

SWAD-Europe deliverable 11.1: Recommendation systems as a basis for trust on the Semantic Web

Project name:

W3C Semantic Web Advanced Development for Europe (SWAD-Europe)

Project Number:

IST-2001-34732

Workpackage name:

Distributed trust systems

Workpackage description:

<http://www.w3.org/2001/sw/Europe/plan/workpackages/live/esw-wp-11.html>

Deliverable title:

11.1: Recommendation systems as a basis for trust in the Semantic Web

URI:

<http://www.w3.org/2001/sw/Europe/reports/trust/issues.html>

Author:

Jan Grant, ILRT, University of Bristol

Abstract:

The Semantic Web promises the distribution of interchangeable machine-readable data across the web. In particular, technologies like FOAF are being used as the basis for many promising social networking applications. As those applications increase in perceived value, the need for a system for establishing and expressing trust will also increase. This document discusses the issues involved with such systems and the use of Dempster-Schafer opinion systems for evaluating trust metrics. A strawman vocabulary for expressing subjective logic opinions (and those expressed using other trust metrics) is given.

STATUS:

Completed, last modified 2004-07-15.

Contents

-
- 1 [Publishing assertions](#)
 - 1.1 [Importing identity from PGP](#)
 - 1.2 [Identity does not suffice](#)
 - 2 [Evaluation mechanisms](#)
 - 2.1 [Dempster-Schafer](#)
 - 2.2 [Problems](#)
 - 2.3 [Non-probabilistic metrics](#)
 - 3 [Subject contextualisation](#)
 - 3.1 [Contextualisation with subjective logic](#)
 - 3.2 [Requirements contextualisation beings to vocabulary mapping systems](#)
 - 4 [Publishing recommendations and opinions](#)
 - 4.1 [Describing a recommendation](#)
 - 4.2 [PICS](#)
 - 5 [Issues with trust in social networks](#)
 - 5.1 [Cliques](#)
 - 5.2 [Libel](#)
 - 5.3 [Peer pressure](#)
 - 5.4 [Making objective assessments](#)
 - 6 [Applications of recommendation systems](#)
 - 7 [Conclusion](#)
 - A [References](#)
 - B [Summary of operators in subjective logic](#)
-

1 Publishing assertions

The Semantic Web is founded upon a small number of machine-processable metalanguages with well-defined semantics [\[RDF-SEMANTICS\]](#), [\[OWL-SEMANTICS\]](#).

Rather than inventing new syntaxes on a per-application basis, a user publishes assertions in terms of application-specific (or general application-neutral) vocabularies specified in terms of the semantics of the underlying metalanguages. Providing the use of these vocabularies is consistent with the semantics of the underlying metalanguages, a level of semantic interoperability can therefore be obtained. Furthermore, the data so published may be used and reused in ways not originally envisaged by the original publisher.


Although the envisioned scope of the Semantic Web goes far beyond simple assertions, there has been a great deal of interest in applications that are effectively based around the publishing of simple assertions. One such application is FOAF [\[FOAF\]](#), a vocabulary for describing social networks. At a "grass roots" level, there is a great deal of community interest in novel applications to which FOAF and extensions thereof may be put. Furthermore, FOAF exemplifies the distributed interoperability promised by the Semantic Web. Much of this document, therefore, adopts FOAF and similar social networking technologies as motivational examples.

FOAF assertions are normally expressed using RDF: FOAF usage today typically relies on the RDF schema semantics plus a small subset of OWL (functional and inverse functional properties). In this document, the term "RDF" is used informally to indicate the semantics of the family of Semantic Web languages built upon RDF.

1.1 Importing identity from PGP -

Since FOAF is used to make assertions about (amongst other things) people and the relationships between them, it requires a mechanism to identify the individuals that assertions concern. The method adopted by FOAF is to identify a person via association with a unique address (a mailbox). (In fact, cryptographic hashes of mailbox addresses are normally used, but the principle remains the same.)

Of course, a person may be associated with multiple mailboxes, and FOAF permits the "unification" of the individual associated with multiple mailbox addresses via a simple assertion.



The FOAF model is that an individual can make any number of assertions, both about themselves and about others. These assertions are typically published on the web, where they may be aggregated and processed by FOAF-aware software agents.

Although FOAF has a notion of identity, and can be used to make claims about document authorship, what it does not directly support is a mechanism for *assuring* the authorship of a particular published set of RDF assertions. Instead, typical FOAF use can import authorship assurance from other trust networks, such as PGP [\[FOAF-SIGN\]](#).

PGP identifies individuals via association with an email address (or addresses), so the principal identifiers for individuals in both PGP and FOAF map well onto each other. As an additional point in its favour, PGP has a distributed model (in the case of key servers one might say "decentralised") which fits nicely with the FOAF and Semantic Web views. One disparity between the FOAF and PGP approaches to identity is that FOAF uses cryptographic hashes of mailbox addresses as identifiers, in an attempt to obscure valid email addresses from web crawlers. PGP exposes these addresses. (The author's personal opinion is that attempting to obscure email addresses in this fashion (in an attempt to combat spam) is misguided.)


A FOAF application can therefore "import" trust from a system like PGP in order to assure document integrity.

1.2 Identity does not suffice -

However, in the context of a social networking tool, there is a great deal more to *trust* than establishing identity. While the *author* of a particular document might be established to a high degree of reliability, that does not directly imply that the *semantic content* of that document is reliable.

This is where recommendation networks come into the picture. A social network can be viewed as a set of agents. Those agents may have opinions of the reliability of assertions made by other agents. A particular agent (our "user", A) may have direct opinions about other users, B and C, and can therefore make direct judgements about the reliability of assertions made by those agents.


However, A may have no direct opinion of the agent D. If A wishes to form an opinion of D's assertions, he must look for a chain of evidence (the published opinions of other agents) that links A to D. If such a chain exists, A can attempt to evaluate the chain of opinion to form a derived opinion of D, and armed with this make his best value judgement as to the level of trust he places in the assertions of D.




