Enhancing Subjective Ontologies with Social Tagging Systems

Dennis Hooijmaijers[1] **Markus Stumptner**[1]

[1] Advanced Computing Research Centre
University of South Australia,
Mawson Lakes Blvd, Mawson Lakes, South Australia 5095,
Email: {dennis,mst}@cs.unisa.edu.au

Abstract

Social computing websites provide a framework for users to submit (and discuss) content, while often providing a user created categorization system, dubbed folksonomies. Additionally users may have the ability to specify friends (and foes) within the social network, populated by registered users of the social website. Integrating folksonomies with ontologies provides the ability to capture the semantic relationships between terms. This allows for greater reasoning to refine searches.

Most categorisation for social computing is currently achieved by a meta-tagging system dubbed folksonomies, which fails to take into consideration the semantics that occur between terms. By integrating social networks with an ontology we aim to provide mechanisms to capture agreement in opinions and to capture and enhance the semantics of the social network, allowing for greater reasoning to refine searches.

Keywords: Ontology, Subjective Logic, Uncertainty, Folksonomy

1 Introduction

Knowledge is subjective and this can be seen by the many and varied uses of terminology in on-line communities. User generated content has expanded from comments on forums to include, reviews, synopsis, blogs, and multimedia. Where originally the website provider would post the content for discussion, current sites provide frameworks for users to upload content, or synopsis of on-line content, and discussion boards to allow for interaction between users. Categorizing content allows users to find objects of interest within a particular domain.

Some sites (e.g. Digg[1], Slashdot[2]) provide content areas for a user to select when uploading, while others augment this with user defined tags (metadata), such as Flickr[3] and YouTube[4]. This allows the users to select, or create, appropriate descriptive terms for their content. These user defined tags are collectively known as folksonomies (Al-Khalifa & Davis 2006), and can provide insight into popular terminology within the domain of a community.

To further assist users to sort through on-line submissions a user-based rating system is applied. This allows users to find the most popular submissions providing a collaborative voting mechanism. These rating systems can be as simple as a 5-star rating system where each user provides a rating between 0 and 5 (YouTube, Flickr), or a positive (*thumbs up*) and negative (*thumbs down*) rating approach where users state whether they like or dislike an object (Digg). Additionally the ratings can be categorised (Slashdot, Digg[5] with additional terms (*informative*, *insightful*, *troll* etc.) to assist in clarity of ratings.

These frameworks also allow users to specify other users as friends (and foes), allowing for users to refine the popularity ratings to find submissions that their friends have found popular.

RIPOSTE is a knowledge management framework, that allows for user ratings to be captured within context. It is based on OWL DL, as defined by the w3c[6], a knowledge representation language for the semantic web and Subjective Logic(Jøsang 2002), an extension of probabilistic logic. RIPOSTE provides a meta-ontology which captures authors, and their opinions with the ontological resources they provide. The combined ontology and meta-ontology is a *subjective ontology.* RIPOSTE provides mechanisms to capture, integrate, manipulate and reason over subjective ontologies. This work extends the RIPOSTE framework to provide additional functionality for capturing social networks and to integrate folksonomies with the subjective ontologies.

Adding semantic value to folksonomies (Angeletou et al. 2007, van Damme et al. 2007, Szomszor et al. 2007) has been identified as a solution for improving *recall* for queries (i.e. finding objects related to Sydney when querying for cities). The identification of relations is achieved by using background information from multiple sources. Most current techniques for integration use on-line lexicons (e.g. WordNet[7]) and Semantic Web technologies. Lexicons are used for discovering homonyms, meronyms and synonyms, while the Semantic Web for richer semantic relationships (i.e. object properties, disjoint, hierarchical). Within the Semantic Web relationships between terms may exist within a single ontology or multiple ontologies (Angeletou et al. 2007). Our approach also uses online lexicons and the Semantic Web, although by including probabilistic approach we can capture all ambiguities of a term within the resultant ontology.

Disambiguation of the terminology, can then be discovered by dynamic filtering for a specific query (Hooijmaijers & Stumptner 2006). Additionally by including the social network within the ontology the ability to capture the author's opinion for specific con-

[1] http://digg.com
[2] http://slashdot.org
[3] http://www.flickr.com
[4] http://www.youtube.com
[5] Digg only provides categorisation for negative ratings. i.e. the object is considered *lame* or a *duplicate* etc.
[6] http://www.w3.org/2004/OWL/
[7] http://wordnet.princeton.edu/

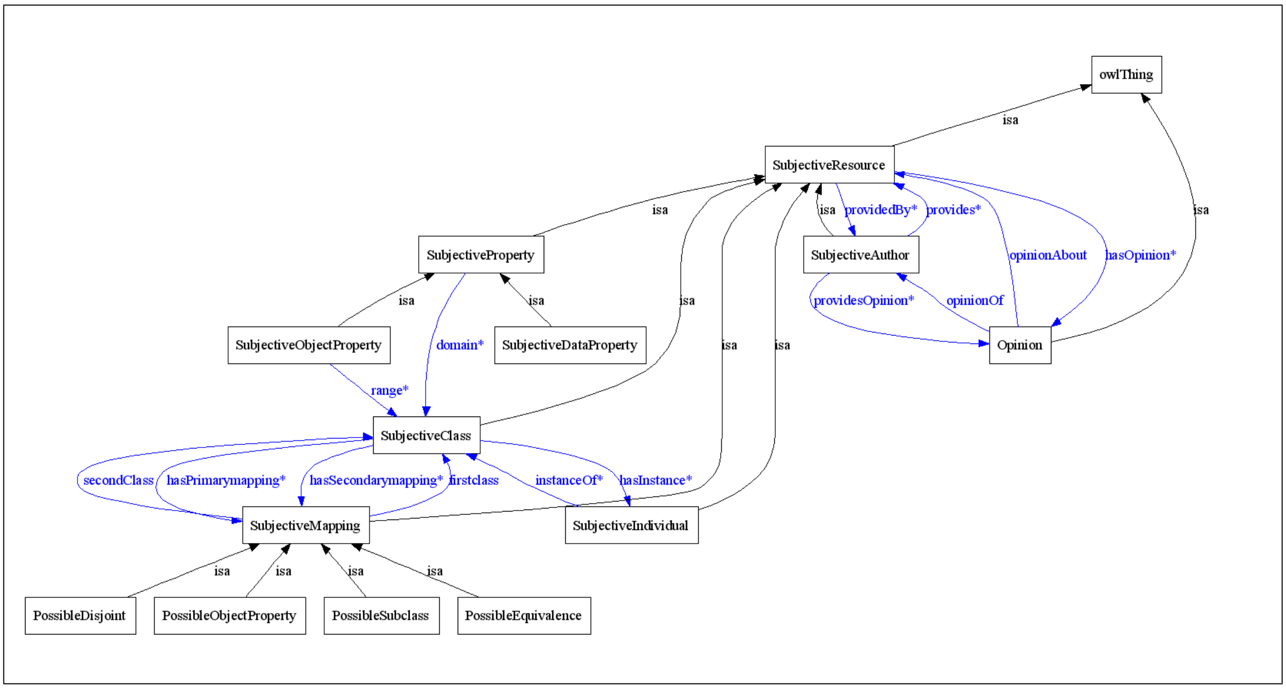

Figure 1: The RIPOSTE Meta-Ontology

cepts and easily discover the group of authors that agree and/or disagree with that opinion.

