2016-01-01

A Semantically-Enabled Trust Model For Collaborative Environments

Joaquin Andres Reyna
University of Texas at El Paso, jareyna2@miners.utep.edu

A SEMANTICALLY-ENABLED TRUST MODEL FOR COLLABORATIVE ENVIRONMENTS

JOAQUIN ANDRES REYNA CRUZ

Master's Program in Computer Science

APPROVED:

____
Natalia Villanueva, Ph.D., Chair

____
Deana Pennington, Ph.D., Co-Chair

____
Olac Fuentes, Ph.D.

____
Charles Ambler, Ph.D.
Dean of the Graduate School

to my

MOTHER

thank you for everything

A SEMANTICALLY-ENABLED TRUST MODEL FOR COLLABORATIVE ENVIRONMENTS

by

JOAQUIN ANDRES REYNA CRUZ

THESIS

Presented to the Faculty of the Graduate School of

The University of Texas at El Paso

in Partial Fulfillment

of the Requirements

for the Degree of

MASTER OF SCIENCE

Master's Program in Computer Science

THE UNIVERSITY OF TEXAS AT EL PASO

May 2016

Acknowledgements

I would like to express my deep-felt gratitude to my advisor, Dr. Natalia Villanueva Rosales of the Computer Science Department at The University of Texas at El Paso, for her advice, encouragement and help this past year.

I also wish to ackowledge Dr. Pennington for her lead and the Cyber-ShARE Center of Excellence for their support as this work used resources from their center. This work was supported in part by NSF grant HRD-0734825 and OCI-1135525. Additionally, I want to thank the Computer Science professors and staff for all their hard work and dedication, they provided me the means to complete my degree and prepare for a career as a computer scientist. I would like to give a special thanks to Blanca Cruz my mother, to whom I dedicate this work. It because of her hard work through all my life that I have achieved any and all objectives, including this work.

Abstract

The World Wide Web has revolutionized the way we find, exchange, reuse and integrate information at a rate inconceivable only a generation ago; but for all its perks it has one big weakness: trusting that the information retrieved is accurate. Web technologies and applications allows anyone to share information. This work aims to assist humans (and eventually autonomous agents) in the task of trusting Web resources, a very important step when reusing information found on the Web. In this research, we propose a new algorithm to tackle the challenge of assigning trust to Web resources. Our algorithm is inspired on Google's famous algorithm PageRank and generates values of trustworthiness of Web resources based on dynamically adjusted user credentials on different areas of study and history of interactions. The proposed algorithm was implemented in the Virtual Learning Commons, a bookmarking system, developed and hosted at the Cyber-ShARE Center of Excellence. The proposed algorithm was evaluated in an academic environment but can be used in other applications. The evaluation shows that the algorithm produced outperforms two out of the three algorithms we benchmarked against, while matching the results of the third. The evaluation results also show that the algorithm provides adequate trust assessments for Web resources. Future work includes the extension of this algorithm to consider additional variables (i.e., beyond user credentials) and a more comprehensive evaluation settings, including a more diverse subject population.

Efforts that facilitate Web users to assign values of trustworthiness to Web resources will contribute to the overall utility of the world wide web, it will leverage the knowledge of the world and will formalize said knowledge so that machine agents may use to our benefit.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

The World Wide Web (WWW) enables any person or institution to publish information as Web resources. Trusting in a Web resource, which typically involves considering the source of the Web resource, guides the user on how to make use of the information found. Therefore, assigning a value of trust to Web resources is a task that more and more users are facing, especially those who rely on the information found to derive new knowledge or make decisions. For example, a researcher might use the information retrieved on the web as a building block for his work, or a student might use the information as complementary material to learn the subject. In both cases trust is very important, albeit to different degrees.

1.2 Hypothesis and Objectives

Our hypothesis is that an algorithm based on user credentials with dynamically adjusted weights and a history of human endorsements can accurately classify Web resources as trustworthy or not. It is also our hypothesis that the classification accuracy will be higher than that of simpler algorithms that do not include this information or make static use of the same. The primary objectives of this research are:

1. Define an algorithm based on user preferences (i.e., credentials and endorsements) that calculates a value of trustworthiness of Web resources.

2. Compare the proposed algorithm with simpler algorithms that do not consider user preferences or consider them in a limited fashion(more on this on the testing chapter) to evaluate its performance.

In this thesis, we propose an algorithm that calculates the trust of a Web resource based on the credentials of the user that shared said Web resource (source), and how this user is perceived by the community, that is, given the sharer's reputation and finally we consider the communities' stated trust in the Web resource. We enable community members to provide individual input (an "endorsement") on Web resources, determining whether they trust the content of the same or not. Weights are given to these endorsements. Endorsements and credentials in our algorithm are domain dependent, that is, our trust in a person's opinion for Web resources in the field of Computer Science does not necessarily translate into trust in other fields such as Geology. The algorithm makes this distinction because it is the way credentials are naturally handled in our society, after all we might trust the opinion of a psychology doctorate on psychoanalysis theory but intuitively we would not trust his opinion on the best metallurgy procedure, the user is still a Ph.D. but the subject area in which his credentials apply is an essential and inescapable part of the context. In a nutshell the algorithm propagates trust when a new endorsement is created. When this happens the total trust of the user is redistributed equally among all of their endorsements. The trust that we have in a user (which is distributed to his endorsements) is a combination of his credentials (whose weight is adjusted dynamically) and the sum of the endorsements given to her shared Web resources. The process is recursive, more specifically it falls within the category of graph and wave algorithms, and as such the readjustment of trust propagates from all Web resources to their owners to the Web resources this person has endorsed through all connected nodes in the subgraph.

Two things must be noted. First, everything is initiated by a new endorsement. In our algorithm a share/upload is an implicit endorsement by the Web resource owner, therefore the process also happens when a new Web resource is shared. Secondly our algorithm makes the assumption that the credentials of the user are correct. We make the assumption that

a fully-fledged implementation of the system will require verification of an email address. For example, in an academic setting, the system would verify that an email ends in ".edu" and would have access to appropriate academic records to verify that such an address corresponds to the user.

The proposed algorithm was implemented on a Web-based system that uses Semantic Web technologies and was evaluated in an academic environment, the Virtual Learning Commons system (VLC). Credentials of the user considered their area of study (e.g. Computer Science, Geology, etc.) and their classification (e.g. undergraduate student, professor, etc.). The main piece of information elicited from a user was whether she considered a Web resource's content as trustworthy (as previously stated users were given controls to "endorse" Web resources). The algorithm was evaluated with a group of 30 users from the areas of Computer Science and Geology (15 from each area). Each group was directed to a page where they found Web resources both in their area of study and out of it, where half of those resources were trustworthy and the other half were non-trustworthy (defined as such by an expert in the field). Participants were asked to click a link next to the ones whose content they deemed accurate. The evaluation study was approved by the Institutional Research Board (IRB), and no questionnaire was required of the participants, the IRB approval can be found in Appendix C.The results of the evaluation study show that the algorithm performs better or on par with three other algorithms. The first was a simple voting algorithm where a constant value (in our case 1) is given to every Web resource for every endorsement given to it. The second algorithm is one that does not consider credentials but is otherwise equivalent to our algorithm. The last algorithm we compare against considers credentials but does so in a static manner; their weights are never readjusted. After analysis of the results, we discovered that the algorithm might provide more accurate results given certain conditions.

1.3 Contributions

The major contributions of this thesis are:

1. The design of an algorithm to compute a value of trustworthiness of a Web resource.
2. The creation of an ontology to annotate the information required by the algorithm.
3. The implementation and evaluation of the proposed algorithm.

In addition, this work contributed to the design and implementation of features for the VLC, a bookmarking and data sharing system developed and hosted by the University of Texas at El Paso (UTEP)[Pennington et al., 2012]. The features deal specifically with the sharing and annotation of Web resources in an academic context.

Although our algorithm was tailored to an academic context, it can be applied to other contexts, such as those addressed by other algorithms in the related work section. In this chapter we provided an overview of the algorithm, its use and our hypotheses. The rest of the thesis is divided as follows: Chapter 2 provides background information on the core concepts used in this thesis, i.e., trust and semantic web technologies. Chapter 3 provides a definition of the proposed trust algorithm. The implementation of the algorithm in the VLC system is presented in Chapter 4. Chapter 5 presents the evaluation of the algorithm along with the results. Chapter 6 presents an overview of related work, Chapter 7 presents the conclusions of this work and Chapter 8 outlines future directions. Appendix A includes a Manchester serialization of our ontology and Appendix B includes the approval document (IRB) by UTEP.

Chapter 2

Background

In the previous chapter we introduced an overview of this work including our hypothesis and algorithm. In this chapter we introduce the core concepts in which our work is based.