2 Evaluation mechanisms

By and large, recommendation systems all use the same basic operations for the evaluation of opinion networks; although the exact detail of these operations may vary.

For combining the opinions of multiple agents, the basic operations are *consensus* and *discounting* (these terms are taken from Dempster-Schafer theory, below). The *consensus* operation combines the opinions of two agents on a single proposition. The *discounting* operation is used where a chain of opinions need to be evaluated. Agent A can *discount* B's opinion of D (the first proposition) by A's opinion of B (the second proposition) to produce a discounted opinion.



Once A has produced the two discounted opinions, they can be combined using the consensus operator to derive a final opinion of D.



In the example above, A finally forms an opinion thus:
 $o(A,D) = \text{consensus}(\text{discount}(o(A,B), o(B,D)), \text{discount}(o(A,C), o(C,D)))$.

Where an agent has opinions of several propositions, the compound opinion of logical combinations of those propositions (conjunction, disjunction, conditional relationship, etc.) may also be formed.

2.1 Dempster-Schafer Theory -

A popular theoretical framework for opinion networks is based around Dempster-Schafer theory [\[IDS-INTRO\]](#). This framework is called *subjective logic*.

In this theory, the opinion of an agent on a particular proposition in DS theory is described by a 4-tuple (b,d,u,a). Here *b* represents the agent's *belief* that the proposition is true, and is a probabilistic measure. Similarly, *d* represents the *disbelief* in a proposition. The value of *u* represents the agent's *uncertainty* about the proposition, which is constrained by the relationship $b + d + u = 1$. Finally, *a* represents the *relative atomicity* of the agent's belief in the proposition [\[IDS-INTRO\]](#).

It is important that an agent's opinion is formed about a the truth of a specific proposition. Example propositions might be, "agent C is telling the truth", "C is an expert in aromatherapy", etc. The consensus operation in DS theory only operates to combine opinions of the same proposition. The consensus operator is both commutative and associative.

Chains of related opinions may be combined via the *discounting* operation. For example:

- Agent B publishes an opinion of the proposition, "Agent C is an expert in aromatherapy".
- Agent A has an opinion of the proposition, "Agent B has reliably judged C's expertise in aromatherapy (and is telling the truth)".
- Then Agent A can *discount* B's stated opinion of C against his own opinion of B, to derive an opinion of the proposition, "Agent C is an expert in aromatherapy".

Note that the discounting operation is associative, but not commutative.

An online demonstrator permitting a user to experiment with the consensus and discounting operators is available at [\[IDS-DEMO\]](#).

It is important that the exact proposition that an opinion is being offered on is understood for DS theory to be applicable.

Let us return to FOAF as a motivating example. Many of the FOAF use-cases involve locating individuals based upon shared interests or expertise. Individuals may wish to make assertions about the expertise of others in particular fields; however, such recommendations are not by necessity constrained to an academic field. For instance, a user may discover a published opinion about a film or a restaurant; in attempting to evaluate that opinion, a user is interested in the critical ability or expertise of the author of that opinion, and can seek out published opinions on the expertise of the author.

For simplicity's sake, we adopt for a simple template of a proposition that FOAF users may use to make recommendations about other FOAF users: that a particular individual "is an expert in subject *S*". Here *expertise* is used to mean that an expert, E, in a subject, S:

- has a thorough knowledge of the subject;
- publishes interesting, truthful, accurate documents on that subject;

- is a good judge of subject expertise in others.

The last condition is effectively a recursive definition of expertise. Whether this is particularly realistic and what implications it has are discussed below.


2.2 Problems -

When attempting to evaluate an opinion network, D-S Theory (and opinion theories in general) have to deal with several problems.

Multiple identities and double voting


The first issue concerns the consensus operation. When an agent, A, forms an opinion from the consensus of several input opinions, it is important that those inputs are independent. The simplest way that this condition can be broken is when A takes the consensus of several opinions, more than one of which have the same originator. That is to say, that $\text{consensus}(\text{opinion}(A,B), \text{opinion}(A,B))$ does not generally equal $\text{opinion}(A,B)$: counting the "vote" of a particular individual more than once biases the derived opinion.

As the value of manipulating a trust network increases, the incentive to do so also rises. (An example of a manipulated network for profit exists in the "link trading" networks used in an attempt to generate search engine "rank". Such networks are effectively gaming the trust network used by the search engine in order to highly-place particular web sites, either for profit or as a joke.) Anecdotal reports of the prevalence of "identity theft" are rife; it is a simpler matter to generate the facsimile of an independent "online identity" - which may be established under conditions where physical presence is not required - than to generate a separate fake identity in the physical world. Consequently, it is expected that multiple identities associated with an individual will be seen as the impact of social networking tools increases.



In the diagram above we see an unsuccessful attempt by Agent E to subvert A's opinion of D via multiple voting. However, since A will discount all the opinions of E's alter-egos by his opinion of E, the attempt will fail.

In the next diagram, however, E has successfully established different identities with B and C. A is therefore unwittingly fooled into counting E's opinion twice. In the example below, E's planted opinions are still discounted against A's opinion of B and C.



FOAF has a mechanism for asserting that the individuals associated with two mailboxes are in fact one and the same. An agent must choose whether or not to 'believe' such an assertion, when it finds one. It would be unwise to blindly accept such assertions since a denial-of-service attack against a FOAF network (as evaluated by a naive FOAF application) might consist of publishing false assertions that many mailboxes correspond to the same individual.


(An acceptable purely technical solution to the problem posed by the establishment of multiple online "identities" by an individual - one that preserves legitimate benefits to the user at the same time - is unlikely to exist. Since society at large is facing the same issue we might defer this matter until a solution to the problem at large is found. It is possible that an increase in the amount of machine-processable information concerning a particular individual that is normally available in the Semantic Web is likely to raise the cost of establishing multiple convincing online identities as time goes on: producing a convincing "fake" may become harder.)

Dogmatic opinions

The second problem that occurs is when dogmatic opinions are encountered. These are opinions that have zero uncertainty: that is, no admission of flexibility in opinion. Dogmatic opinions cannot be combined via the consensus operator. A FOAF tool for evaluating trust networks must have some tactic for dealing with dogmatic opinions. Options are presented in [\[DS-DOGMATISM\]](#).