2 RIPOSTE

RIPOSTE (Hooijmaijers & Stumptner 2006) is a knowledge management framework, based on OWL DL and subjective logic, that provides mechanisms for capturing, integrating and evolving ontologies with the *authors* that provide the ontological resources. The term author is used to differentiate between members of a community who provide resources and the users who only utilise content. We allow both authors and users to provide opinions and a user will become an author when they provide any resource. It uses a meta-ontology to provide additional information about each resource. The resources are subjective in that each author may have a different opinion on how a resource should be used and this allows for filtering and reasoning based on similar opinions. The subjective nature of the knowledge is captured using subjective logic, an extension of probabilistic logic (Jøsang 2002). This allows for opinions to be compared to capture and manipulate the probability that two authors agree (or disagree) on the usage of a term.

To assist with capturing ontology mappings RIPOSTE allows for the capture of the mappings between ontologies by creating a set of individuals of SubjectiveMapping. This allows for opinions to be formed about the mapping and for the author (which can be an automated mapping agent) to be captured. This can be used to generate a final integrated ontology based on multiple mappings (Hooijmaijers & Stumptner 2008).

2.1 Subjective Logic

Subjective Logic is an extension to probabilistic logic that represents an author's (A) beliefs about an object (x) as *opinions* (w_x^A). An opinion is the combination of belief (b_x), disbelief (d_x) and uncertainty (u_x) in a given resource (x), where

$$b_x + d_x + u_x = 1$$

and the opinion,

$$(w_x^A) = (b_x, d_x, u_x, a_x)$$

where a_x represents the size of the state space, atomicity, from which x is taken, Figure 2 and is used as the *a priori* probability in the absence of evidence. The expectation value represents the likelihood of a positive belief given no evidence and is calculated by:

$$E(x) = b_x + a_x u_x \tag{1}$$

For resources within the ontology we use the sense count from WordNet (e.g. Figure 6) for the opinions, which represent the possible choices for use of a keyword for the state space. For classes and tags we use a_x to work out the possible relationships which for *unrelated* and *equivalent* classes is given by:

$$a_x = \frac{1}{senseCount + 1} \tag{2}$$

For trust between authors we use the state space of {trust, distrust} $a_x = 0.5$.

In this work we use the subjective logic operators, *Consensus* and *Discount*, as defined by Jøsang (Jøsang 2002). Consensus ($\oplus$) is used to, fairly and equally, combine two possibly conflicting opinions about a given resource and Discount ($\otimes$) is used to capture reputation (trust transitivity).

Consensus (Jøsang 2002) Let $w_x^P = (b_x^P, d_x^P, u_x^P, a_x^P)$ and $w_x^Q = (b_x^Q, d_x^Q, u_x^Q, a_x^Q)$ be opinions respectively held by P and Q about the same proposition x. Let $w_x^{PQ} = (b_x^{PQ}, d_x^{PQ}, u_x^{PQ}, a_x^{PQ})$ be the opinion such that:

1. $b_x^{P,Q} = (b_x^P u_x^Q + b_x^Q u_x^P)/k$
2. $u_x^{P,Q} = (u_x^P u_x^Q)/k$
3. $d_x^{P,Q} = (d_x^P u_x^Q + d_x^Q u_x^P)/k$
4. $a_x^{P,Q} = \frac{(a_x^P + a_x^Q - (a_x^P + a_x^Q) u_x^P u_x^Q)}{u_x^P + u_x^Q - 2 u_x^P u_x^Q}$

where $k = u_x^P + u_x^Q - u_x^P u_x^Q$ such that $k \neq 0$ and $a_x^{P,Q} = (a_x^P + a_x^Q)/2$ when $a_x^P, a_x^Q = 1$. Then w_x^{PQ} is called the consensus between w_x^P and w_x^Q, representing an imaginary agent $[P,Q]$'s opinion about x, as if she represented both P and Q. By using the symbol '$\oplus$' to designate this operator, we define $w_x^{PQ} \equiv w_Q^P \oplus w_x^Q$.

Discount (Jøsang 2002) Let P and Q be two agents where $w_Q^P = (b_Q^P, d_Q^P,\ u_Q^P, a_Q^P)$ is P's opinion about Q's advice, and let x be the proposition where $w_x^Q = (b_x^Q, d_x^Q, u_x^Q, a_x^Q)$ is Q's opinion about x expressed in an advice to P. Let $w_x^{PQ} = (b_x^{PQ}, d_x^{PQ}, u_x^{PQ}, a_x^{PQ})$ be the opinion such that:

1. $b_x^{PQ} = b_Q^P b_x^Q$
2. $u_x^{PQ} = d_Q^P + u_Q^P + b_Q^P u_x^Q$
3. $d_x^{PQ} = b_Q^P d_x^Q$
4. $a_x^{PQ} = a_x^Q$

By using the symbol '$\otimes$' to designate this operator, we define $w_x^{PQ} \equiv w_Q^P \otimes w_x^Q$.

Given two authors, Allan and Bob, if Bob uses the term 'Australia' to describe an object (x) from Figure 4, and *Allan* trusts that *Bob* is right 66% of the time then:

$$w_{Bob}^{Allan} = (0.66, 0.34, 0, 0.5)$$

and Bob believes he is correct:

$$w_x^{Bob} = (1, 0, 0, 0.17)$$

Then Allan's opinion that Bob has tagged x correctly is calculated using discount $\otimes$:

$$w_x^{Allan} = (0.66, 0, 0.34, 0.17)$$

If a third author Carl also provides the tag 'Australia' for x ($w_x^{Carl} = (1, 0, 0, 0.17)$) and Allan trusts Carl 50% using discount gives:

$$w_x^{Allan} = (0.5, 0, 0.5, 0.17)$$

This will be combined using consensus ($\oplus$):

$$w_x^{Allan} = (0.75, 0, 0.25, 0.17)$$

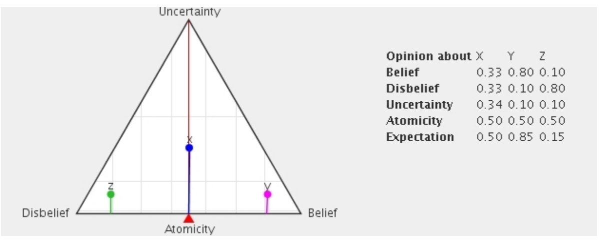

Figure 2: Subjective Logic opinion triangle

2.2 Meta-ontology

To capture the additional information about subjectivity and the author with a resource we have introduced a meta-ontology layer (see Figure 1). RIPOSTE is a framework that provides the ability to capture author information with the meta-data provided. It provides mechanisms for integrating author meta-data and the ability to filter the resultant ontology to allow for queries based on different opinions. The key concept of RIPOSTE is that all resources of an ontology are subjective to the author(s) that provides it. RIPOSTE meta-ontology, shown in Figure 1 provides the framework for capturing the author's opinion and the relationship between authors and the resources they provide. Each SubjectiveResource has an opinion, a Boolean filter isAccepted and an object property relating the resource to the providing authors. To annotate an ontology, O for Riposte it is necessary to make all classes, $C \in O$ an instance of SubjectiveClass. Additionally all $C \in O$ that have only *owlThing* as a super-class become a *subClass* of SubjectiveClass, as shown in Figure 3. By creating each class as an instance of SubjectiveClass we can apply the credibility attribute, filter and relationship to the author to the class, while the sub-classing allows the additional credibility information to be provided to the individual instances.