2.1 Trust

Trust has been defined in many ways and contexts. The authors in [Mui et al., 2002] state that "[Trust is] a subjective expectation an agent has about another's future behavior based on the history of their encounters". The authors in [Grandison and Sloman, 2000] state that "[Trust is] the firm belief in the competence of an entity to act dependably, securely, and reliably within a specified context". In the literature the terms agent and entity are used somewhat interchangeably and can mean anyone or anything be it a human user or a program, in our work we only consider human users who have verifiable credentials, but as a subclass of entities/agent any knowledge in the literature applies to them.

Previous research for creating and managing trust is presented in [Artz and Gil, 2007] using the following categories:

- Policy-based trust - which focuses on exchanging user credentials and enforcing access policies.
- Reputation-based trust - which uses past interactions of an agent to predict its future behavior.
- General models of trust - modeling trust in different domains.

- Trust in information resources - trust on the Web has varying uses and meanings which include capturing ratings from users about the quality of resources, propagating trust over links, and so on.

Our work does not ascribe to either reputation or competence definitions entirely. Instead it falls between the categories of reputation-based trust and trust in information resources. The proposed algorithm can be used as a stand-alone reputation systems or integrated into reputation systems that works with Web resources, user interactions and propagation of trust between Web resources and users.

2.2 Semantic Web

The Semantic Web is an extension of the current Web where the information is published for machine consumption(computer agents) whereas the standard web is populated with documents for human consumption [Berners-Lee et al., 2001]. The Semantic Web allows the automatic answering of questions, because in addition to data, the context of the data is also published. This context is encoded in ontologies (more in a following subsection) that define the rules and relationships pertaining to the domain to which the data belongs to or is relevant in. One such piece of context is the **provenance of the information** (i.e., the origin of the data), which is of vital importance to trust. Credentials of the agents(machine or human) and a history of past interactions is what allows any other agent to make a decision of whether to trust the information provided. In addition to keeping provenance information about where the Web resource originated we also keep provenance for the trust (where the trust for a particular Web resource comes from).

2.2.1 Semantic Web Framework

The Semantic Web framework is composed of languages, tools and methods depicted in Figure 2.1.

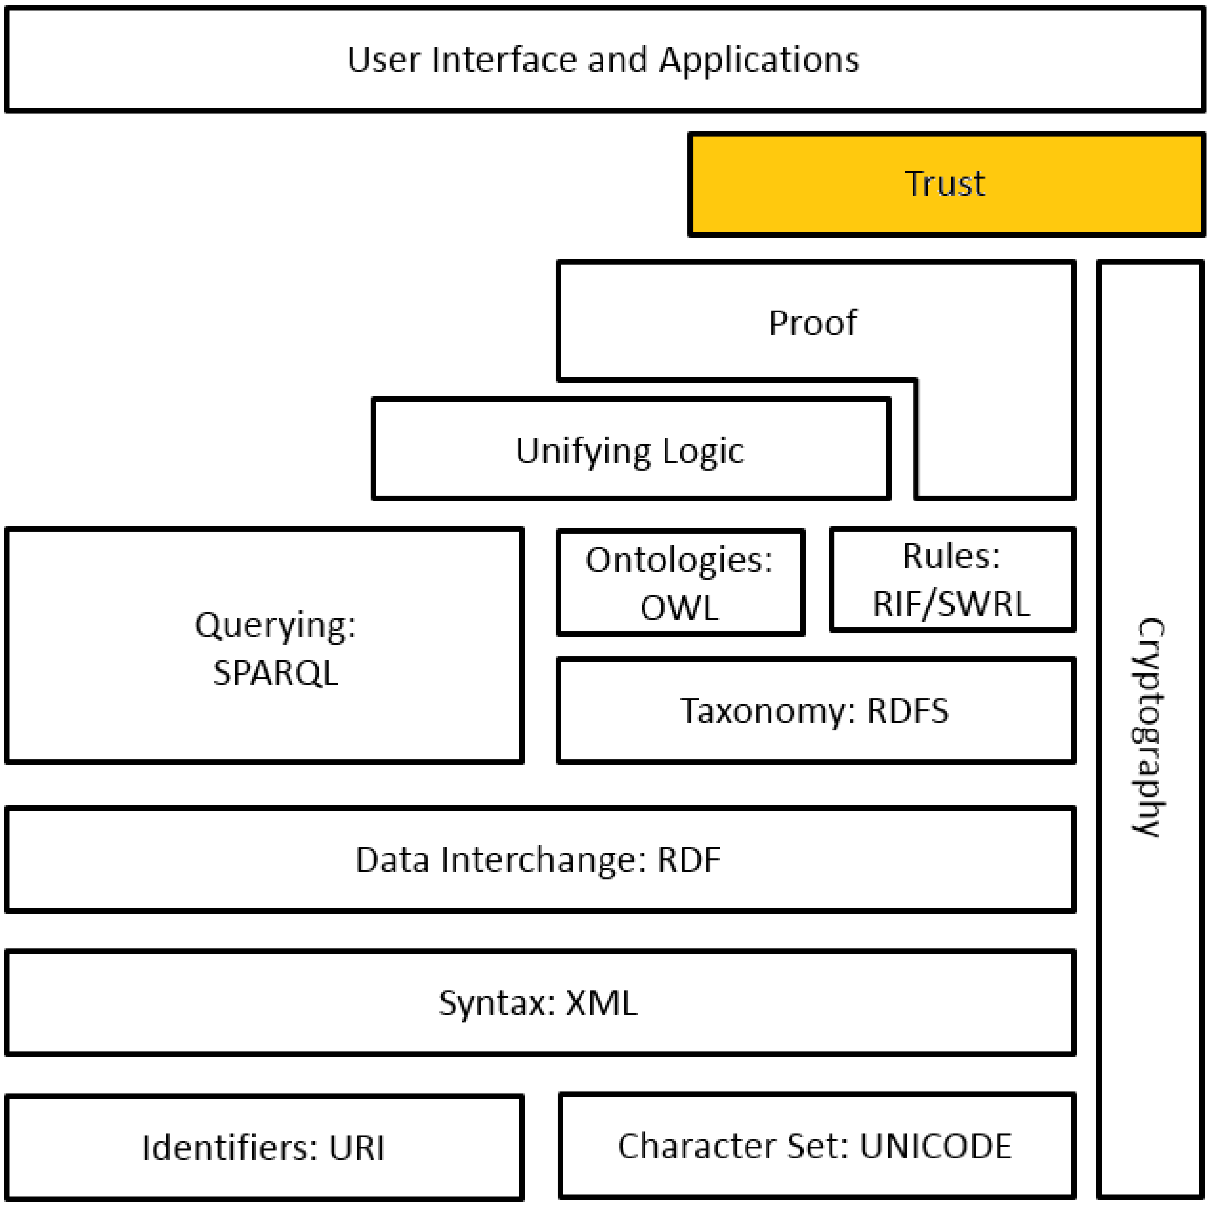

Figure 2.1: The Semantic Web Stack modified from [Berners-Lee et al., 2001]. illustrates the collection of components and languagues that allow the realization of the Semantic Web vision. The yellow layer indicates the trust component, which is the focus this thesis.

Applications using the Semantic Web framework use Unicode language for character rep-

resentation [Aliprand et al., 2003] and the Universal Resource Identifier [Berners-Lee et al., 1994] as the naming syntax. On top (in the stack) of this encoding format and naming syntax is the eXtensible Markup Language (XML) [Bray et al., 2000]. This language can be used to create documents that store semi-structured data in the form of elements, and attributes of the elements. Since anyone can create an XML document and therefore name collisions are a probabilistic certainty, disambiguation is performed through the use of namespaces and the use of a schema language (XMLS) [Fallside and Walmsley, 2004]. The guarantee of unambiguity allows for data exchange; integration and the use of schemas allow for document validation.

One level above in the stack, sitting on top of XML is the technology which allows the generation of data models and representation of information, the Resource Description Framework (RDF)(explained more below) and the RDF Schema (RDFS)[Brickley and Guha, 2003]. The latter defines the notion of class and provides a means of defining class hierarchies with an *is-a* relationship as well as property hierarchies by the use of the *sub property* relationship.

The next level in the stack, above RDF/RDFS, is the Web Ontology Language(OWL)(explained further in the following section).

At the same level as OWL is the SPARQL Query language [PrudHommeaux et al., 2008] that represents RDF queries (which can query across diverse data sources) and the Rule Interchange Format (RIF) [Kifer, 2008]. Both are W3C recommendations.

The last two levels are proof and trust and are meant to enable the work of the agents described in the Semantic Web vision [Berners-Lee et al., 2001].