Cycles of recommendations

There is an unfortunate consequence of adopting the third clause in the definition of *expertise* as the basis for propositions that opinions might be proffered on ("E is a good judge of subject expertise in others"). The universal quantification in this clause means that a network of published opinions may well contain cycles.



The possibility of such cycles does not occur when opinions are instead offered on explicitly identified propositions, for example:

- *E* has an accurate opinion of the expertise of *G*;
- *E* has a poor opinion of the expertise of *H*

However, by requiring such explicit opinions of opinions, the number of assertions required in order to bootstrap a recommendation network is vastly increased.

The problem of feedback loops in belief networks is normally tackled by adopting a heuristic approach to finding a cutset of the network that removes cycles. Since any social network is likely to contain pairs of mutually recommending agents, any realistic approach to weighting opinions must have some mechanic (even if it is a naive approach that can only generate approximate evaluations) for dealing with this situation.

Facets of expertise

When considering real-world trust networks, it becomes clear that a full picture of trust, even subject-contextualised, may divide trust into several facets. For instance, we may consider someone to have a high degree of knowledge in a subject area, but to draw conclusions that we consider to be erroneous.

For the sake of simplicity, this nuance is ignored in the treatment here. Adopting a pragmatic position, reputation built from evidence (ratings of actual assertions/outputs of a particular agent) used to form a single opinion of "expertise" will likely suffice for many applications.

2.3 Non-probabilistic metrics -

Although we do not explore them further here, not all opinion systems are based upon strictly probabilistic measures. Other approaches to voting and evaluating recommendations can certainly be formulated. However, they will almost certainly provide equivalents to familiar operations (consensus, discounting) as well as offering an ordering operation on the opinion metric.

A trivial example might be a trust metric based around a fully-ordered set of trust levels. *Consensus* between opinions might be implemented as the *maximum* of two opinions. *Discounting* might similarly be implemented by taking the minimum trust level of the chained opinions.

3 Subject contextualisation

The typical use-case for a FOAF application is for a user to ask:

find me experts in subject S (ranked according to the recommendations of those people I trust, those they consider expert, and so on)

Typically in a Semantic Web context the *subject* that contextualises this question is considered to be a resource. It may well be selected from a formal vocabulary, ontology, or thesaurus. It will almost certainly be named (in the Semantic Web manner) with a URI reference.

If all users publishing their recommendations were to draw their subject classifications from the same vocabulary, then the question of subject contextualisation would be relatively simple. However, it is not typically the case that this happens. For various historical, political, technical or even just accidental reasons, subject identifiers across the Semantic Web will be drawn from many vocabularies. In fact, recognising this is a key feature of the Semantic Web. Rather than expend intellectual and political effort to produce a single, unified and blessed preferred vocabulary, effort instead is directed towards technology that permits terms to be mapped from one vocabulary to another [\[SKOS-Mapping\]](#).

Consequently, it is possible that a FOAF agent, acting on behalf of a user, who is looking for opinions on expertise in subject *S*, finds a published opinion on the expertise in a subject *S'*. Where no relationship between the two identified subjects is known, the agent can make no inference (and can discard the published opinion). However, if there is some measure of "semantic overlap" between the two terms, then the agent may instead be able to give some weight to the published opinion as it tries to answer the user's original question.

3.1 Contextualisation with subjective logic -

In terms of subjective logic, we seek an opinion on one proposition, *P1* ("Agent *E* is an expert in subject *S*"), when we are given an opinion on a second proposition, *P2* ("Agent *E* is an expert in subject *S'*"). If the relationship between *P1* and *P2* is known: that is, we have opinions on (*P1* | *P2*) "*P1* is true given that *P2* is true" and (*P1* | ~*P2*) "*P1* is true given that *P2* is false", then [\[IDS-CONDITIONAL\]](#) gives a formula for calculating *P1*.

We can approximate opinions for $(P1 \mid P2)$ and $(P1 \mid \sim P2)$ if we have a measure of the *semantic overlap* of the two concepts, S and S' . For example, a thesaurus mapping service may use probabilities derived from the analysis of subject-classified documents as follows:

For documents classified as belonging to concept S' , a proportion, p of them are also classified as belonging to S . Then in the case that S and S' are related terms with some overlap, we approximate the opinion of $(P1 \mid P2)$ as $(p, 1-p, 0, 0.5)$ and the opinion of $(P1 \mid \sim P2)$ as $(0, 0, 1, 0.5)$ - that is, completely unknown. Similar calculations can be done when S and S' are related by the *broader term* or *narrower term* relationship.

3.2 Requirements contextualisation beings to vocabulary mapping systems -

There are some requirements that the desire to cope with subject contextualisation brings to the consideration of the design of vocabulary languages (eg, [\[SKOS-Mapping\]](#)) and vocabulary mapping services.

For this vocabulary mapping to be possible, it must be possible to locate a service which can map terms between vocabularies.

Assertions of opinion may be contextualised by a term taken from a thesaurus or other formal vocabulary. It should be possible to locate the originating vocabulary, given the term. There are various alternative approaches to this that might be taken. Possibly the simplest is to utilise the *rdfs:isDefinedBy* relationship to link the term to its containing vocabulary, as is suggested in [\[SKOS-Mapping\]](#). This at least would let an agent identify source vocabularies, in order to search for a service advertising a vocabulary-mapping capability.

Finally, it would be helpful for vocabulary mapping services to be able to offer a measure of *semantic overlap* between terms. Simple notions have already been proposed (for instance, distinguishing between terms that loosely correspond and those that largely correspond). It is the author's belief that being able to offer probabilistic measures of these subject overlaps - derived from corpus analysis - would be useful in the future.

4 Publishing recommendations and opinions

Assuming a mechanism for publishing these assertions of individual opinions, a recommendation network may be built on top of a FOAF system as follows:

FOAF documents are published by their authors. These may contain all kinds of claims; some of those claims may be assertions of the document's author's opinion on subject-contextualised expertise of other FOAF users.

