As their name betokens, these are individuated so as to reflect all the data that derive from a single act of tagging. Yet, functional identification and access, relations that characteristically link URIs and resources, are absent from this schema. Of course, if our model was to be represented, this could be accomplished through a "traditional" tag, one that is embedded in some hyperlink. It could even, thanks to the information encapsulated in the named graph, grant access to the one `nt:WebResource` which prompted the tag action (a feature yet unseen) and would at the same time assess: (a) what kind of resource a user is referring to; (b) how exactly it is related to a chosen label; and (c) what this label means and what it refers to.

This is achieved by adding specific range and domain constraints on each sub property of `nt:isRelatedTo` (for instance, the range of `nt:relevantToSb` could be `foaf:Agent`, while that of `nt:sentTo` could be `sioc:UserAccount`). That way relations aptly shed some light on both the referent of the sign and the status of the `nt:TaggedResource` (either a `irw:WebResource` just as the one that initiates the act of tagging, the representation of this `irw:WebResource` subject to many informational hazards, the information object that the `irw:WebResource` is a realization

of, a non-information resource like a thing or an abstract concept that the `irw:WebResource` aims to represent, etc.).

It would frankly be unrealistic to expect users to tag resources with only other resources. At the same time, we do not wish to simply shift to literals because that would make it impossible to add range constraints other than `rdfs:Resource`. To cope with this difficulty, and give users the liberty to choose whatever label they see fit, we split the triple into two triples, one with a blank node (bnode) as its value instead of a literal, to keep the range of the resource that corresponds to the label. The other triple links that bnode with a literal thanks to `rdf:value` as follows¹⁹:

Listing 3 Modeling property ranges with bnodes and `rdf:value`

```
1. nt:sentTo rdfs:range sioc:UserAccount .
2. nt:TaggedResource nt:sentTo _:v .
3. _:v rdf:value "aamonnz".
```

This increased expressivity comes with a high price to pay since SPARQL queries get more complicated. From:

```
<Select ?x ?y { ?x nt:sentTo ?y}
```

we now have to ask, to get the same answer:

Listing 4 SPARQL query with range and "proxy" value:

```
1. Select ?x ?y where {
2. {
3.   ?x nt:sentTo ?y .
4.   FILTER (isURI(?y))
5. }
6. UNION
7. {
8.   ?x nt:sentTo ?z
9.   ?z rdf:value ?y.
10. FILTER (isBlank(?z))
11. }
12.}
```

Further work is thus needed to determine whether expressivity is to be favored over simplicity of use.

5.1.2 The identity of tags in delicious

Examples of tagging on delicious.com, for example, do not appear to follow the same criteria to define what a tag is. These are indeed considerably loosened. On delicious a community tag identifies a resource as much as it gives access to its representation, the latter being constantly evolving over time (a collection of bookmarks and tags is bound to grow). It might be added, thus departing slightly from the interpretations given in the W3C document entitled *Architecture of the World Wide Web, Volume One*²⁰, that while a URI identifies an information resource in a functional way, such as a periodically updated report on the weather in the State of Oaxaca, which can be represented online in such and such a way, it may also refer (`irw:refersTo`) to a *non-information resource* that corresponds to the current (real-world) weather in this part of Mexico. Similarly, a URI like http://delicious.com/fabien_gandon/web in delicious.com, would (a) *identify* an information resource (a document that consists in a list of resources picked up by the community users and assembled thanks to their sharing a common label), (b) grant access (`irw:accesses`) to a constantly changing representation of this `irw:WebResource` (as users add new bookmarks falling under a given label), and (c) *refer to* a community tag, understood as a non-information resource (a collective social act, considered in all its dimensions) (cf. Figure 3). The aforementioned URI,

¹⁷ <http://microformats.org/wiki/Rel-Tag>.