Our work adheres to the Semantic Web guidelines by defining an OWL ontology that describes all the elements required to calculate a value of trustworthiness of a Web resource. The implementation of the algorithm in the VLC exchanges and stores data using RDF and a triple store (more on this on the next subsection). Specifically we used RDF triples to annotate the Web resources shared in the VLC and to encode the endorsements made by the users when they interacted with the VLC. RDF-encoded metadata about the Web

resources also included the score of trustworthiness after being calculated.

2.3 Resource Description Framework

The Resource Description Framework (RDF) is "a language to represent information about resources in the World Wide Web"[Manola et al., 2004][Cyganiak et al., 2014]. It represents this information through 'statements'. The statements encode two resource/entities and the relation between them. Each statement has a subject (the first entity), a predicate (a relationship) and a predicate(second entity). These statements are called triples. An example of a triple would be John is Human (John, is, Human), with John as the subject, "is" the predicate and Human the object.

In practice triples are stored in what is called triplestores; specialized databases that organize the data not in tables and rows but rather in graphs containing triples. Instead of SQL they use a special type of query language, SPARQL, and use ontologies (more in next subsection) instead of schemas. One example of a triplestore is the one we used for our work, arc2. The main difference between arc2 and other triplestores such as virtuoso is that it supports SPARQL+, a distinct SPARQL dialect with its own subset of commands.

2.4 Ontologies

An ontology is a formal and axiomatic definition of entities, their properties and relationships. Ontologies are normally encoded using RDF or the Web Ontology Language (OWL). Using the previous example we can use the triple (John,is, Human) and define in an ontology the class Human and have it represent the class of all humans.

As part of our work we created an ontology that defines the entities and actions in our system and we used this ontology to provide a formal, logical framework for the data generated and managed by our algorithm. This means that every triple we generate and store describes the transactions in our implemented system and adhere to our ontology.

Ontologies are meant to provide a formal definition of a domain but they are designed so that knowledge in one domain can be combined and compatible with that in similar domains. This is done by mapping the entities and properties of one ontology to that of another though equality and subclass relationships. For example if we have an ontology for a bookstore and for a library, it is likely that we will find the concept of book in both but the information associated with it in each ontology is necessarily different because of the needs of the domain. For the bookstore's book (let's call it bkst:book) there may be information on price of the book while a library book (lib:book) might not have that information, but it is likely that both have the book's name and author. By mapping/-combining both entities, that is by stating relationships between them, such as bkst:book is-a lib:book we adapt information in different contexts, answer more complex questions and leverage heterogeneous data to build more dynamic applications. Our work's domain is trust in Web resources, reputation and the generation and origin of trust for resources in our system; because of this we mapped our ontology to the PROV-O ontology, the W3C recommendation. We give an overview of PROV-O and go a little bit into more detail on how it maps to our ontology in the next subsection. Ontologies make use of different formats to be defined and serialized. One format is RDF/XML but others exist that have different advantages. In the figure below you can see the example of one of our classes in the more readable Manchester syntax [Horridge and Patel-Schneider, 2009].

```
Class: carp:Resource

    Annotations:
        rdfs:label "Resource"@en

    SubClassOf:
        carp:hasStartDate some xsd:dateTime,
        carp:hasEndDate some xsd:dateTime,
        carp:hasQualityRating some xsd:double,
        carp:hasKeyword some xsd:string,
        prov:Entity,
        carp:hasRelevanceRating some rdfs:Literal,
        carp:hasName some xsd:string
```

Figure 2.2: The Resource Class from our ontology in Manchester syntax.

2.5 PROV-O

The PROV-O ontology [Lebo et al., 2013] is the most relevant ontology for our work. It is a World Wide Web Consortium (W3C) recommendation and represents the Provenance Data Model [Moreau and Missier, 2013], a conceptual model that captures information about entities, activities and people, and how they relate to the creation of pieces of data (their provenance or origin).

The information captured by PROV-O can be used to assess trustworthiness (if we know the origin we can make a decision whether to trust the source). Since it is a W3C recommendation it is likely to be adopted by the community at large. Our ontology extends the PROV-O ontology (you can find the ontology in Appendix A). Our users (Person in the ontology) are an extension of Prov's "Agent", our endorsement process an extension of

Prov's “Activity” and our Resources (Web Resources) an extension of PROV-O's “Entity”.

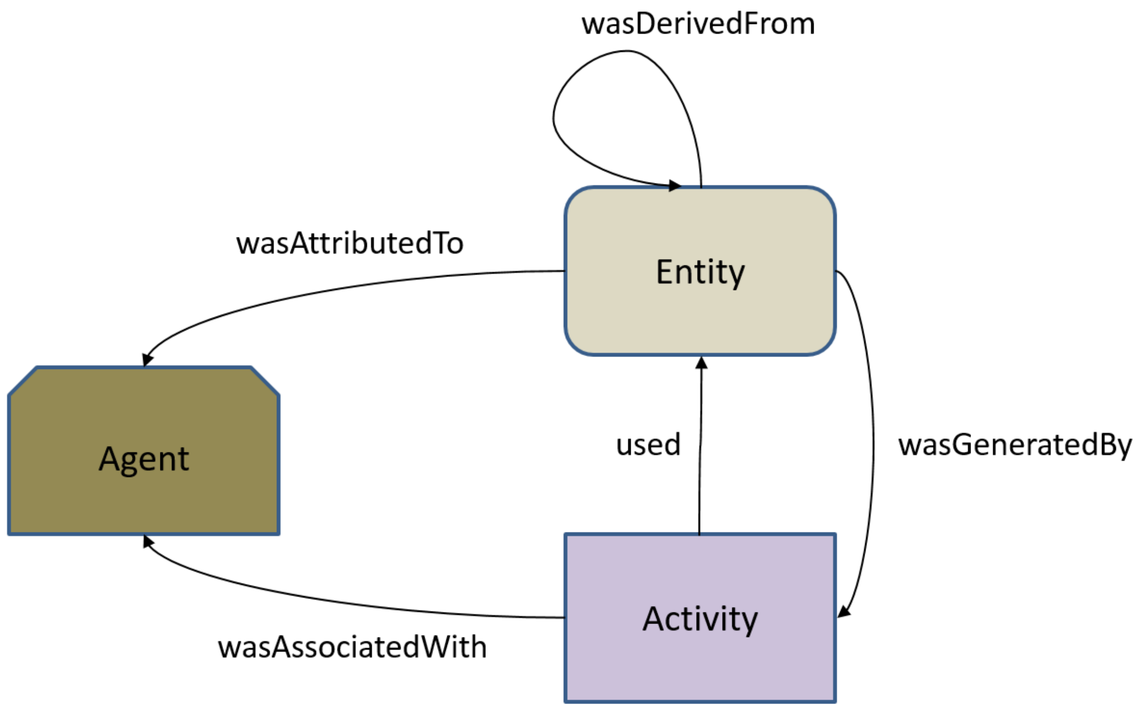

Figure 2.3: A simple representation of the PROV-O ontology modified from [Lebo et al., 2013] that illustrates the three main components of PROV-O: Agents, Entities and Activities.

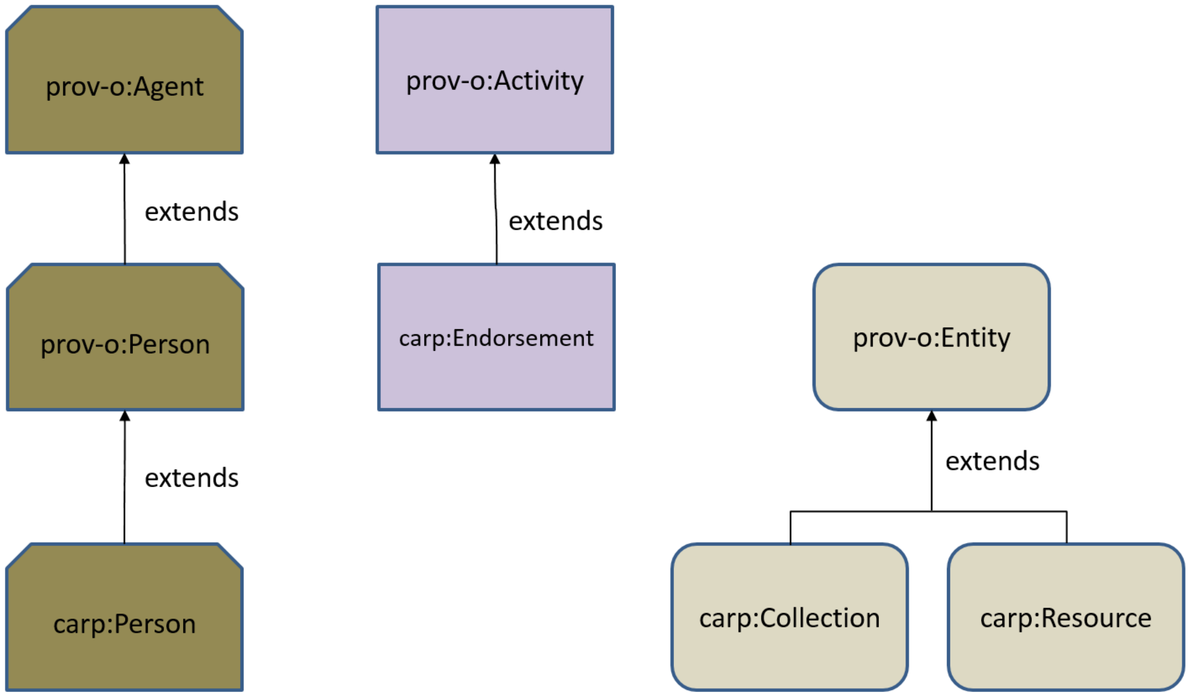

Figure 2.4: Classes in the ontology used in this research project were mapped to PROV-O ontology classes.