The documents themselves may link to an external PGP signature for the purposes of authorship verification. (Alternative signature systems, such as XML Signature [\[XML-SIG\]](#), might also be used.)

The user of a FOAF agent may direct it to locate "experts", that is, those highly-recommended in a particular subject area.

A FOAF agent, acting on the behalf of its user, can crawl a web of linked FOAF documents, verifying the authorship of documents as it goes. (Unsigned documents might be discarded in this scenario.) It extracts published opinions on a particular topic and uses these to build an opinion network. Where an opinion is published about an unknown subject, the FOAF agent may attempt to discover a vocabulary-mapping service that relates the term to the one of interest to its user. If such a mapping is found, the opinion can be reexpressed in terms of the user's desired subject; otherwise it is considered irrelevant.

Since FOAF users can publish signed assertions directly, and these can be linked to using RDF's "seeAlso" mechanism, the need to quote the embedded RDF assertions of another user (a notoriously complex problem) can be sidestepped, instead linking directly to the source documents containing those assertions.

The FOAF agent, once the opinion network has been constructed, can evaluate it (actually, this evaluation can normally be done piecemeal as the network is constructed) to produce derived opinions of the expertise of other FOAF users. By using the normal subjective logic ordering on opinions, highly-regarded experts can be located.

4.1 Describing a recommendation -

What remains, therefore, is to consider what the concrete expression in RDF of a recommendation should look like. The following is offered as a strawman. It is inspired largely by the PICS [\[PICS\]](#) system, with some distinctions.

A published recommendation is a member of the class Recommendation. A Recommendation has the following attributes:

- the recommender, which may often be the FOAF:Person who originated the recommendation;
- the recommendee, which is similarly unconstrained, but in our scenario may often be the FOAF:Person who the recommendation is about;
- the *field* is the subject that contextualises the recommendation. (The property name here was chosen to avoid the overloaded term "subject" which already is common in RDF usage). The field of a Recommendation is unconstrained;
- the *opinion*. The value of this attribute is an encoding of opinion; we use the class *SLOpinion* as a carrier for *belief*, *disbelief* and *uncertainty*. (We do not treat relative atomicity here.)

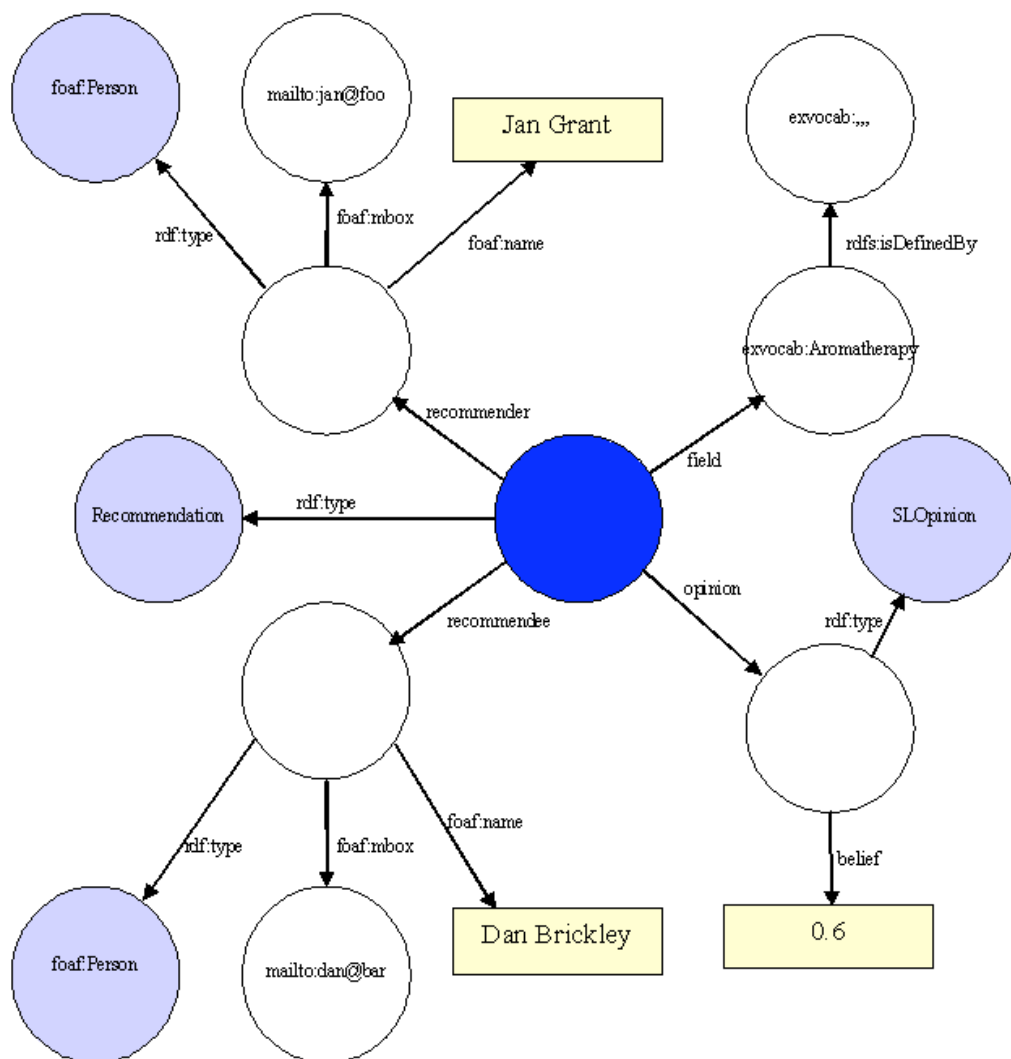
This example schema can be found at [\[REC-SCHEMA\]](#).

It should be noted that this vocabulary permits the expression of opinions of opinions, and so on; if a particular Recommendation (published as RDF/XML) is named by a URI reference, then other agents may publish opinions of *that* opinion.