3 Classification

Classification can be performed by either creating an explicit semantic representation, such as an ontology, or by allowing the semantics to evolve from authors creating their own tags. The tags from all authors in a community are collated into a folksonomy.

An ontology can be described as a collection of resources that explicitly and formally conceptualise a domain model (Guarino 1997). Figure 3 shows an example class structure of an ontology for photographs taken in locations, the ellipses represent the classes in the ontology and the arcs the semantic relations between the classes. The white box (SubjectiveClass) shows the replaced owlThing. For simplicity axioms and instances have not been shown.

Definition An OWL ontology, O, is considered to be a 4-tuple of resources $< C, S, I, \chi >$ where C is a set of *classes*, S is a set of *semantic relations* which relate two or more classes, I is a set of *instances* of a class, and χ is a set of *axioms*.

An ontology in this work is a subjective view of a domain which we call a **subjective ontology**. A subjective ontology provides credibility ratings for all resources and links all resources within the ontology to the authors that provided them. This provides a mechanism to manage an integrated ontology by each author that has contributed resources.

Definition Let an author A_i be a provider of a resource R for an ontology $R \in O^D$ then their opinion is given by $\omega_R^{A_i}$. We define a subjective ontology as, O_{A_i}, is an extension of an ontology that relates an author, A_i, and their opinion $\omega_R^{A_i}$ to the provided ontological resource R. $(\omega_R^{A_i}, A_i, R)$

A folksonomy is a term used to describe a tag-based system that has evolved over time through the direct input by the user base (Al-Khalifa & Davis 2006). Flickr, a photograph sharing site, and del.icio.us, a social bookmarking site, are examples of implemented folksonomies. A folksonomy is created by users supplying objects and selecting, or creating, tags that describe the object. Tags are assumed to represent a

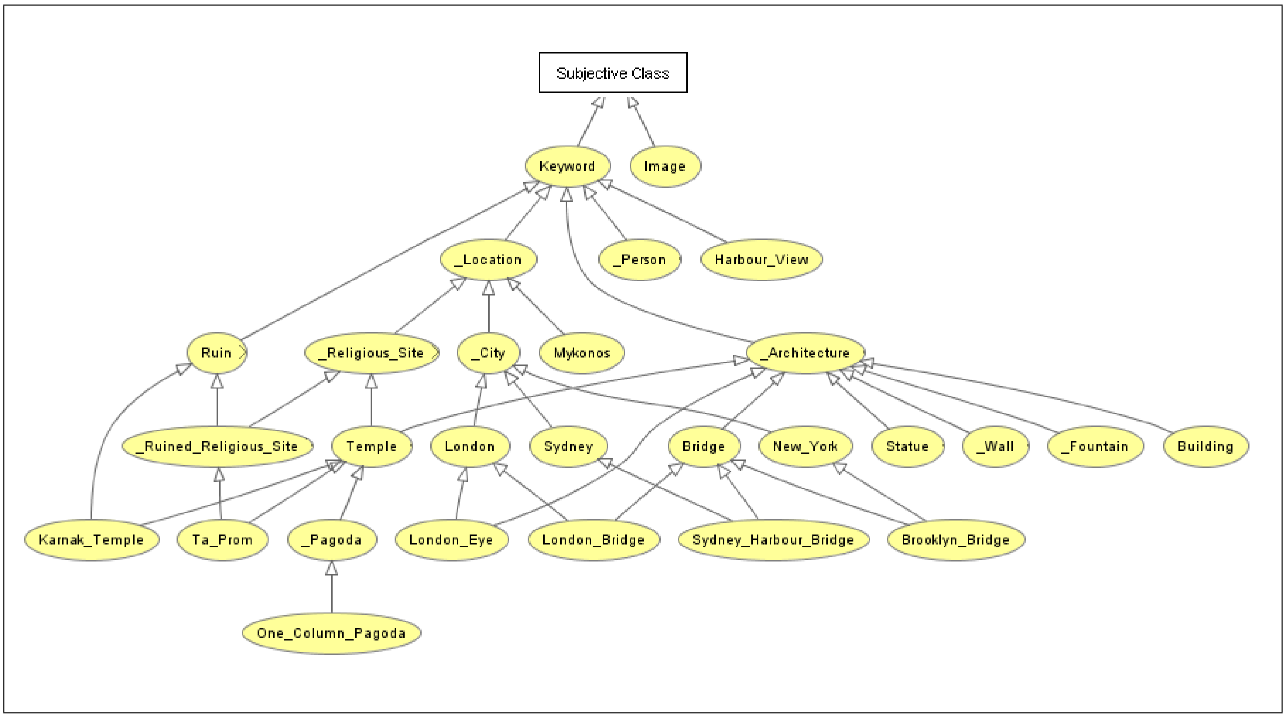

Figure 3: Section of example ontology created by Domain expert

collection of similar objects. This classification allows for the creation of a vocabulary to represent all tags that an author uses.

In this work we capture the importance of tags as they relate to both objects and authors to allow for the calculation of the probability that a tag is useful in another author's view of a domain.

Definition We define T_{ij} to be the set of tags $\{\alpha_1, \alpha_2, \ldots, \alpha_m\}$, applied to an object, O_j, as a description provided by an author A_i.

By combining all of an author's tags we capture the vocabulary of that author.

Definition A vocabulary, V_i of an author, A_i is the combination of all tags used by that author for all objects:
$V_i = T_{i1} \cup T_{i2} \cup \ldots \cup T_{ij} \cup \ldots \cup T_{ik}$

We assume that all objects within the environment have at least one tag, $\exists i \forall j T_{ij} \neq \{\varnothing\}$, thus a vocabulary can be used to access all of an authors categorised objects.

Folksonomies have difficulty when a term's meaning in another domain is different. By lacking the ability to apply a context to a tag there is no ability for an application to differentiate the content on the tags alone. Redundant tagging or use of synonyms creates an additional difficulty that must be overcome. An effective tagging system could require the integration of a thesaurus (such as WordNet) to ensure that all appropriate synonyms are returned as the result of a search query. Spelling errors and use of plural form also cause difficulty and need to be addressed by the framework.