In this chapter we provided an overview of the background work on which our algorithm is based upon. Armed with this knowledge we define and go into details on our proposed algorithm in the next chapter.

Chapter 3

Trust Algorithm

The previous chapter provided an overview of the background and basic concepts on which our algorithm is based. This chapter describes our algorithm, starting with an example to develop intuition and proceeding into a proper definition.

The proposed trust algorithm is based on user credentials with dynamically adjusted weights and a history of human endorsements tied to specific contexts, that is, different areas of study. In this manner, the values that affect the trustworthiness of a Web resource shared by a user with expertise in Geology does not change the calculation of trustworthiness of Web resources shared by the same user in the field of Computer Science. As previously mentioned, the trust algorithm was implemented inside the VLC. We will use the VLC context to better describe the algorithm using a specific scenario. The credentials for the user are elicited upon account creation and the system provides the framework to keep track of user interactions across knowledge domains. These interactions capture the user's subjective measure of trust out of which we hypothesize we can create an accurate trust metric. These interactions take the form of a unary decision of whether the user "endorses" the Web resource, meaning they trust the accuracy of its content. A history of endorsements is kept for every Web resource and is combined with the user credentials to generate the trust metric. The trust algorithm proposed in this thesis will be described at three different stages of user interaction with the VLC: 1) a user shares a Web resource, 2) a user endorses a Web resource, and 3) a user searches and finds trustworthy Web resources.

3.1 Intuitive Example

The intuition behind the proposed algorithm is explained using a simple example based on the credentials and interactions of a user with the VLC system in three different stages defined in the following subsections.

3.1.1 Scenario 1. User shares a Web resource

Consider two users of the VLC system, Paul, a Geology professor, and Simon, a Geology student. Both of them sign up to VLC and provide their credentials, which includes their title and discipline. Both start sharing Web resources in the system (Figure 3.1).

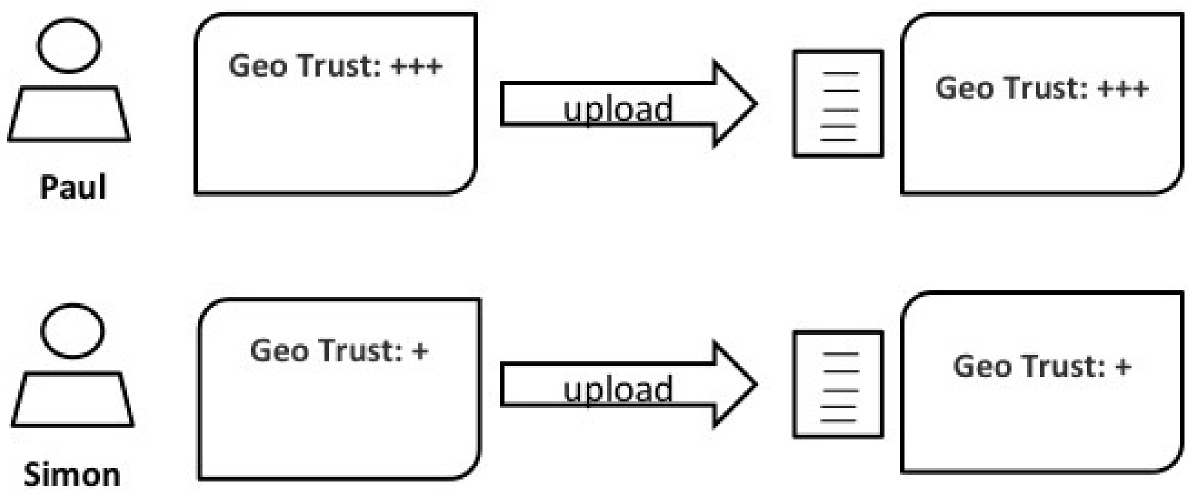

Figure 3.1: Both Paul and Simon share Web resources in the VLC system. Paul and Simon have different values in Geology trust to reflect the impact of having different credentials in a specific discipline, i.e., Paul is a professor and Simon is a student.

3.1.2 Scenario 2. A user endorses a Web resource in the VLC system

A third user, Ariel enters the VLC system. Ariel is a Computer Science master's student who is currently researching Geology material in preparation for an upcoming project. He searches the system with the "Geology" tag in the system. The VLC system displays the most trustworthy Web resources at the top of the list, therefore, Paul's contributions are displayed first (Figure 3.2).

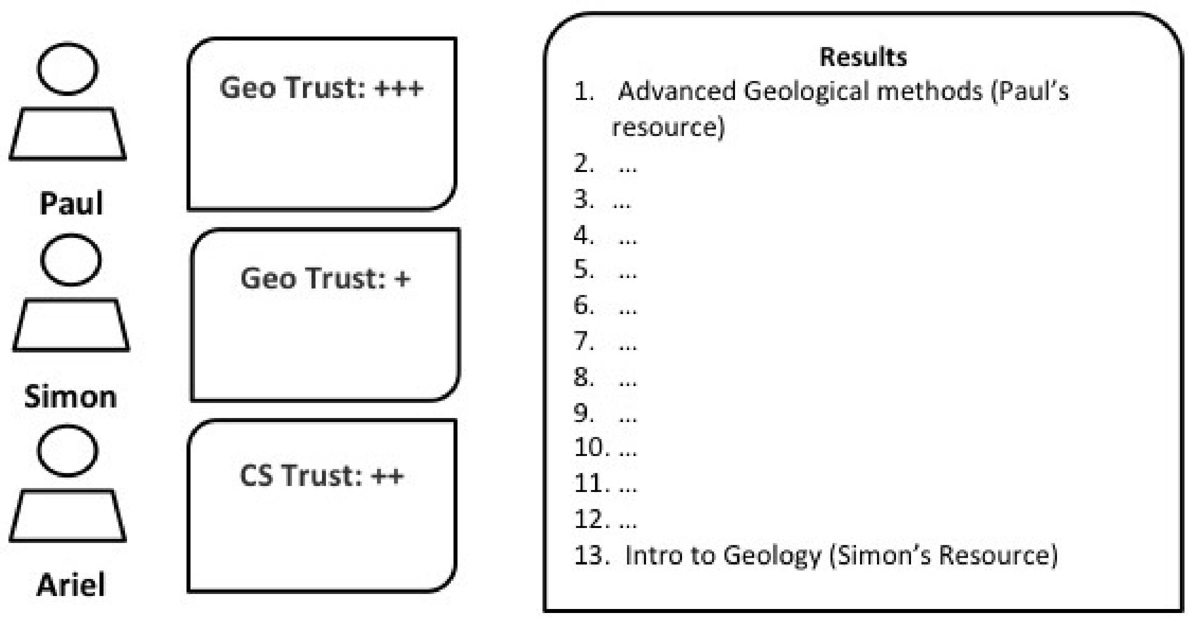

Figure 3.2: The results page shows Paul's Web resource in the first few hits and is likely to be found. Notice also that Ariel's trust credentials are within the field of Computer Science area and not Geology.

Ariel reviews the Web resources returned by querying the system for Geology-related Web resources but none of them contains the information that he was looking for. Later that day Paul reviews the Web resources shared by his students, one called "Introduction to Geology for other disciplines" (the one shared by Simon) catches his eye. Paul assesses

the information contained in the Web resource and determines that it is correct and useful for novices in the domain so Paul decides to endorse it (Figure 3.3). Paul's endorsement causes the system to transfer a portion of his trust (that is, initially weights more due to his professor credentials) onto the Web resource. This makes it more trustworthy, which in turn promotes Simon's Web resource to show up near the top in searches with the tag “Geology”.