The semantics of the Recommendation class can only be fully defined in terms of subjective logic; however, the open world nature of RDF assertions must be kept in mind. That is to say, RDF assertions are essentially existentially qualified. As an example: to omit the *field* from a published recommendation does not automatically mean that the recommendation applies equally to all fields (and that assumption is not valid). Rather, to make such an (unlikely) assertion, a value of the *field* must be supplied which encompasses *all possible subjects*. (And vocabulary mapping services would need to reflect this universality.)

It would be a reasonable heuristic for an agent evaluating a recommendation network to supply missing values where *disbelief* and *uncertainty* have been omitted, by setting *disbelief* to zero and *uncertainty* to preserve the $b + d + u = 1$ constraint. Similarly, published opinions that did not preserve this constraint should be discarded (and optionally reported) by a robust agent.

Extensions and other attributes are of course possible; however the essentially conjunctive existential semantics of RDF assertions need to be kept in mind whilst doing so.



The example diagram above indicates the structure of a simple recommendation. The *foaf:Person* named "Jan Grant" offers an opinion on the expertise of "Dan Brickley" in the contextualising field of "Aromatherapy", as defined by a typical vocabulary.

4.2 PICS -

In many regards, PICS \perp [PICS] would seem to offer a ready-made format for expressing recommendations. However, there are a number of issues with adopting PICS as it currently stands for the purposes of publishing recommendations.

"Documents", not "resources"

The first issue, although minor, is that PICS is specified only in terms of ratings for documents, rather than any more general notion of "resource" such as that used by the Semantic Web. An revision of the PICS semantics that adopts a broader scope for the target of a rating would address this.

Numeric attribute values

PICS content labels are consist of attribute-value pairs (each attribute being considered a separate *dimension* of the rating). The values of these pairs are constrained to be only integer or IEEE single-precision floating-point values, ranges of such values, or sequences or sets of such values. While this suffices for encoding a subjective-logic opinion, it does not directly suffice for encoding the subject contextualisation of a recommendation, barring extreme measures of encoding character sequences numerically. The adoption of URI references as identifiers is widespread. Therefore, it would seem that to adapt PICS to represent values named by URI references (as opposed to the numeric enumeration scheme it currently uses) would be a step forward in this regard.

Independent rating dimensions

The various dimensions of a PICS content label are considered independent according to the PICS semantics. That is, by simply encoding "subject contextualisation" as an attribute value in a rating, and a subjective logic opinion as another attribute value, PICS semantics do not support the association of one value with the other. The subject contextualisation is therefore lost. This leads on to the final point.

Closed sets of dimensions

In order to retain the subject-contextualisation of an opinion, the opinion could be represented as the value of a label that corresponds to the contextualising subject. And this would seem to solve the matter: a PICS ratings service can define a rating system that contains any number of "dimensions", each of which correspond to the opinion contextualised by a particular subject classification. Unfortunately, the rating system is effectively a "closed world" predefined by the rating service, and is not appropriate for the expression of arbitrary terms taken from disparate vocabularies.

Conclusion

The expression of PICS labels was one of the motivating requirements behind the original design of RDF. RDF is adequate to express PICS labels. PICS itself is a useful technology with existing implementations that support a specific problem domain. However, the open world of the Semantic Web brings a new requirement to a recommendation system (the need to cope with terms drawn from arbitrary numbers of vocabularies). RDF permits the natural expression of such recommendations, and a general recommendation vocabulary built upon RDF would seem to be a better solution than to attempt to retrofit the requirement onto the current version of PICS.

5 Issues with trust in social networks

Any discussion of trust in social networks cannot ignore the social aspects of such networks. Issues that have arisen during this investigation are presented below.

5.1 Cliques -

Social networks tend to have many tightly-interconnected subnetworks, which one might term "cliques". The members of a clique are all known to each other. Often they are drawn together through a shared interest (particularly in the FOAF world as it currently stands) and consequently may all have reasonably high degrees of expertise in their shared interest.

It is the case, then, that members of a clique are often legitimately in the position of making recommendations of each other contextualised by the same subject classification. Such circular interrecommendations are often problematic to evaluate using probabilistic belief networks. At best, an approximation of an evaluated opinion may be derived.

However, since the social cause of cliques is unlikely to disappear, further study is required to locate efficient ways to accurately approximate evaluations of opinion networks with cycles.

5.2 Libel -

An entirely different problem is this: by publishing a signed opinion of the expertise of another person, one may be held liable (or at least threatened) for the recommendation network equivalent of "defamation of character". Many people welcome critical peer review, and consider feedback to be important. There are many popular discussion systems that use feedback and user ratings already (eg, Slashdot's *Karma*), but the feedback is often anonymous.

Consequently, it is plausible that people would refrain from publishing direct recommendations of others for fear of causing offence.

5.3 Peer pressure -

The situation described above is one example of a general phenomenon, that of *peer pressure*. Groups of people have social motivations to offer artificially-inflated opinions, either for ostensibly innocent reasons (simply mutual back-patting) or for more obvious gain.

An anecdotal example of peer pressure in a reputation system might be the feedback system of the online auction service, eBay [\[EBAY\]](http://www.ebay.com). Since both parties to a transaction are able to offer feedback on each other, which is amalgamated into a numerical score, there is a tendency to mitigate negative feedback on a transaction in the fear of inviting a retributory negative opinion from the person with whom one is dealing. This is a simple example where external considerations mean that it is not in the interests of the user to offer as objective an opinion as that can make.

5.4 Making objective assessments -

We have thus far couched a recommendation network as being composed of published subjective logic opinions on propositions such as, "E is an expert in the field of S". However, it is hard for a user to directly produce an objective assessment of expertise in another; and it is certainly also hard to attempt to reduce that assessment to a probabilistic set of values.

In order to fully complete the picture of a recommendation network, then, it is necessary to consider how a user might form an opinion of the expertise of another. An opinion is formed by weighing evidence. For instance, suppose a person publishes articles, or other documents on a particular topic. Rather than directly attempt to rate that person, one might assess and assign a rating to each of the articles, taken as evidence of expertise; and modifying one's opinion that "this person is an expert" by considering if the document appears to have been written from a position of expertise.