¹⁸ Moreover, this definition, summarized by Tantek Çelik on microformat.org raises other problems: "By adding `rel="tag"` to a hyperlink, a page indicates that the destination of that hyperlink is an author-designated "tag" (or keyword/subject) for the current page. Note that a tag may just refer to a major portion of the current page (i.e. a blog post). e.g. by placing this link on a page, `tech` the author indicates that the page (or some portion of the page) has the tag "tech". The linked page SHOULD exist, and it is the linked page, rather than the link text, that defines the tag. The last path component of the URL is the text of the tag, so `fish` would indicate the tag "tech" rather than "fish." From this definition it appears that tagging is choosing a URI instead of a label which seems to contradict just about everything we know about tagging. Consequently we depart from [21]'s Upper Tag Ontology (UTO) since it rests on similar assumptions: "Tags are nothing more special than a typed hyperlink. We can use "rel" attribute to type hyperlinks". Actually, tags do exist without being embedded in hyperlink like most machine tags on Flickr do, as was the case with hash tags on Twitter at first.

¹⁹ For clarity purpose we temporarily shift to the Turtle notation.

²⁰ <http://www.w3.org/TR/webarch/>

http://delicious.com/fabien_gandon/web, gives access only to a Web representation of several tag actions, in other words a web page listing all the bookmarks of user fabien_gandon tagged with the sign “web”. We can see that URIs given on delicious.com do not allow for a fine-grained individuation of tags at the level of tag actions.




Figure 3 Abstract representation of the interplay between URIs, resources and tags (labels).

One of the aim of the NiceTag model is to account for the continuum between individual tag actions and tags in their current implementations (`nt:CollectiveTagActions`). Thanks to its fine-grained model that individuates each tag action thoroughly, one can end ambiguity when two tags have the same sign by specifying the relation that links this sign to the tagged resource using the panel of subproperties of `nt:isRelatedTo`. On the one hand, when the choice does not exist, this is clearly a limitation for a given act of tagging prompted by a `irw:WebResource X` where a sign `Y` attached to a tagged resource may express many different things depending on the relation holding between the two. On the other hand, similarity of signs is a good *source of serendipity* because two tag actions with similar signs, prompted by the same `irw:WebResource`, may in fact express different things or different aspects of the same thing. In any case, *delicious* never intended to represent such relations explaining why once a sign has been applied in relation to a given `irw:WebResource` it cannot be reused in the same context.

Suffices to get rid of the typed relations (sub properties of `nt:isRelatedTo`), the cardinality constraints that limit the number of resource assigned to a Tag Action and then the burden of identity necessarily rests on the sign (and possibly the `sioc:Item` value). Then, with SCOT property `nt:spellingVariant`, signs themselves need not be strictly identical to be considered equivalent (e.g. "Paris", "paris", two different words in French). This explains the transition from *tagging* to *folksonomies* (personal or collective), from single well-individuated Tag Actions to Tag Actions which associate for undisclosed reasons (i.e., the missing relations) a variety of resources to a sign which admits a plurality of variants. This is how serendipity is accidentally fostered. Within the NiceTag framework we aim to make this feature an intentional one whose parameters – the criteria of identity that hold for tags – can be specified at will.

6. SCENARIOS & USES-CASES

A model would not be complete if no indication was provided of its potential and actual uses. We can distinguish at least three potential scenarios involving NiceTag.

The first is consistent with current implementations of tagging and requires no further addition. Suffice to lower the identity criteria of a tag by dropping the relations between tagged resources and

labels and sticking to the `nt:Assert` tag action, and no additional complexity will be added to what exists. The use of named graphs and a theoretically robust model will simply makes things less quirky an act in favor of a unifying model of tagging.

The second scenario involves one or more typed tag actions with additional relations. Such relations (especially the factual ones, consistent with `nt:Assert`) are for the greatest part found outside of NiceTag. They are relevant in direct proportion with the main topic of the Web service they are a part of: social bookmarking website with an emphasis on books, like LibraryThing, would be well-inspired to tap into book-related properties and ontologies while their picture-orientated equivalents (such as Flickr) could reuse Hayes and Warren’s work on images²¹. NiceTag, being entirely open as regards the relations subsumed under `nt:relatedTo`, was design to adapt to as many contexts as possible. Other websites, where communication between users is an important feature and where tagging is a success, could decide to implement tag actions like `nt:Share` to benefit from the union of both dynamics, communicating and tagging (the aim and the method). Actions like sharing or sending a message do no depend on tagging and this does not necessarily represent the most efficient way to accomplish them. Yet, people may favor a method that has proven successful in the past (plus those communication acts realized through tagging may at the same time play the role of metadata used for indexing).