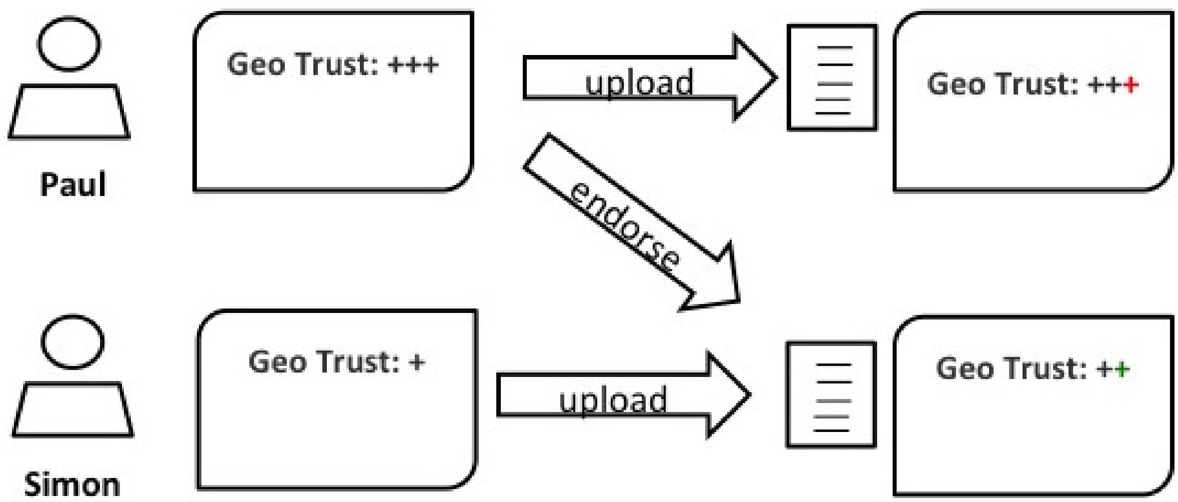

Figure 3.3: Paul finds a Web resource shared by Simon very useful and decides to endorse it. This action causes the system to transfer some of the Paul's trust to Simon's Web resource. This makes it more trustworthy, denoted by the red plus sign, which in turn will make said resource to appear closer to the top of the list of Web resources tagged with “Geology”.

Scenario 3. User searches and finds trustworthy Web resources

The next day Ariel searches the system again for Geology-related resources and this time Simon's “Introduction to Geology” document (the Web resource endorsed by Paul) appears close to the top of the list of results (Figure 3.4). This Web resource contains the information that Ariel was looking for and he uses it to learn in preparation for his new

project.

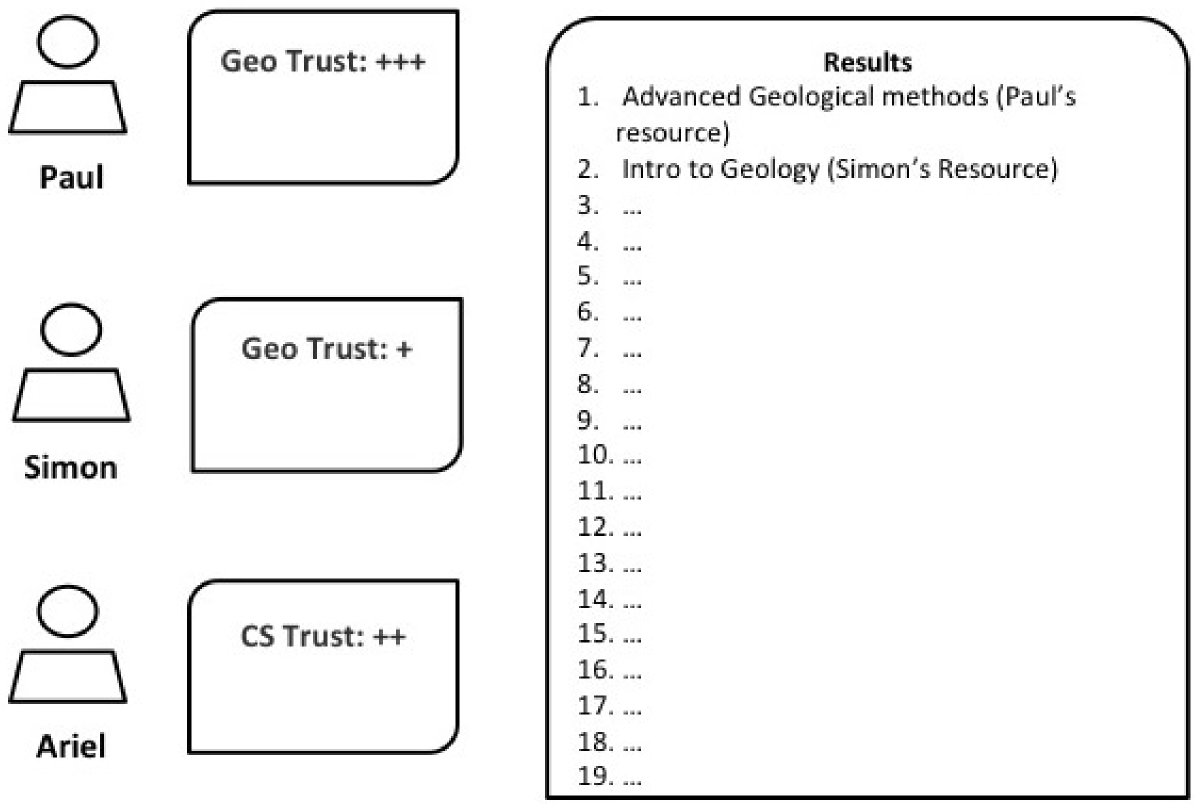

Figure 3.4: The results page shows both Paul and Simon's resource in the first few hits therefore Ariel finds it in his second search.

Example Analysis

The above workflow illustrates the common use of the system. Two users create accounts, Paul and Simon, a Doctor in Geology and an undergraduate in the same field respectively. When they each share Web resources the system knows that given Paul's credentials, his opinion of the field of Geology is more informed than Simon's, so it assigns a proportionally large amount of trust to Paul's Web resources compared to Simon's. This trust difference is what drives Paul's Web resources higher up the list. But when Paul decides to endorse one of Simon's Web resources the system transfers some of the trust it has for Paul onto

Simon's Web resource. Now intuition would tell us that since Paul trusts Simon's Web resource, he (Simon) is more knowledgeable of the field than previously thought and we can give more value to his opinion. The general intuition behind our algorithm is that the trust in a user and the Web resources he or she has shared in the VLC system is dependent on the credentials the user has and on the number and quality of the endorsements received. Also one of our assumptions is that credentials are domain-dependent, therefore we would trust Ariel's opinion in “Computer Science” more than we trust Paul's who has a higher classification but in the “Geology” field. Building on the intuition developed by the three presented scenarios we define some concepts and formalize the algorithm. The concepts are:

1. Web resources - A Web resource is any reference to a website or file over the web. Note that the VLC system enables a user to share a resource on the system and provides a URL to access the resource.

2. Users - A user is any person that uses the VLC system. A user can provide ratings or metadata on Web resources.

3. Tags - Each tag represents a field of study or a specialization of one.

In our algorithm trust is transferred from users to Web resources and vice versa. Trust is transferred through endorsements from a user to a Web resource and from the Web resource to its sharer. Trust is updated every time a new endorsement is made. In the case of users their trust on a particular tag is calculated twofold. One value is based on their credentials and the other based on the value accumulated from their endorsements. The final trust score of a Web resource is an aggregation of both values. For calculating trust based on the credentials, the system multiplies a value representing their academic level (1- no classification, 2- undergraduate student, 3 - masters student, 4- Ph.D. student, 5- professor) times the average endorsement value on the system. This gives the credentials a variable weight and prevents the credentials from dominating or being insignificant for the final calculation.

The second part is the aggregation of the trust of the Web resources he has shared on the same tag. For example, if a user shared three Web resources, one tagged as Geology with a value of 2 and all of them tagged as Computer Science with a value of 1, the sharer will get a trust of 2 in Geology and of 3 in Computer Science.

In order to determine the trust of Web resources we must first take a few things in consideration. In our algorithm, a sharer implicitly trusts the shared resource. The new trust of the endorsement is transferred to the endorsed Web resource and to the Web resources sharer (unless the endorser and sharer are the same user). Afterwards the trust is propagated to each of the uploader's endorsements that update the trust of connected Web resources and then to their sharers, etc.

Every endorsement made by a user has the same trust value, thus when creating a new endorsement the first thing we must do is obtain the user's (endorser) trust and recalculate the trust of existing endorsements when adding an extra one.

To update preexisting endorsements we add the difference between the previous endorsement value and the new endorsement value to the trust of the Web resource.