By focussing on documents and other outputs as evidence, some of the social aspects of peer pressure may be defused. Informally, it is less "rude" to criticise an article than to directly attack the expertise of the author.

6 Applications of recommendation systems

Typical use scenarios given for FOAF begin with the location of a set of people related via some shared characteristics. For example:

- all people who are related by the *knows* relationship - that is, immediate friends, their acquaintances, and so on up to some degree;
- people who share a particular *self-declared* interest or expertise.

With the use of a recommendation system, we can add to this a third mechanism for selecting people, namely:

- people who are recommended as experts in a particular field.

Once a set of target people have been identified, the claims they (and others) make can be queried

further, for instance to look for other recommendations they might make; or conferences they are visiting; for documents they've authored, and so on.

It is a difficult task to explicitly enumerate use cases for social networking tools. The flexibility of FOAF and RDF means that many potential applications are possible. The ability to formulate general queries over aggregated RDF is a key enabler in these applications. By adding an application layer of recommendations, potential user scenarios are given an additional dimension.

One other possibility that merits some future investigation is as follows. The recommendation system described above is built as an application on top of RDF assertions, with additional semantics. Opinions are formed of particular propositions. Now, RDF triples can be seen as the assertion of particular propositions. Therefore it is at least conceivable that we might ask what the measured opinion of a collection of agents is, about the assertion encoded by a particular triple or set of triples.

Aside from the technical difficulties of evaluating such an opinion network (since the triples that assert recommendations may themselves be the subject of varying opinions, such a system is nonmonotonic at best - and it is reasonably simple to create a system which is unstable), there seems to be a conceptual difficulty when talking about the level of trust in a particular triple. At an application level, we are used to viewing a collection of RDF triples as a whole. For instance, we might say informally that a particular FOAF document "is about" a person; and such an assertion is understandable. However, typical FOAF documents describing people are full of triples that have no obvious relationship to "the subject of the document", when viewed in isolation.

7 Conclusion

A simple recommendation system based around subjective logic has been presented, including a simple vocabulary for expressing recommendations using RDF.

Vocabulary-mapping systems have a key role to play in the semantic web. By including quantitative measurements of semantic overlap (taken from corpus analysis), subject-contextualisation in opinion networks can be approached using subjective logic.

PICS has been briefly considered as a mechanism for expressing recommendations. While it is suitable for simple opinion networks, the open-endedness of subject contextualisation means that PICS is not suitable for such a task. For this application, the use of general RDF assertions should be seen as a natural "successor" to PICS.

The particular approach taken has a problem due to the over-general nature of the definition of "expertise" adopted, which leads to cycles in the opinion network. However, the vocabulary suggested permits the direct assertion of opinions of opinions (and so on). In this case, the opinion networks resemble more closely the cycle-free networks of subjective logic.

Due to the algebraic properties of the subjective logic operators, opinion networks can be evaluated "piecemeal"; the evaluation of these networks is not overly expensive.

Appendix A. References

[DS-CONDITIONAL]

<http://security.dstc.edu.au/papers/JG2003-Fusion.pdf> Conditional inference in subjective logic: Audun Jøsang, Tyrone Grandison. (Online copy of preprint.)

[DS-DEMO]

<http://security.dstc.edu.au/spectrum/trustengine/demo2.html> Audun Jøsang, Matthew Davey. A Java applet demonstrating discounting and consensus in a simple opinion network.

[DS-DOGMATISM]

<http://security.dstc.edu.au/papers/JDV2003-Fusion.pdf> Strategies for combining conflicting dogmatic beliefs: Audun Jøsang, Milan Daniel. (Online copy of preprint.)

[DS-INTRO]

<http://security.dstc.edu.au/papers/logunprob.pdf> A logic for uncertain probabilities: Audun Jøsang. (Updated copy of preprint.)

[EBAY]

<http://pages.ebay.co.uk/help/feedback/feedback-scores.html> A description of the simple recommendation/feedback system used by eBay.

[FOAF]

<http://www.foaf-project.org> The main FOAF project site.

[FOAF-SIGN]

<http://usefulinc.com/foaf/signingFoafFiles> PGP signing FOAF files: Useful Information Company. A simple HOW-TO guide to publishing FOAF documents assured by PGP signature.

[OWL-SEMANTICS]

<http://www.w3.org/TR/owl-semantics/> OWL Web Ontology Language Semantics and Abstract Syntax: Peter F. Patel-Schneider, Patrick Hayes, Ian Horrocks, eds.

[PICS]

<http://www.w3.org/PICS/> Platform for Internet Content Selection.

[RDF-SEMANTICS]

<http://www.w3.org/TR/rdf-mt/> RDF semantics: Patrick Hayes, ed.

[REC-SCHEMA]

<http://www.w3.org/2001/sw/Europe/reports/trust/recSchema.rdf> A strawman schema for expressing contextualised recommendations in RDF: Jan Grant.

[SKOS-Mapping]

<http://www.w3.org/2001/sw/Europe/reports/thes/rdfthes.html> An RDF vocabulary for inter-thesaurus mapping: Alistair Miles. Prerelease version at <http://www.w3c.rl.ac.uk/2003/11/21-skos-mapping>.

[XML-SIG]

<http://www.w3.org/TR/xmlsig-core/> XML-Signature Syntax and Processing: Donald Eastlake, Joseph Reagle, David Solo, eds.

Appendix B. Summary of operators in subjective logic

The definitions below are taken from [\[DS-INTRO\]](#) and [\[DS-CONDITIONAL\]](#) and summarised for convenience.

Opinion as (b, d, u, a)

An opinion of a proposition, P , is given by the four-tuple (b, d, u, a) [\[DS-INTRO\]](#). The *belief* in the proposition is b . The *disbelief* is d , defined as the belief in the proposition being false. The *uncertainty* is u , satisfying $b + d + u = 1$. The *relative atomicity* of the assertion is a . For a simple subjective logic application, relative atomicity can be omitted from consideration (taken as $a=0.5$). We include the effect of subjective logic operators on relative atomicity for completeness here.