The main difficulty focused on in this work is the use of incorrect tags (tagging a photo about cities with rural landscapes), or conflicting tags (tags that can be classified as disjoint). Another problem is the lack of attribute information. This would prevent querying a folksonomy for any results based on an attribute (e.g. cities with a population $\geq$ 3 million).

4 Social Networking

The SubjectiveAuthor concept in Figure 1 is a sub-class of SubjectiveResource and inherits the attributes and relationships. This allows an author to be rated by credibility, to be accepted or rejected based on the isAccepted property, and to provide other authors as resources within their ontology. Additionally we extended the RIPOSTE framework with the additional semantic relationships between authors labeled 'friend' and 'foe'. These additions allow for a greater ability to replicate existing social networks such as a *Web of Trust* (Guha 2004), *Friend of a Friend* (FOAF) (Dumbill 2002), or an existing social computing site, such as Digg, or Slashdot. Initially we assume that an author is categorised as a friend if there is a high level of agreement on a previous topic and a foe if actively sabotaging the current author on a previous topic of discussion.

4.1 Rating System

In a distributed environment, such as the Semantic Web, any author can supply information. It can be difficult to discover credible information and reliable authors. To assist in this discovery social modelling approaches (Resnick et al. 2000, Dumbill 2002, Guha 2004) rate users and their information providing additional feedback to help establish trust.

Numerous trust models have been proposed to rate or quantify the credibility of authors and the information they provide. In this work trust is defined as:

> Trust is the subjective probability by which an individual, A, expects that another individual, B, performs a given action on which its welfare depends. (Jøsang et al. 2005)

Trust is often calculated using social modelling to create networks that examine reputation (Resnick et al. 2000) captured in a directed graph with nodes (authors) and arcs (trust relationships) (Dumbill 2002). Additionally, when creating a social model, it is necessary to model trust and distrust between objects (ontologies, tags or instances), users, reviews, and ratings of reviews (Guha 2004). Each review can be ranked and allows for the propagation of trust changes. When a user's trust rank is altered, their reviews of other users will also alter rating within the model to reflect their modified reputation. This is often referred to as a *Web of Trust*(Guha 2004). In this work trust is represented as an opinion. That is if an author (A_1) is said to trust another author, or object, (x) then A_1 has an opinion $\omega_x^{A_1}$ about the other author.

Tag clouds, shown in Figure 5 and Figure 8 are used to show the popularity (Γ) of a tag. A tag cloud can be separated into time periods (Δ)(i.e. last 24 hrs or week) and is generated by using size to represent the number of authors ($|A|$) who use a tag (α). We calculate the popularity by the percentage of authors who use a given tag within a given time period:

$$\Gamma_\Delta(\alpha) = \Delta \frac{|A^\alpha|}{|A|} \; where \; \alpha \in V_i \tag{3}$$

where A^α is the set of authors that have $\alpha \in V$.

By introducing a time period, Δ, we are able to capture the evolutionary aspect of tags within the environment.

Content rating $\Gamma(\alpha)$ is based on user opinions and in its simplest form is a popularity measure, where users vote whether they like, or dislike, a submission and an aggregated value is applied to the content. This value is often used calculate popularity, which can provide a prominent position on the website (e.g. Digg).

We use the ratings of the tag cloud to provide a value for the expectation of a positive belief of the tag:

$$E(\alpha) = \frac{\Gamma(\alpha)}{max(\Gamma(\alpha))} \tag{4}$$

where $\max(\Gamma(\alpha))$ is the highest rated content. This can be combined with a_x calculated in equation 2, and an initial assumption that with the absence of evidence the opinions for all tags are completely uncertain (i.e. $u_x = 1$). By combining equations 1, 4 we calculate an initial belief b_x:

$$b_\alpha = \frac{\Gamma(\alpha)}{max(\Gamma(\alpha))} - a_\alpha \tag{5}$$

As shown in the tag cloud in figure 5 tag 'tower' has a popularity $\Gamma(tower) = 5$, while $\Gamma(building) = 10$,

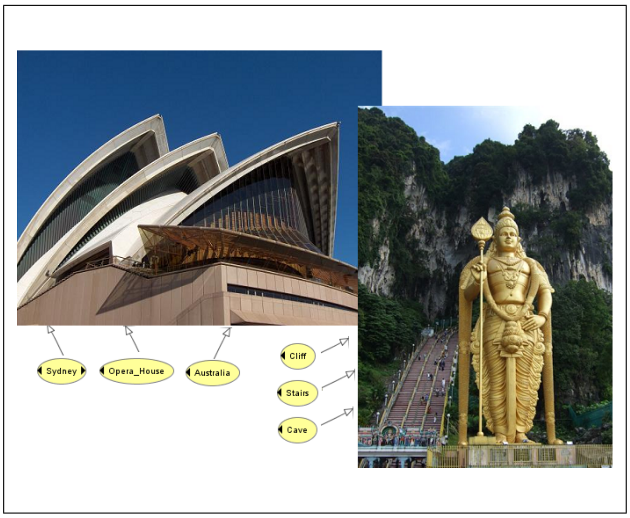

Figure 4: Example of user supplied content with user defined tags

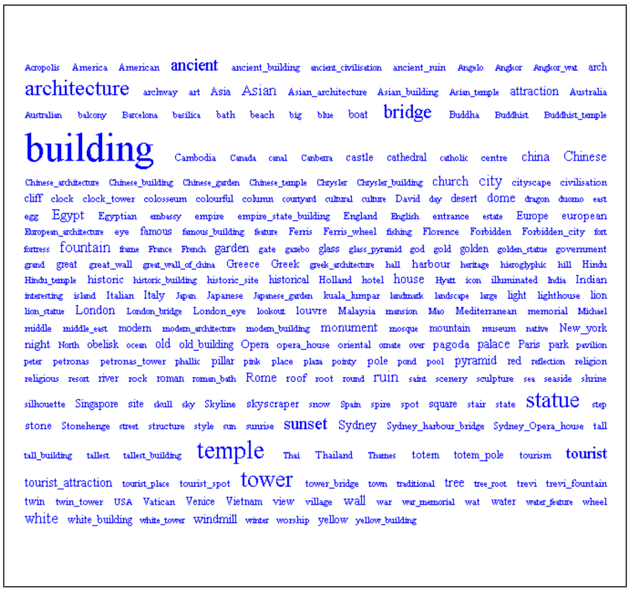

Figure 5: Resultant Tag-Cloud from study with a threshold of 5

while figure 6 shows 'tower' to have a $senseCount = 3$. Assuming 'tower' is unrelated to existing ontology classes we use equation 2 to calculate $a_{tower} = 0.25$. The initial belief for 'tower' is calculated by:

$$b_{tower} = \frac{5}{10} - 0.25 = 0.25$$

This becomes evidence reducing our uncertainty from 1 (initial uncertainty) to 0.75 and disbelief = 0.