Readers should note that trust value is dependent on the tag, therefore the process must be repeated for every tag shared by the user and Web resource. Making an endorsement only redistributes the existing trust; the only way that the total trust in the system can increase is if a user get new credentials (e.g. a Ph.D. student becomes a professor) or a new user enters the system. Also worth noting is the fact that we keep track of the visited entities (i.e. Web resources and users) to prevent cycles. Since every endorsement shares the endorser's trust equally the value of individual endorsements naturally decreases over time as the user makes new endorsements and as a consequence the trust in Web resources and their sharers/uploaders also decreases. Although every endorsement is susceptible to this decrease, our assumption is that in general the ratio of trust between resources should maintain stable and thus the system where we implement the algorithm would remain useful in determining trust of a user/Web resource relative to others. One last thing to note is that our algorithm closely resembles the traversal of an n-ary tree. Our root node

would be the endorser and we must recursively traverse and update his endorsed resources then move to the users that shared those Web resources, then on to their own endorsements and so on. Therefore time complexity of our algorithm is O(n), where n is the number of nodes we must traverse and can be no greater than the number of nodes in the system.

3.2 Entities and Concepts

Definition 1 A Web resource is any reference to a website or file over the Web. Web resources can be documents, video, presentations, etc.

$$r \in R \tag{3.1}$$

Where r is a particular Web resource in the system. It is the trust the Web resource has gained on the system's tags due to the history of its endorsements and R is the set of all Web resources in the system

Definition 2 A collection is an aggregation of Web resources. It is a way to group Web resources that might or might not be related.

$$c \in C \tag{3.2}$$

Where c is a particular collection on the system and C is the set of all collections on the system.

Definition 3 Time in our algorithm is handled in an incremental, discrete manner and it does not refer to the actual time and date but rather to the iterations of the system. In this fashion t_1 refers to the first time the algorithm has executed and t_2 refers to the 2nd time it has run, independent of the time between each particular run. Current time will be referred in this document as t_k where the previous time would be represented as $t_{(k-1)}$ and the next time as $t_{(k+1)}$.

$$i \in I \tag{3.3}$$

Where i is a particular time/iteration of the algorithm and I is the set of all times/iterations in which the algorithm has run.

Definition 4 The system allows Web resources to be tagged. Each of these tags represents a matter of interest, a field of study or a specialization of one. Examples:

- Computer Science
- Machine Learning
- Carlsbad Mountains

Web resources and users (defined below) have copies of T due to the fact that each will have separate trust values for each tag. To denote a copy of the tag set associated with a particular user or Web resource we will write it as T'unless otherwise noted.

$$t \in T, t = (name, value) \tag{3.4}$$

Where T is the set of all tags in the system and t is a particular tag in the system and a tuple of name (a human-readable label of the tag's subject matter) and value (the trust value associated with this tag).

Definition 5 A user is any agent that uses the system. A user can be a person or a program (although in our implementation we only considered human users). A user can upload, read or provide ratings or metadata on Web resources. Although users can be anonymous in this document when we refer to a user we assume it is an authenticated user (with associated credentials) unless stated otherwise.

$$u \in U, u = (cT', hT') \tag{3.5}$$

Where U is the set of all users in the system, u is a particular user in the system and a tuple of cT'and T'. cT'is a copy of the set of all tags in the system with modified trust values. It is the trust the user has gained on the system's tags due to his credentials. hT'is

a copy of the set of all tags in the system with modified trust values. It is the trust the user has gained on the system's tags due to the history of his interactions.

Definition 6 Trust in a Web resource is defined as:

$$trust(resource, tag, time) = \sum_{w=1}^{uert} endorsement(w, tag, time - 1) \tag{3.6}$$

Where uetr is the # of users that endorsed the tag on the Web resource

Algorithm 1 Web resource trust algorithm

```
procedure getTrust(resource)
    trust = 0;
    for each endorsement ∈ resource.endorsements do
        trust+=endorsement.trust;
    end for
end procedure
```

Definition 7 The trust in a user is a combination of his credentials and the aggregation of the trust on his Web resources (the ones he has shared).

$$trust(user, tag, time) = credentials(user, tag, time) + \sum_{k=1}^{rust} trust(k, tag, time) \tag{3.7}$$

Where rust is the # of Web resources uploaded by user with tag.

Definition 8 An endorsement is how much of a user trust on a particular tag is transferred to a Web resource.

$$endorsement(user, tag, time) = \frac{trust(user, tag, time)}{numOf(endorsements, user, time)} \tag{3.8}$$

Definition 9 The user credentials is the numeric trust value we assign to a user based on their official knowledge on the field (Ph.D, MS, Bachelors). Rank has a numeric representation for the aforementioned field (MS=3,PhD=4,etc).

$$credentials(user, tag) = rank(user, tag) * basicTrustUnit(tag) \tag{3.9}$$

Algorithm 2 User trust algorithm

```
procedure GETTRUST(user)
    trust = user.getCredentials();
    for each upload ∈ user.uploads do
        trust+=upload.trust;
    end for
end procedure
```

Definition 10 The BTU (basic trust unit) is the average trust that a single endorsement can provide in the system for a particular tag. This is the element that provides the dynamic credentials element to the algorithm. Adding this element to the credentials calculation is important because if credentials are kept static their value can disproportionally determine the trust of the endorsed Web resources or on the other hand be so small as to render any user credential insignificant.

$$\frac{average(trust(r_1, tag), ..., trust(r_{numOfResources}, tag))}{average(numberOfEndorsements(r_1, tag), ..., numberOfEndorsements(r_{numOfResources}, tag))} \tag{3.10}$$

3.3 Endorsement Algorithm

From the example we can observe that there are two basic actions in the system that implements the proposed algorithm: share and endorse. We assume that if a user shares a Web resource (the same as with anything), he trusts the resource. It is for this reason that in our system a share is an implicit endorsement from the sharer. In this manner the whole trust process revolves around the endorsement process, specifically when a new endorsement is made either by users other than the sharer or by the sharer when the Web resource is entered into the system. Please see figure 3.5 for a graphical representation of the endorsement process and algorithms 3 and 4 for the pseudocode.

In this chapter we defined our algorithm. In the following chapter we talk about the

implementation we did on the VLC system.