Finally, the last scenario involves the entire model. A legitimate question to ask is whether interfaces will make it easy to shift from a paradigm where users type their own tags, to a paradigm where users benefit more and more from the work of algorithms that present automatic recommendations. A similar evolution has already been witnessed in the past in the field of video games (a domain where the lack of sound interfaces is a no-go). In the first adventures games of the 80s, actions were performed by typing keywords and matching them against a pre-defined vocabulary. Then, when the SCUMM (Script Creation Utility for Maniac Mansion) scripting language appeared, the command line disappeared and instead the player had to choose between a set of actions, differing from games to games. Now that the interplay between serious games and tagging has become a matter of academic study²² with wild appeal, the lessons learned in other fields, could help us overcome the envisioned difficulties.

7. CONCLUSION

Though the success of tagging systems is due to their extreme simplicity and immediacy of use, the limitation of dealing with unstructured content appears straightforward, and users have been shown to long for more efficient and creative way of using tags to perform a wild variety of actions (a process reminiscent of what is sometimes called “function creep”). Likewise, the “killer application” of semantic tagging could very well actually be tagging itself, but tagging as fully implemented for the first time thanks to a comprehensive model accounting for its versatility in forms and social uses. We believe that this project, already partly undertaken by current models of tagging, will strongly benefit

²¹ <http://www.slideshare.net/carmapro/warren-hayes-ai-symposium>

²² It should never be forgotten that Flickr emerged over the remains of Game Neverending, a Passive Multiplayer Online Role Playing Game as noted in [22].

from a shift of focus on the manifold relations between a tagged resource and a tag. This is what NiceTag is targeted at: to be able to express the various uses and forms of tagging by maximizing both the structure of data and freedom of users. Our proposal to harness the Semantic Web technologies to tagging interfaces comes as a complement to other approaches (such as SCOT, SIOC, CT or MOAT, see [11], [23] and [24]) which tend to rely on the user's will to specify the meaning of their tags. The NiceTag framework aims at introducing little steps of semantics in existing interfaces, according to the taggers' needs, thus enriching current tagging systems and keeping their essential simplicity. Tools such as Twitter offer new and exciting use cases and call for a thorough analysis to enrich the list of tag actions presented here. NiceTag will move on to reflect and leverage the creativity displayed every day by users in concrete situations, a program currently being implemented in the ISICIL project²³.

8. ACKNOWLEDGEMENT

Alexandre Monnin would like to thank INRIA Sophia Antipolis and the Edelweiss team. This work was funded in part through a visiting scholarship that took place between March and April 2010.