4.1.1 Combining Ratings

One of the key issues of social computing is combining ratings. In traditional sites (Digg, Slashdot) a single positive value is provided for all content that friends rate as good, while all comments by foes are filtered out. To provide the ability to see both opinions it is necessary to combine ratings, using consensus ($\oplus$) and discount($\otimes$), from both friends (ϕ) and foes (ψ), of the deciding author (D), to reflect that a foe may be correct. To capture the opinions of friends we use:

$$\omega_x^{D\phi} = (\omega_{A_1}^D \otimes \omega_x^{A_1}) \oplus \cdots \oplus (\omega_{A_n}^D \otimes \omega_x^{A_n}) \forall A \in \phi$$

and foes:

$$\omega_x^{D\psi} = (\omega_{A_1}^D \otimes \omega_x^{A_1}) \oplus \cdots \oplus (\omega_{A_n}^D \otimes \omega_x^{A_n}) \forall A \in \psi$$

This is combined with a weighted opinion ω_ϕ^W and ω_ψ^W which allows for friends to be held in higher ($\omega_\phi^W > \omega_\psi^W$),or lower ($\omega_\phi^W < \omega_\psi^W$), regard than foes:

$$\omega_x^D = (\omega_\phi^W \otimes \omega_x^{D\phi}) \oplus (\omega_\psi^W \otimes \omega_x^{D\psi})$$

Given the 3 authors Allan, Bob and Carl introduced earlier, let Allan and Bob be friends ($Bob \in \phi_{allan}$) and Allan and Carl be foes ($Carl \in \psi_{Allan}$). Then using an arbitrary weight for friends and foes of:

$$\omega_\phi^W = (0.8, 0.2, 0, 0.5)$$
$$\omega_\psi^W = (0.2, 0.8, 0, 0.5)$$

showing that we believe our friends 80% of the time they provide us with a tag and disbelieve them 20%, while foes are only believed 20% and disbelieved 80%. Then for the keyword 'Australia' the opinion is calculated as:

$$\begin{aligned}\omega_{Aus}^{Allan} &= (0.8, 0.2, 0, 0.5) \otimes (0.66, 0, 0.34, 0.17)\\ &\oplus (0.2, 0.8, 0, 0.5) \otimes (0.5, 0, 0.5, 0.17)\\ &= (0.55, 0, 0.45, 0.17)\end{aligned}$$

5 Integrating Folksonomies and Ontologies

The main advantage of integrating a folksonomy is to utilise the expressiveness of the ontology to improve querying (Angeletou et al. 2007) (i.e. querying for photos of cities should return images tagged as being photographs of Barcelona). If we query the vocabulary of Figure 4 for images of cities in China we would get no result. Once a mapping to the ontology, shown in Figure 3 is performed, the resultant images from User1 will be returned. To integrate folksonomy to ontology it is necessary to convert them to the same format. Since this work uses OWL DL, which is a richer language than the meta-tag folksonomy, the folksonomy is converted to OWL DL without loss by encoding the folksonomy.

5.1 Encoding Folksonomies

A folksonomy can be defined as an ontology that has no semantic relations, $S = \{\varnothing\}$, and no axioms $\chi = \{\varnothing\}$. This allows the folksonomy to be encoded in OWL DL as an ontology, O', by creating a concept, C_α, $\forall \alpha \in V_i$, such that $C_{alpha} \sqsubset owlThing$. Finally each object, O_j, described by the vocabulary, V_i should be encoded as an individual, Instance(O_j, C_{alpha}), $\forall \alpha \in T_ij$.

This conversion allows for the integration of the folksonomy with an ontology. The difficulty for integration is that the folksonomy has no semantic relations allowing for no graph matching (Noy & Musen 2000). This limits the integration algorithm to pattern matching on the labels (Noy & Musen 2000, McGuinness et al. 2000) or by examining external sources (Sabou et al. 2006).

Examining existing ontologies for semantics between concepts can assist in discovering mappings (Sabou et al. 2006). The Semantic Web provides a semantically rich environment that can be used as background knowledge. By utilising existing ontologies that contain a concept from each of the ontologies to be merged, it is possible to derive a mapping. To deal with contradictions provided by multiple background ontologies that contain the two concepts any multiple results from a Semantic Web search engine (such as Swoogle[8]) need to be compared to derive the best mapping. Currently this comparison is achieved by ranking the results by the authors reliability if it exists or by selecting the ontology with the shortest path between the concepts. An additional necessity of integration is the need to add an additional resource to describe semantic relationships between two concepts.

5.2 Integration Process

Our integration process involves finding similarities between tags in a folksonomy vocabulary and classes in an ontology. The integration is subjective based on the opinions of the user performing the operation.

1. If $A_i \notin O$ we add A_i to O and then examine syntactic equivalence for each $C \in O$ and for each $\alpha \in V_i$:

2. If $syntax(\alpha) = syntax(C)$ the two terms are possibly equivalent. Then the opinion is calculated using equation 5 where:
$a_x = \dfrac{1}{senseCount(C) + 1}$ where senseCount is calculated as the number of possible senses returned from a WordNet query of the term and example is shown in Figure 6

3. Else if $syntax(\alpha) = synonym(syntax(C))$ the two terms are possibly equivalent (e.g. city and metropolis shown in Figure 6). Add α as an individual of SubjectiveClass and create an individual of PossibleEquivalence linked with both α and C. The opinion is calculated using equation 5 where:
$a_x = \dfrac{simSense(C, \alpha)}{senseCount(C) + senseCount(\alpha) + 1}$
where simSense pertains to the number of senses in which both words occur (if they are synonyms this will be $simSense(C, \alpha) \geq 2$).

4. Else if $syntax(\alpha) \in syntax(C)$ the two terms are possibly in super-class relationship where the approximation is based on syntactic closeness (McGuinness et al. 2000)(e.g. $\alpha =$ "'tower'" and $C =$ "drum tower"). Add α as an individual of SubjectiveClass and create an individual of PossibleSubClass linked with both α and C where C as the superclass. Then the opinion is calculated using equation 5 where:
$a_x = \dfrac{1}{maxSense(\alpha, C) + 1}$ where maxSense is the maximum number of senses of either α or C.

5. Else if $syntax(C) \in syntax(\alpha)$ the two terms are possibly in sub-class relationship. Add α as an individual of SubjectiveClass and create an individual of PossibleSubClass linked with both α and C where α is the superclass. Then the opinion is calculated using equation 5 where:
$a_x = \dfrac{1}{maxSense(\alpha, C) + 1}$ where maxSense is the maximum number of senses of either α or C.