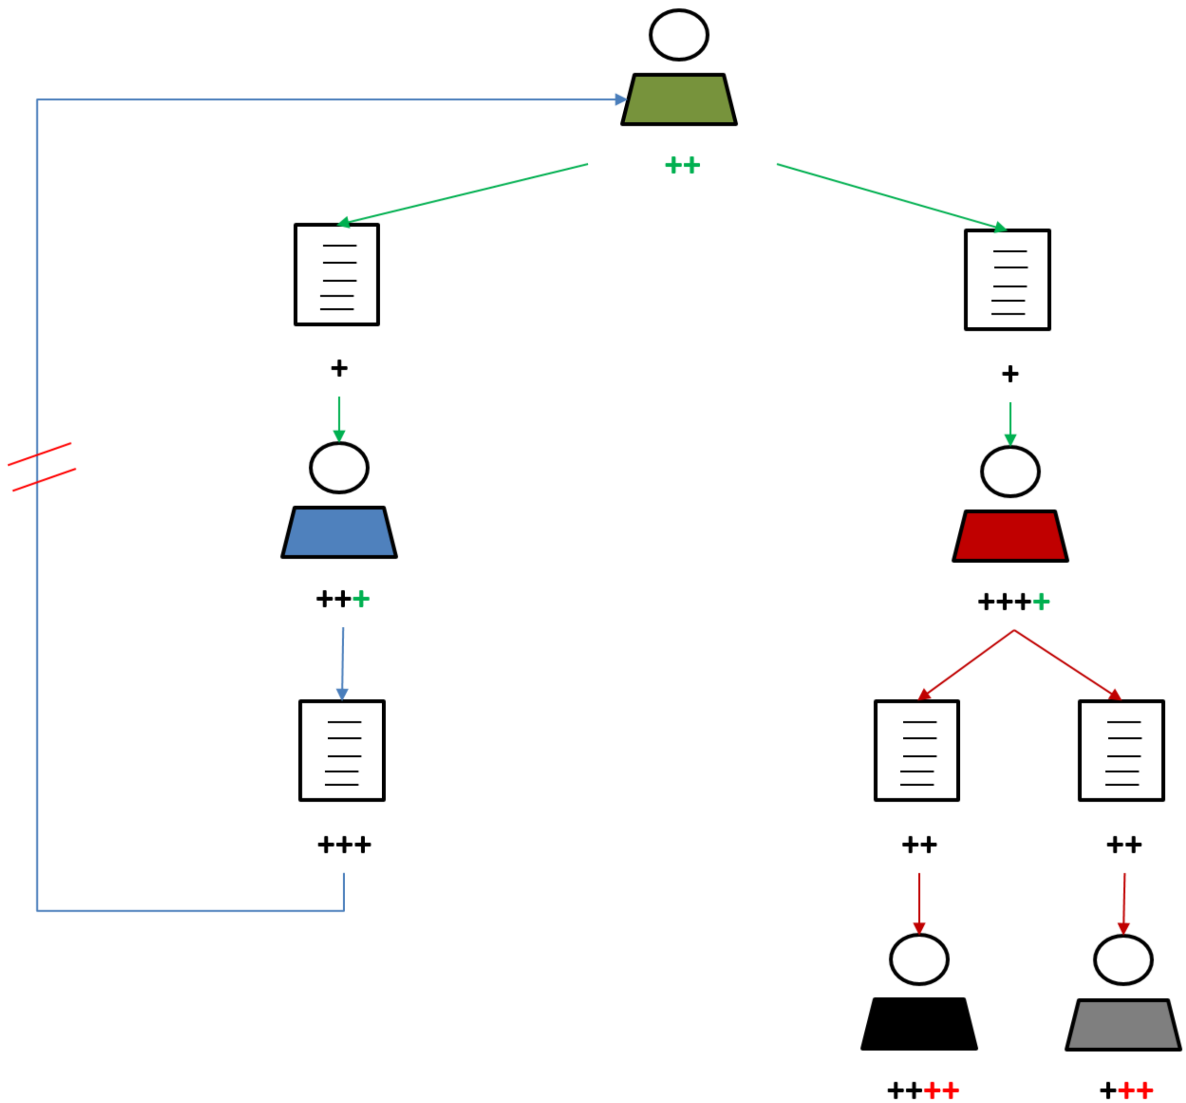

Figure 3.5: A graphical representation of trust propagation in the endorsement process. Trust transfers from users to resources (through endorsements) and from resources to their users (to their sharer, e.g. The green user endorsed a resource shared by the blue user). Trust is divided equally among endorsements and transferred directly to the original sharer. As you can see we keep track of visited entities to prevent cycles (the crossed out line).

Algorithm 3 endorse algorithm

```
procedure ENDORSE(user,resource,endorsement)
    user.visited = true;
    endorsement.visited = true;
    oldEndTrust = getMode(user.endorsements);
    newEndTrust = user.getTrust()/count(user.endorsements);
    newEndUpdate = abs(oldEndTrust - newEndTrust);
    oldEndUpdate = newEndTrust - oldEndTrust;
    resource.trust += newEndTrust;
    endorsement.trust = newEndTrust;
    endorseRec(resource.owner,newEndUpdate);
    user.visited = true;
    endorsement.visited = true;
    for each e ∈ user.endorsement do
        if(e.visited){continue;}
        e.visited =true;
        e.resource.trust += oldEndUpdate;
        e.trust += oldEndUpdate;
        endorseRec(e.resource.owner,oldEndUpdate);
        e.visited = false;
    end for
    user.visited = false;
    endorsement.visited = false;
end procedure
```

Algorithm 4 endorse recursive algorithm

```
procedure ENDORSEREC(user,trustUpdateForAll)
    if(user.visited){return;}
    user.visited =true;
    trustUpdatePerEndorsement = trustUpdateForAll/count(user.endorsements);
    for each e ∈ user.endorsement do
        if(e.visited){continue;}
        e.visited =true;
        e.trust += trustUpdatePerEndorsement;
        e.resource.trust += trustUpdatePerEndorsement;
        if(!(e.resource.owner.visited)){
        e.resource.owner.trust += trustUpdatePerEndorsement }
        endorseRec(e.resource.owner,trustUpdatePerEndorsement);
        e.visited = false;
    end for
    user.visited = false;
end procedure
```

Chapter 4

Implementation

In the previous chapter we described the algorithm. In this chapter we describe the implementation and design choices associated with it. The algorithm was implemented inside the VLC, a bookmarking system developed by the Cyber-Share group at the University of Texas at El Paso. The system is running on a Drupal installation [Drupal.com, 2015].Drupal is a Content Management System written in PHP. It was selected because it handles basic functionality out of the box such as user accounts and system administration. It runs on top of a WAMP or LAMP server stack (Windows/Linux, Apache, MySQL, PHP). Drupal works through the concepts of modules and hooks. Modules are code packages that alter Drupal's functioning, the way that these modules can interact with Drupal is through the use of hooks which are PHP functions that can be overwritten with a naming convention of the form of {module-name}_{hook-name}.

For this thesis we extended the functioning of the VLC Drupal module, a module created by Aida Gandara that provides the customizations required by the VLC system. We modified the VLC menu hook to add pages to manage the trust system execution and to provide a page for users to endorse Web resources. We also modified the VLC block view to add controls for users to navigate and combine Web resources from different collections.

For the RDF triplestore we used the arc2 database [GitHub, 2015]. We created a new triplestore to save all the trust information pertaining the elements in the system. The table below contains the triples we used as part of our trust algorithm calculation.

Table 4.1: This table presents the structure of triples generated in the VLC to represent user credentials or endorsements.

Triples			
Subject	Predicate	Object	Description
carp:Person	carp:hasAreaOfKnowledge	xsd:string	Defines an area of knowledge for the person
carp:Person	carp:hasInstitution	xsd:string	Defines which institution the person belongs to
carp:Person	carp:hasTitle	xsd:string	Defines the academic title of the person (e.g. masters,phd,etc.)
carp:Person	carp:hasAreaOfKnowledge	xsd:string	Defines an area of knowledge for the person
carp:Resource	carp:hasKeyword	xsd:string	Defines an area of knowledge for the Web resource
carp:Endorsement	carp:hasKeyword	xsd:string	Defines an area of knowledge for the endorsement
carp:Person	carp:hasRank	xsd:string	Defines an area of knowledge for the person
carp:Resource	carp:isPartOf	carp:Collection	Establishes a composition relationship between Collection and Web resource
carp:Resource	carp:hasUploader	carp:Person	Establishes the upload relationship between Web resource and Person

carp:Endorsement	carp:hasEndorsee	carp:Resource	Establishes the endorsement relation between a particular endorsement and a resource being endorsed
carp:Endorsement	carp:hasEndorser	carp:Person	Establishes the endorsement relation between a particular endorsement and the Person doing the endorsement
carp:Resource	carp:isPartOf	carp:Collection	Establishes a composition relationship between Collection and Web resource
carp:Any	carp:hasValue	xsd:double	Establishes a composition relationship between a Collection and Web resource

The VLC's architecture contains two main components, a Google Chrome browser extension and a server implementation on top of Drupal. Both elements follow our ontology (see Appendix A) to store and manage their data and as previously mentioned in the background section, and by extension our data, maps to the PROV-O ontology.

4.1 Browser extension

In order to share Web resources to the system a Chrome browser extension was developed. Chrome was selected due to its ubiquity. The browser extension provides an interface to bookmark Web resources. The extension makes REST calls to an endpoint to create triples that describe a Web resource, these triples are stored in the arc2 triple store.

The extension was developed by Matthew Rister on top of work by Ariel Garcia. For this thesis we made modifications to the extension to match our ontology.

Figure 4.1: Chrome's extension interface to enter user credentials.

Figure 4.2: Chrome's extension interface to bookmark Web resources.

4.2 Drupal Module

The server component provides a Web-based interface that allows account creation (figure 4.4), the visualization of Web resources and handles the endorsement process (i.e. the execution of our algorithm, figure 4.5). Since the VLC was a pre-existing system we were able to leverage existing code to achieve part of the functionality of the web interface (i.e. some of the functions necessary to manage and query the database). This code was developed by Aida Gandara. In addition to provide a foundation for this part of the implementation, her Dissertation [Gandara, 2013] provided an inspiration for our semantic components, i.e. our ontology resuses some of the CARP ontology concepts, the carp short-hand (table 4.1) in our ontology is actually a reference to her Collect-Annotate-

Rene-Publish Methodology (CARP).

In order to verify the implementation of the system a preliminary test was conducted with manually entered data and manually generated interactions (see figure 4.3).

iteration	element	value
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_1_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#endorsement_55c9210cc6ca7	5
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_1_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#paul	5
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_1_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#endorsement_55c921f483f05	5
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_1_record	http://advgeo.com/methods	5
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_2_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#endorsement_55c922c6aaae5	2
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_2_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#simon	2
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_2_record	http://advgeo.com/intro	2
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_3_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#paul	5
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_3_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#endorsement_55c922c6aaae5	2
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_3_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#simon	2
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_3_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#endorsement_55c92347b1382	2.5
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_3_record	http://advgeo.com/intro	4.5
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_3_record	http://ontology.cybershare.utep.edu/CARP/carp.owl#endorsement_55c921f483f05	2.5
http://ontology.cybershare.utep.edu/CARP/carp.owl#iteration_3_record	http://advgeo.com/methods	2.5

Figure 4.3: Triplestore record of algorithm execution with a testing example.