The probability expectation of an opinion, $E((b,d,u,a))$ is given by $E=b+ua$.

Ordering opinions

From [\[DS-INTRO Def 9, 10\]](#), opinions can be ordered on the basis of:

- greatest probability expectation;
- least uncertainty;
- least relative atomicity.

Conjunction

Given opinions (b_1, d_1, u_1, a_1) on proposition P_1 , and (b_2, d_2, u_2, a_2) on proposition P_2 , then the combined opinion that $(P_1 \text{ AND } P_2 \text{ are true})$ is given by (b, d, u, a) where:

- $b = b_1 b_2$
- $d = d_1 + d_2 - d_1 d_2$
- $u = b_1 u_2 + b_2 u_1 + u_1 u_2$
- $a = (b_1 u_2 a_2 + b_2 u_1 a_1 + u_1 a_1 u_2 a_2) / (b_1 u_2 + b_2 u_1 + u_1 u_2)$

Disjunction

Given opinions (b_1, d_1, u_1, a_1) on proposition P_1 , and (b_2, d_2, u_2, a_2) on proposition P_2 , then the combined opinion that $(P_1 \text{ OR } P_2 \text{ is true})$ is given by (b, d, u, a) where:

- $b = b_1 + b_2 - b_1 b_2$
- $d = d_1 d_2$
- $u = d_1 u_2 + d_2 u_1 + u_1 u_2$
- $a = (u_1 a_1 + u_2 a_2 - b_1 u_2 a_2 - b_2 u_1 a_1 - u_1 a_1 u_2 a_2) / (u_1 + u_2 - b_1 u_2 - b_2 u_1 - u_1 u_2)$

Negation

Given opinion (b,d,u,a) on proposition P , the opinion that $(P$ is false), (b',d',u',a') is given by:

- $b' = d$
- $d' = b$
- $u' = u$
- $a' = 1 - a$

Consensus

Given two independent opinions of the truth of a proposition P , (b_1,d_1,u_1,a_1) and (b_2,d_2,u_2,a_2) , then the consensus between those opinions of the truth of the proposition is (b,d,u,a) given by:

- $b = (b_1u_2 + b_2u_1) / k$
- $d = (d_1u_2 + d_2u_1) / k$
- $u = u_1u_2 / k$
- $a = (a_1u_2 + a_2u_1 - (a_1 + a_2)u_1u_2) / (u_1 + u_2 - 2u_1u_2)$

where

- $k = u_1 + u_2 - u_1u_2$ such that k is not zero, and
- $a = (a_1 + a_2) / 2$, when $u_1 = u_2 = 1$

Discounting

Given two opinions, (b_1,d_1,u_1,a_1) of Agent A_1 that some proposition P_1 is true, and (b_2,d_2,u_2,a_2) of Agent A_2 that " A_1 is knowledgeable and will tell the truth" (ie, an opinion of A_1 's opinion), then A_2 can form a *discounted* opinion of P_1 , (b,d,u,a) , given by:

- $b = b_2b_1$
- $d = b_2d_1$
- $u = d_2 + u_2 + b_2u_1$
- $a = a_1$

Conditional inference

The following definition is taken from [DS-CONDITIONAL]

Given an opinion in proposition P , (b_p,d_p,u_p,a_p) ; an opinion on the proposition $(P$ implies $P')$ of $(b_{p|p},d_{p|p},u_{p|p},a_{p|p})$; and an opinion on the proposition $(NOT P$ implies $P')$ of $(b_{p|\sim p},d_{p|\sim p},u_{p|\sim p},a_{p|\sim p})$, then we can form an opinion on the proposition P' , (b,d,u,a) , according to the following cases:

- Case I: $((b_{p|p} \geq b_{p|\sim p}) \text{ AND } (d_{p|p} \geq d_{p|\sim p})) \text{ OR } ((b_{p|p} \leq b_{p|\sim p}) \text{ AND } (d_{p|p} \leq d_{p|\sim p}))$
 - $b^I = b_p b_{p|p} + d_p b_{p|\sim p} + u_p (b_{p|p} a_p + b_{p|\sim p} (1 - a_p))$
 - $d^I = b_p d_{p|p} + d_p d_{p|\sim p} + u_p (d_{p|p} a_p + d_{p|\sim p} (1 - a_p))$
 - $u^I = b_p u_{p|p} + d_p u_{p|\sim p} + u_p (u_{p|p} a_p + u_{p|\sim p} (1 - a_p))$
 - $a^I = a_p$ (relative atomicity of P' is assumed to be constant, so $a_p = a_{p|p} = a_{p|\sim p} = a_{p|\sim p}$)
- Case II.A.1: $((b_{p|p} > b_{p|\sim p}) \text{ AND } (d_{p|p} < d_{p|\sim p})) \text{ AND } (E \leq b_{p|\sim p} + a_p(1 - b_{p|\sim p} - d_{p|p})) \text{ AND } (b_p + a_p u_p \leq a_p)$
 - $b^{II.A.1} = b^I - (u_p a_p a_p (b^I - b_{p|\sim p})) / ((b_p + a_p u_p) a_p)$
 - $d^{II.A.1} = d^I - (u_p a_p (1 - a_p) (d^I - b_{p|\sim p})) / ((b_p + a_p u_p) a_p)$
 - $u^{II.A.1} = u^I + (u_p a_p (b^I - b_{p|\sim p})) / ((b_p + a_p u_p) a_p)$
 - $a^{II.A.1} = a_p$
- Case II.A.2: $((b_{p|p} > b_{p|\sim p}) \text{ AND } (d_{p|p} < d_{p|\sim p})) \text{ AND } (E \leq b_{p|\sim p} + a_p(1 - b_{p|\sim p} - d_{p|p})) \text{ AND } (b_p + a_p u_p > a_p)$
 - $b^{II.A.2} = b^I - (u_p (1 - a_p) a_p (b^I - b_{p|\sim p})) / ((d_p + (1 - a_p) u_p) a_p)$
 - $d^{II.A.2} = d^I - (u_p (1 - a_p) (1 - a_p) (d^I - b_{p|\sim p})) / ((d_p + (1 - a_p) u_p) a_p)$