6. Else search for an existing ontology, O_n (using Swoogle (Sabou et al. 2006)) that contains both $\alpha, C \in O_n$ and import the minimum set of resources that semantically relate the terms, $O' \subseteq O_n$ where
$S(C', C'_1)$,..., $S(C'_{j-1}, C'_j S(C'_j, \alpha')$ and
$syntax(C') = syntax(C)$ and $syntax(\alpha') = syntax(\alpha)$ and $\forall C''$ from C'_1 to $C'_{j-1} syntax(C'')$ $\neq syntax(C_k \in O)$ or $syntax(\alpha_k \in V_i$
   (a) if $C'' \in O$ then remove all $S(C', C'_1)$ ..., $S(C'_n, C'')$ from O' as this will be the shortest path between the compared concepts in O and V_i.

This process will introduce additional authors and require the integration process to be performed between O and O' as above then add each R $\in O'$ to O.

If no existing ontology is found then α is added to O' as a class C_alpha such that $subclass(C_\alpha, owlThing)$ and $providedBy(C_\alpha, A_i)$ and the opinion is calculated using equation 5 and $a_x = \frac{1}{senseCount(C)+1}$

[8] http://swoogle.umbc.edu/

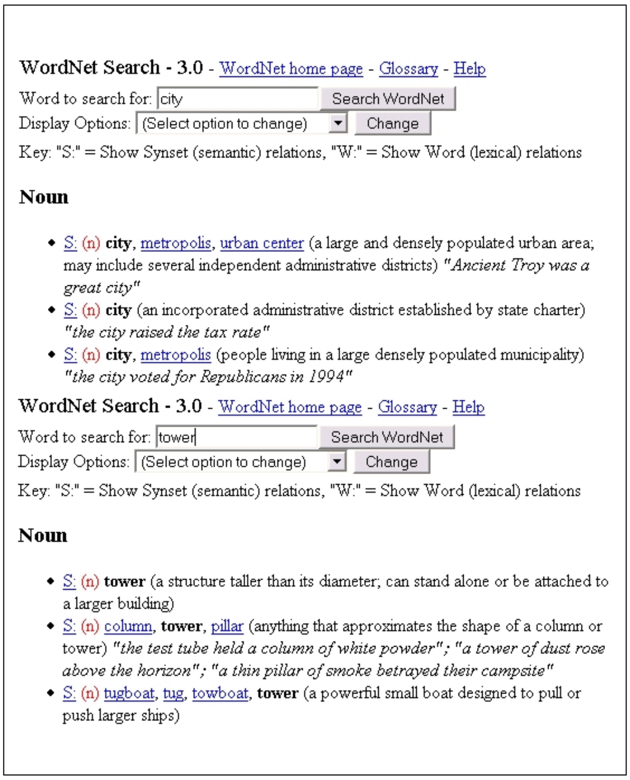

Figure 6: Results of WordNet queries for 'city' and 'tower'

6 Evaluation

To evaluate the RIPOSTE Framework, an ontology (O_{A_0}) was created using Protégé [9] to capture a domain of landmarks for photography. This included locations and types of landmarks. The ontology, partially shown in Figure 3, used simplified classification and focused on taxonomy of classes and meronyms, providing a simplified ontology for classifying photos. Additionally the domain was reduced to only include concepts that related to a set of specific photographs. Additional concepts were included for clarification where necessary (preceded by "-").

A set of 60 photographs (Y), taken of famous landmarks in various locations around the world were suplied to an expert (A_0). The expert then annotated using O_{A_0}. This set of annotations was used as a benchmark for tag correctness.

The photographs were divided into two subsets of 40 photos (Y_1 and Y_2), containing an intersection of 20 photos ($|Y_1 \cap Y_2| = 20$) were supplied to a test group of 50 additional authors ($\{A_1 ... A_{50}\}$) to be annotated with at most 3 tags ($|T^{A_n}_{I_m}| \leq 3$). If a tag contained multiple words a new tag was added for each word of length > 2 and all were added as a superClass for the multi-worded tag.

Expertise of the additional providers were calculated by combing results from 1) a general question based on the previous experience with social tagging applications, and 2) a specific question for each image requiring the provider to state whether they recognised the image.

This created a set of Vocabularies $\{V_{A_1}, \ldots, V_{A_{50}}\}$ with the vocabulary sizes shown in Figure 7. These vocabularies were passed into WordNet to discover the synonyms and hypernyms shown in Figure 7.

[9] http://protege.stanford.edu/

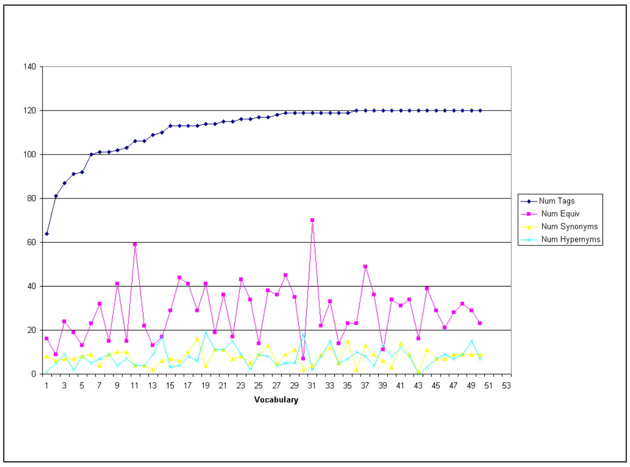

Figure 7: Vocabulary of the people involved in the evaluation

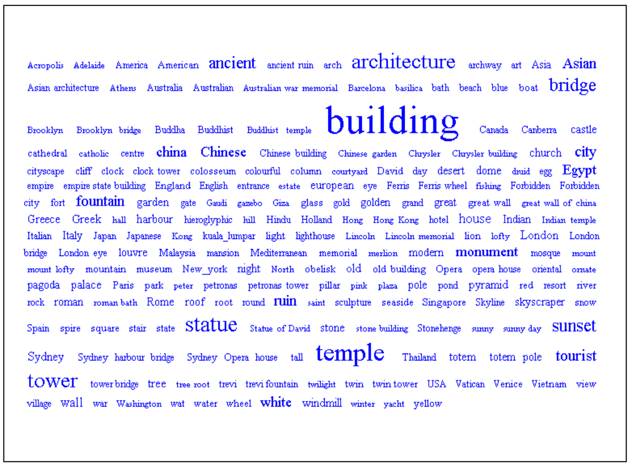

Figure 8: Resultant Tag Cloud with a belief value above 50%

6.1 Comparison of folksonomies

The vocabulary was filtered into two subsets for comparison. The first, F_1 was based on the number of occurrences of a tag within the vocabulary while the second, F_2 took into account the expertise of the user (0.5 or 1.0), based on their knowledge of tagging systems combined with their knowledge of the image. The image knowledge was calculated from whether they recognised the image. A threshold was applied to improve precision of the tags. F_1 was filtered by a threshold value of 5 on tag count, Figure 5, while a threshold value of 0.5 based on the subjective belief of the tag (calculated using consensus and discount described in Section 2.1) was applied to F_2, Figure 8.

The result shows that although the approximate size of the clouds is similar the selection of tags is not. It can be seen that generic terms, such as 'building', still occur in both filters. An advantage of F_2 is that for images that are not as well recognised the tags created by the experts have a greater weighting. This provides the resultant folksonomy with more specialised tags (i.e. the actual names of the landmark "Statue of David") as opposed to more generic tags ("Statue").