9. REFERENCES

- [1] Hayes, P., Halpin, H., In defense of ambiguity.: *International Journal of Semantic Web Information Systems* 4(2), 1-18, (2008)
- [2] Halpin, H. and Presutti, V. An Ontology of Resources: Solving the Identity Crisis. In *Proceedings of European Semantic Web Conference: Research and Applications*, (2009)
- [3] Kripke, S., Naming and necessity, 2nd ed., Cambridge (Massachusetts), Harvard University Press, (1980)
- [4] Gangemi, A. and Presutti, V. The bourne identity of a Web resource. In *Proceedings of Identity Reference and the Web Workshop (IRW) at the WWW Conference*, Roma, Italy, (2006)
- [5] Presutti, V. and Gangemi, A. Towards an OWL ontology for Identity on the Web. In *Semantic Web Applications and Perspectives (SWAP)*, Trento, Italy, (2006)
- [6] Gangemi A. and Presutti V. A grounded ontology for identity and reference of web resources. In *Proceedings of I3 Workshop at WWW2007*, Banff, Canada, (2007)
- [7] Presutti V. and Gangemi A. Identity of Resources and Entities on the Web.: *International Journal on Semantic Web and Information Systems*, 4(2), (2008).
- [8] Bouquet, P., Palpanas, T., Stoermer, H., Vignolo, M., In *Proceedings of the 4th Asian Conference on The Semantic Web*, Shanghai, LNCS 5926, Shanghai, China, (2009)
- [9] Halpin, H., Thompson, Henry, S. Web Proper Names: Naming Referents on the Web. The Semantic Computing Initiative Workshop (SEC, WWW), Chiba, Japan, (2005)
- [10] Recanati, F., Direct Reference: From Language to Thought, Oxford, UK, Cambridge, USA, Basil Blackwell, (1993)
- [11] Passant, A., Laublet, P. Meaning of a tag: A collaborative approach to bridge the gap between tagging and linked data. In *WWW Workshop on Linked Data on the Web (LDOW)*, Beijing, China, (2008)
- [12] Carroll, J. J., Bizer, C., Hayes, P., and Stickler, P. Named graphs, provenance and trust. In *Proceedings of the 14th international Conference on World Wide Web*, Chiba, Japan, (2005)
- [13] Gandon F., Bottolier, V., Corby, O., Durville, P. RDF/XML source declaration. W3C Member Submission. <http://www.w3.org/Submission/rdfsource/>, (2007)
- [14] Reinach, A., The A Priori Foundations of Civil Law, (trans. John Crosby): *Aletheia*, 3, 1–142, (1983).
- [15] Golder, S. and Huberman, B. Usage patterns of collaborative tagging systems.: *Jour. of Information Science* 32, 2, (2006)
- [16] Sen, S., Lam, S. K., Rashid, A., Cosley, D., Frankowski, D., Osterhouse, J., Harper, F. M., and Riedl, J. tagging, communities, vocabulary, evolution. In *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, Banff, Alberta, Canada, (2006)
- [17] Nasirifard, P., Peristeras, V., and Decker, S. An Annotation-Based Access Control Model and Tools for Collaborative Information Spaces. In *Proc. World Summit Emerging Technologies and information Systems For the Knowledge Society*, Athens, Greece, (2008)
- [18] Au Yeung, C. M., Kagal, L., Gibbins, N. and Shadbolt, N. Providing Access Control to Online Photo Albums Based on Tags and Linked Data. In *AAAI Spring Symposium on Social Semantic Web: Where Web 2.0 Meets Web 3.0*, Stanford, USA, (2009)
- [19] Limpens, F., Monnin, A., Laniado, D. Gandon, F.: NiceTag Ontology: tags as named graphs, *Proceeding of SNI09 - Social Networks Interoperability, 1st International Workshop at the 4th Asian Semantic Web Conference (ASWC 2009)*, Shanghai, China, (2009)
- [20] Monnin, A., Limpens, F., Laniado, D., Gandon, F., L'ontologie NiceTag : les tags en tant que graphes nommés. In *Actes de l'atelier Web Social, 10ième Conférence Internationale Francophone sur l'Extraction et la Gestion des Connaissances*, Hammamet, Tunisia, (2010)
- [21] Ding, Y., Toma, I., Kang, S., Zhang, Z., and Fried, M. Mediating and Analyzing Social Data. In *Proceedings of OTM Confederated international Conferences*, Monterrey, Mexico, (2008)
- [22] Kim, H. L., Scerri, S., Breslin, J. G., Decker, S., and Kim, H. G. The state of the art in tag ontologies: a semantic model for tagging and folksonomies. In *Proceedings of the International Conference on Dublin Core and Metadata Applications*. 128-137, Berlin, Germany, (2008)
- [23] Kim, H.L., Passant, A., Breslin, J.G., Scerri, S., Decker, S., Review and Alignment of Tag Ontologies for Semantically-Linked Data in Collaborative Tagging Spaces, *IEEE International Conference on Semantic Computing (ICSC)*, Santa Clara, USA, (2008)
- [24] Monnin, A., From Game Neverending to Flickr. Tagging systems as ludic systems and their consequences. In *Proceedings of the WebSci'09: Society On-Line*, Athens, Greece, (2009).
- [25] Limpens, F., Gandon, F., Buffa, M., Helping online communities to semantically enrich folksonomies, Web Science Conference. In: *Proceedings of WebSci10: Extending the Frontiers of Society On-Line*, Raleigh, USA, (2010)

²³ <http://isicil.inria.fr>, see [25].