- $u^{II.A.2} = u^I + (u_P(1 - a_P)(b^I - b_{P|\neg P})) / ((d_P + (1 - a_P)u_P) a_P)$
- $a^{II.A.2} = a_P$
- Case II.B.1: $((b_{P|P} > b_{P|\neg P}) \text{ AND } (d_{P|P} < d_{P|\neg P})) \text{ AND } (E > b_{P|\neg P} + a_P(1 - b_{P|\neg P} - d_{P|P})) \text{ AND } (b_P + a_P u_P \leq a_P)$
 - $b^{II.B.1} = b^I - (u_P a_P a_P (d^I - d_{P|P})) / ((b_P + a_P u_P) (1 - a_P))$
 - $d^{II.B.1} = d^I - (u_P a_P (1 - a_P)(d^I - d_{P|P})) / ((b_P + a_P u_P) (1 - a_P))$
 - $u^{II.B.1} = u^I + (u_P a_P (d^I - d_{P|P})) / ((b_P + a_P u_P) (1 - a_P))$
 - $a^{II.B.1} = a_P$
- Case II.B.2: $((b_{P|P} > b_{P|\neg P}) \text{ AND } (d_{P|P} < d_{P|\neg P})) \text{ AND } (E > b_{P|\neg P} + a_P(1 - b_{P|\neg P} - d_{P|P})) \text{ AND } (b_P + a_P u_P > a_P)$
 - $b^{II.B.2} = b^I - (u_P(1 - a_P)a_P(d^I - d_{P|P})) / ((d_P + (1 - a_P)u_P) (1 - a_P))$
 - $d^{II.B.2} = d^I - (u_P(1 - a_P)(1 - a_P)(d^I - d_{P|P})) / ((d_P + (1 - a_P)u_P) (1 - a_P))$
 - $u^{II.B.2} = u^I + (u_P(1 - a_P)(d^I - d_{P|P})) / ((d_P + (1 - a_P)u_P) (1 - a_P))$
 - $a^{II.B.2} = a_P$
- Case III.A.1: $((b_{P|P} < b_{P|\neg P}) \text{ AND } (d_{P|P} > d_{P|\neg P})) \text{ AND } (E \leq b_{P|\neg P} + a_P(1 - b_{P|\neg P} - d_{P|P})) \text{ AND } (b_P + a_P u_P \leq a_P)$
 - $b^{III.A.1} = b^I - (u_P a_P a_P (b^I - b_{P|P})) / ((b_P + a_P u_P) a_P)$
 - $d^{III.A.1} = d^I - (u_P a_P (1 - a_P)(b^I - b_{P|P})) / ((b_P + a_P u_P) a_P)$
 - $u^{III.A.1} = u^I + (u_P a_P (b^I - b_{P|P})) / ((b_P + a_P u_P) a_P)$
 - $a^{III.A.1} = a_P$
- Case III.A.2: $((b_{P|P} < b_{P|\neg P}) \text{ AND } (d_{P|P} > d_{P|\neg P})) \text{ AND } (E \leq b_{P|\neg P} + a_P(1 - b_{P|\neg P} - d_{P|P})) \text{ AND } (b_P + a_P u_P > a_P)$
 - $b^{III.A.2} = b^I - (u_P(1 - a_P)a_P(b^I - b_{P|P})) / ((d_P + (1 - a_P)u_P) a_P)$
 - $d^{III.A.2} = d^I - (u_P(1 - a_P)(1 - a_P)(b^I - b_{P|P})) / ((d_P + (1 - a_P)u_P) a_P)$
 - $u^{III.A.2} = u^I + (u_P(1 - a_P)(b^I - b_{P|P})) / ((d_P + (1 - a_P)u_P) a_P)$
 - $a^{III.A.2} = a_P$
- Case III.B.1: $((b_{P|P} < b_{P|\neg P}) \text{ AND } (d_{P|P} > d_{P|\neg P})) \text{ AND } (E > b_{P|\neg P} + a_P(1 - b_{P|\neg P} - d_{P|P})) \text{ AND } (b_P + a_P u_P \leq a_P)$
 - $b^{III.B.1} = b^I - (u_P a_P a_P (d^I - d_{P|\neg P})) / ((b_P + a_P u_P) (1 - a_P))$
 - $d^{III.B.1} = d^I - (u_P a_P (1 - a_P)(d^I - d_{P|\neg P})) / ((b_P + a_P u_P) (1 - a_P))$
 - $u^{III.B.1} = u^I + (u_P a_P (d^I - d_{P|\neg P})) / ((b_P + a_P u_P) (1 - a_P))$
 - $a^{III.B.1} = a_P$
- Case III.B.2: $((b_{P|P} < b_{P|\neg P}) \text{ AND } (d_{P|P} > d_{P|\neg P})) \text{ AND } (E > b_{P|\neg P} + a_P(1 - b_{P|\neg P} - d_{P|P})) \text{ AND } (b_P + a_P u_P > a_P)$
 - $b^{III.B.2} = b^I - (u_P(1 - a_P)a_P(d^I - d_{P|\neg P})) / ((d_P + (1 - a_P)u_P) (1 - a_P))$
 - $d^{III.B.2} = d^I - (u_P(1 - a_P)(1 - a_P)(d^I - d_{P|\neg P})) / ((d_P + (1 - a_P)u_P) (1 - a_P))$
 - $u^{III.B.2} = u^I + (u_P(1 - a_P)(d^I - d_{P|\neg P})) / ((d_P + (1 - a_P)u_P) (1 - a_P))$
 - $a^{III.B.2} = a_P$

Throughout, E (the expected probability of P^I given P) is given by:

- $E = b^I + a^I u^I$

(For an explanation of these apparently complex formulae see the referenced paper; in fact, these cases are all closely related and are motivated by a simple geometric argument.)