6.2 Ontology Creation

The tags in F_2 were then converted into OWL classes with no semantics or axioms except that each class was a subclass of SubjectiveClass to create an ontology O_V. Then each hypernym/hyponym relation was

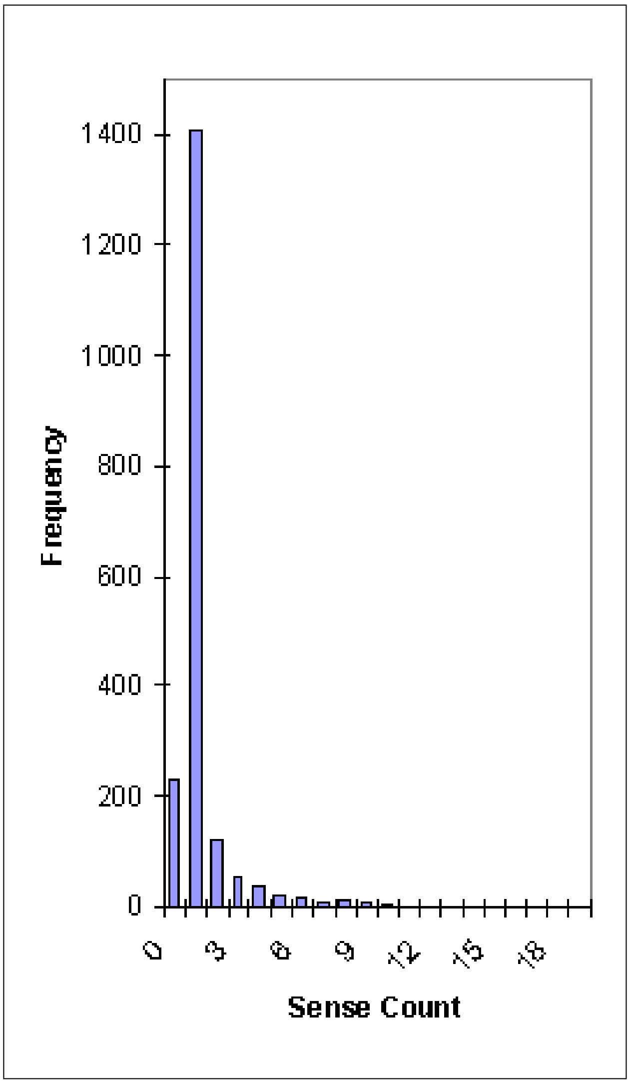

Figure 9: Tag sense count

added such that $hypernym(\alpha_1, \alpha_2) = \alpha_1 \sqsubseteq \alpha_2$ and each synonym relation was added as an equivalent relation. Additionally each author was included as a SubjectiveAuthor and a relation between each author and the tags in their vocabulary was created.

The sense count, also taken from WordNet, was applied to all tags in the folksonomy, shown in Figure 9, and all classes in the domain ontology. Finally each image was added as a SubjectiveIndividual with relationships between the authors and the tags that they were annotated with. The syntax of each class in O_V was then compared with the syntax of each class in O_{A_0} to discover equivalence, synonyms and hypernyms, Figure 11. Of the 199 ontology classes, 145 classes existed in the combined vocabulary. There were 333 synonyms and 89 hypernyms. The ontology is improved by the inclusion of additional terms allowing for greater use of synonyms and hypernyms supplied by the folksonomy to improve querying.

7 Related Work

Most categorisation for social computing is currently achieved by a meta-tagging system dubbed folksonomies, which fails to take into consideration the semantics that occur between terms. To overcome this a folksonomy can be integrated with an ontology allowing for greater reasoning to refine searches.

Current rating systems provide a popularity, or karma, measure to the content provider which is then used to rate the content submitted. This is a simplified mechanism that fails to take into account the expertise of the author. By integrating the folksonomy with an existing subjective ontology a rating of the content domain can be supplied. This can be compared with an authors previous submissions. This allows for examination of an author's consistency within a selected sub-ontology, allowing for a more fine grained approach to ratings ensures an expert in one field is not as highly regarded in a dissimilar field.

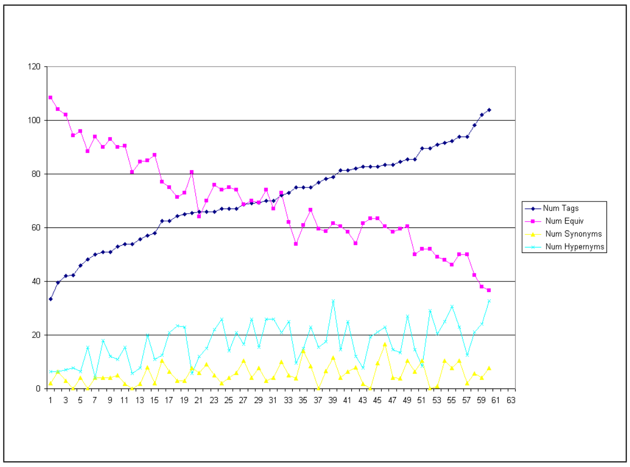

Figure 10: Vocabulary for each image

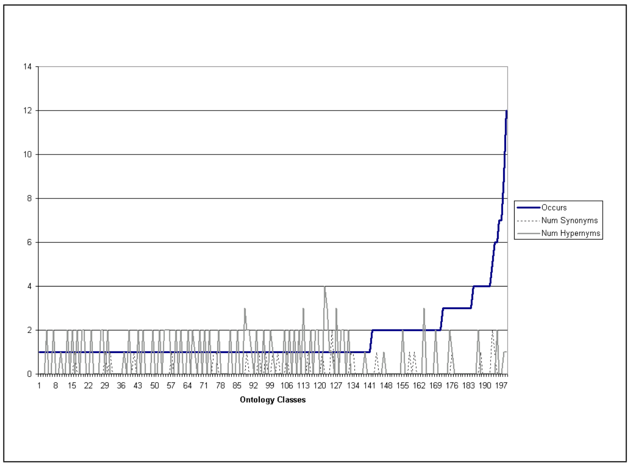

Figure 11: Vocabulary comparison with Ontology

(Schmitz 2006) recommends using a probabilistic model to induce a faceted ontology from Flickr tags. By examining the number of images and the number of users that a tag belongs to a threshold is applied to remove noise. This removes noise but also minimises recall (Schmitz 2006). The thresholded set of tags is then used to discover subsumption pairs, which are manually examined for correctness and synonyms (including multilingual synonyms). It can be seen in our study that the reduction in recall of images will effect those images with more concise, or concrete, tags. This will result in many leaf nodes not being discovered reducing the ability to refine searches within the subsumption hierarchy.

The need to merge Semantic Web technologies with Web 2.0 is necessary to reduce the individual weaknesses of each approach (Heath & Motta 2007, Gruber 2007, Ankolekar et al. 2007). The AI centric nature of the Semantic Web makes it difficult for users while the user centric nature of Web 2.0 provides limited ability for AI assisted reasoning, model checking and reuse. Where as some researchers advocate replacing Web 2.0 with RDF publishing and SPARQL querying (Heath & Motta 2007) others recommend linking solutions to combine Web 2.0 and the Semantic Web (Bojārs et al. 2007, Ankolekar et al. 2007). By providing a protocol to link the tagging systems to an ontology the expressiveness can be utilised for improved querying. It is seen that a threshold is necessary to remove noise (Schmitz 2006) and the preliminary work in linking approaches also shows this. Unfortunately these approaches do not take into account the pedigree of the authors of tags which results in any threshold removing many specific tags.

Providing an ontology combined with a social site (Gruber 2006, Heath & Motta 2007) allows for additional querying that will allow photographs of 'Sydney' to be found in a search about 'Australia'. RealTravel[10], an implementation of (Gruber 2006) provides a static approach to the ontology, which does not take into consideration that as additional tags are added to the site the ontology may need to grow to provide the same reasoning abilities to the newly expanded domain.

By capturing tags in RDF (Heath & Motta 2007) the ontology becomes dynamic allowing for greater flexibility this relies on new interfaces to be created to simplify ontology evolution. By having no mechanisms in place to control the insertion of the new RDF triples it can be easy for users to provide incorrect or misleading information, which may result in a corrupted and unusable ontology.

8 Conclusion

A novel approach to integrating folksonomies to ontologies is discussed. By appending probabilities to all resources within an ontology the ability to estimate the usefulness of a result provides a mechanism for ranking based on this usefulness.

Additionally by incorporating authors with the specific resources they provide a fine grained trust calculation can be used and allows for an author's area of expertise to be taken into account when accepting results. Currently the ability to calculate an area of expertise allows for ranking of results based on the credibility and reliability of the author. Our approach differs from other social network research by considering the semantic relations between provided content. While our approach extends the capabilities of ontology integration research by allowing incomplete or incorrect ontologies to be augmented with probability. This provides a user the ability to see where more credible sources are required to improve the semantic representation of a particular sub-domain within the ontology.

8.1 Future Work

Currently we only provide mechanisms for capturing a relationship between one tag and one class. We aim to examine the possibility of capturing more complex relationships (i.e. multiple tags combined to form the ObjectProperties of multiple classes). Currently we have only evaluated tags that relate to classes and would like to examine tags that will match to semantic relationships (i.e. the tag is a verb).

Additionally we would like to evaluate the Area of Expertise to examine the possible time reduction due to less calculations and how the generalised calculation of trust will effect the precision of result rankings.

[10]http://realtravel.com

References

Al-Khalifa, H. S. & Davis, H. C. (2006), Measuring the semantic value of folksonomies, *in* 'Proceedings IEEE Conf Innovation in IT'.

Angeletou, S., Sabou, M., L.Specia & Motta, E. (2007), Bridging the gap between folksonomies and the semantic web: An experience report, *in* 'Proceedings The International Workshop Bridging the Gap between Semantic Web and Web 2.0'.

Ankolekar, A., Krötzsch, M., Tran, T. & Vrandečić (2007), 'The two cultures: Mashing up web 2.0 and the semantic web', *Journal of Web Semantics* .

Bojārs, U., Breslin, J., Finn, A. & Decker, S. (2007), 'Using the semantic web for linking and reusing data across web 2.0 communities', *Journal of Web Semantics* .

Dumbill, E. (2002), 'XML watch: Finding friends with XML and RDF'.
URL: *http://www.ibm.com/developerworks/x-foaf.html*

Gruber, T. (2006), Where the social web meets the semantic web, *in* 'Keynote at 5th International Semantic Web Conference', Vol. 4273 of *LNCS*, Springer, Athens, GA, USA.

Gruber, T. (2007), 'Collective knowledge systems: Where the social web meets the semantic web', *Journal of Web Semantics* .

Guarino, N. (1997), Semantic matching: Formal ontological distinctions for information organization, extraction, and integration., *in* 'SCIE', Vol. 1299 of *LNCS*, Springer, Frascati, Italy, pp. 139–170.

Guha, R. V. (2004), Open rating systems, *in* '1st Workshop on Friend of a Friend, Social Networking and the Semantic Web', Galway, Ireland.

Heath, T. & Motta, E. (2007), 'Ease of interaction plus ease of integration: Combing web2.0 and the semantic web in a reveiwing site', *Journal of Web Semantics* .

Hooijmaijers, D. & Stumptner, M. (2006), Trust based ontology integration for the community services sector, *in* 'Advances in ontologies. Proceedings of AOW', Hobart, Australia.

Hooijmaijers, D. & Stumptner, M. (2008), Improving integration with subjective combining of ontology mappings, *in* A. An, S. Matwin, Z. W. Ras & D. Slezak, eds, 'ISMIS', Vol. 4994 of *Lecture Notes in Computer Science*, Springer, pp. 552–562.

Jøsang, A. (2002), A logic for uncertain probabilities, *in* 'Int. Journal of Uncertainty, Fuzziness and Knowledge-Based Systems', Vol. 9.

Jøsang, A., Ismail, R. & Boyd, C. (2005), A survey of trust and reputation systems for online service provision, *in* 'Decision Support Systems', Vol. 43 of *LNCS*, Science Direct, pp. 618 – 644.

McGuinness, D. L., Fikes, R., Rice, J. & Wilder, S. (2000), The Chimaera ontology environment., *in* 'Proceedings of AAAI', Austin, Texas, USA., pp. 1123–1124.

Noy, N. F. & Musen, M. A. (2000), PROMPT: algorithm and tool for automated ontology merging and alignment., *in* 'Proceedings of AAAI', Austin, Texas, USA., pp. 450–455.

Resnick, P., Kuwabara, K., Zeckhauser, R. & Friedman, E. (2000), 'Reputation systems.', *Communications of the ACM* **43**(12), 45–48.

Sabou, M., d'Aquin, M. & Motta, E. (2006), Using the Semantic Web as background knowledge for ontology mapping, *in* 'Proceedings of the 1st Int. Workshop on Ontology Matching', Athens, Georgia, USA.

Schmitz, P. (2006), Inducing ontology from flickr tags, *in* 'Proceedings of The International Workshop Bridging the Gap between Semantic Web and Web 2.0', IW3C2, Edinburgh, UK.

Szomszor, M., Cattuto, C., Alani, H., O'hara, K., Baldassarri, A., Loreto, V. & Servedio, V. (2007), Folksonomies, the semantic web, and movie recommendation, *in* 'Proceedings of The International Workshop Bridging the Gap between Semantic Web and Web 2.0'.

van Damme, C., Hepp, M. & Siorpaes, K. (2007), Folksontology: An integrated approach for turning folksonomies into ontologies, *in* 'Proceedings of The International Workshop Bridging the Gap between Semantic Web and Web 2.0'.