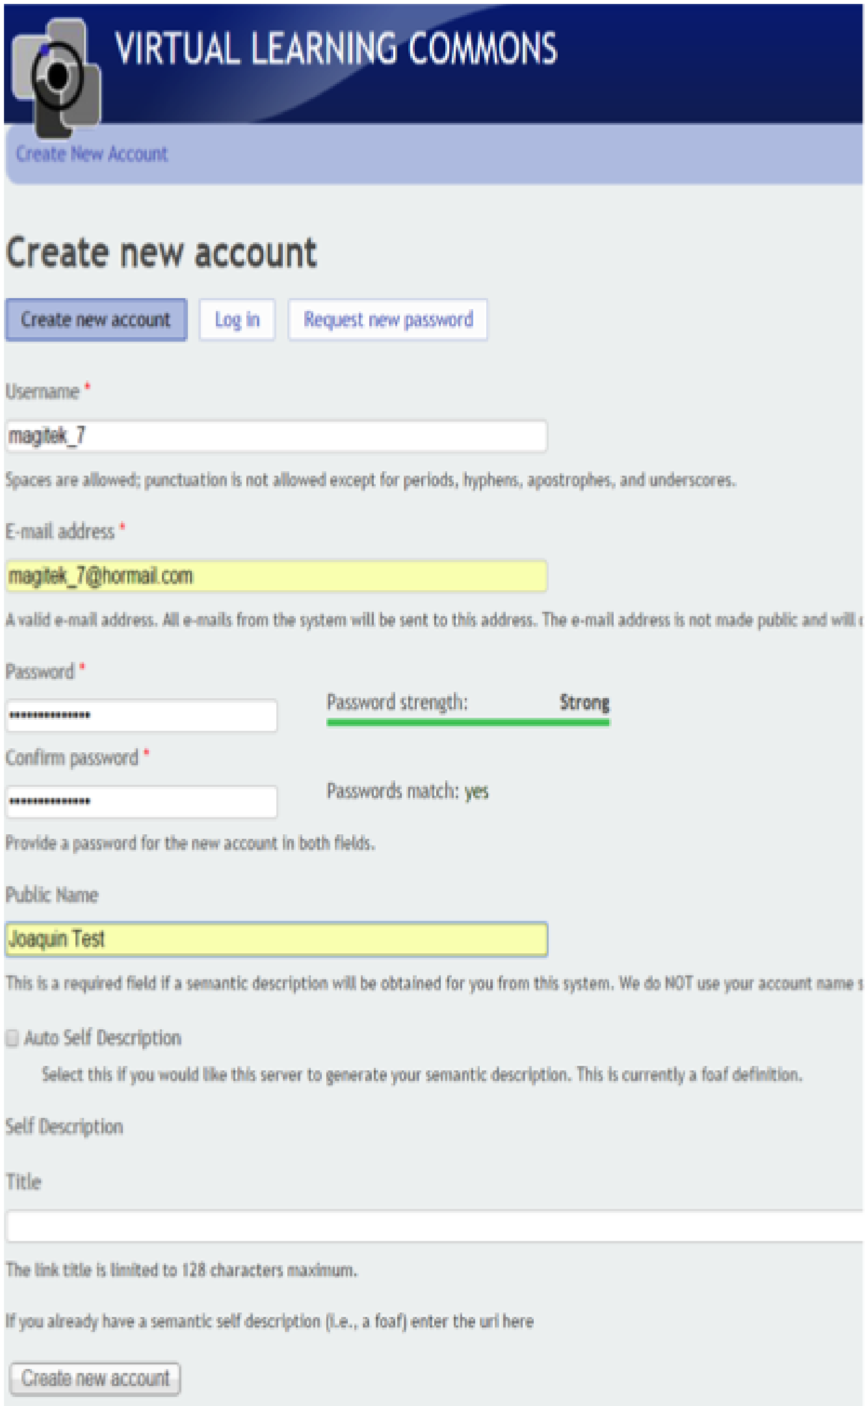

Figure 4.4: The interface to create a new user account was modified from the default to save credential information in the triplestore.

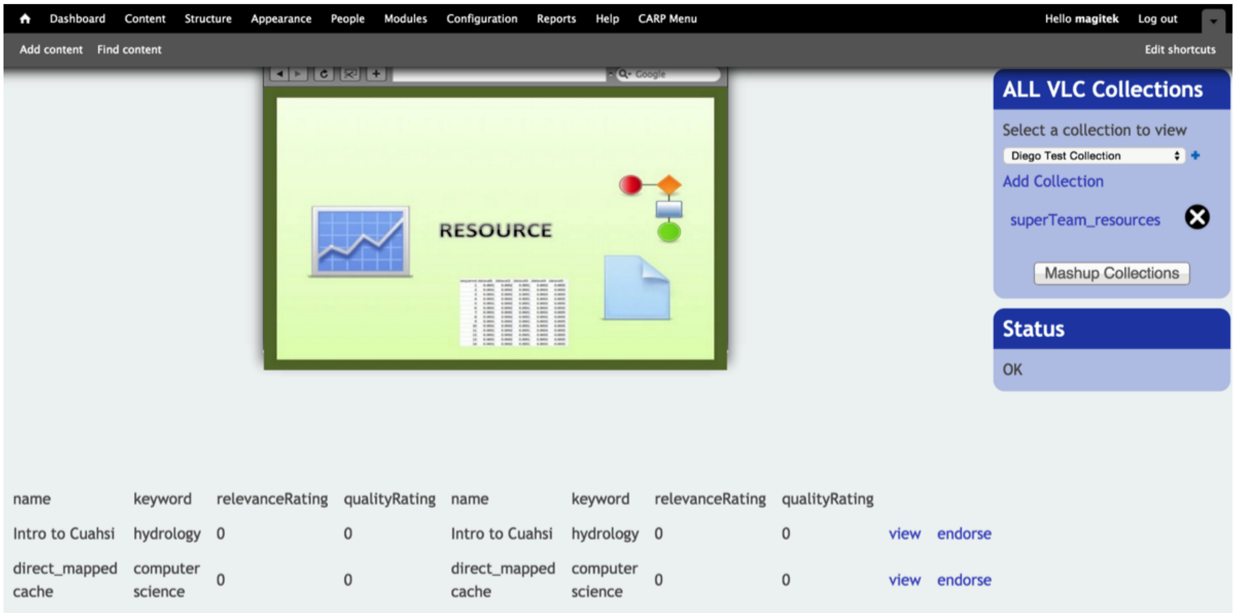

Figure 4.5: The Mashup Overview page. Through this page users can see an aggregate view of Web resources from different collections. The page also provides controls for viewing the Web resource and endorsing it.

The image shows the trust corresponding to the user credentials, the endorsements and the Web resources. "advgeo.com/methods" is the URL of Paul's Web resource and "advgeo.com/intro" is the URL of Simon's Web resource from our previous example. The first iteration corresponds to Paul sharing his Web resource, the second to Simon's sharing and the third to Paul's endorsement. The results of our test yielded the expected results; the endorsement process allows the trust to be redistributed and Web resources shared by users with lower credentials can attain higher trust values in the system (notice how Simon's Web resource gains the highest Web resource trust in the system). We can then use this information to sort it higher in a list or filter results.

In this section we described the implementation of the system (code can be found at http://git.cybershare.utep.edu/agandara1/vlc.git). On the next section we describe the testing scheme and the test results.

Chapter 5

Evaluation

We tested the capabilities of the algorithm by simulating an online environment where users endorsed resources pre-classified by an expert as trustworthy or untrustworthy. In order to test the algorithm and not the proficiency of our users in the classification of Web resources three additional algorithms were implemented:

1. Simple Voting - This is what we consider the base algorithm. A constant value (in our case 1) is added to the trust value of every Web resources with every endorsement.

2. Algorithm w/o Credentials - This variation of our algorithm is equivalent in functioning (propagation of trust) to our full algorithm but the credentials of the user are not considered.

3. Algorithm w/o Dynamic Credentials - This variation of our algorithm is equivalent in functioning (propagation of trust) to our full algorithm but the credential value assigned to users depends on their academic classification (undergraduate student, professor, etc.) and remains constant.

4. Full Algorithm - Our proposed algorithm. This algorithm propagates trust based on tags, making use of credentials in a dynamic way (uses the concept of BTU).

Participants endorsed the Web resources and we applied the four algorithms to the full set of interactions. We then and computed a ranking of Web resources based on each algorithm output. Depending on the Web resource ranking we classified them as trustworthy or not and this classification was compared with the expected/expert classification. We expect that with more participant's input (more endorsements) more trustworthy Web resources

will be correctly classified as trustworthy and less non-trustworthy will be incorrectly classified as trustworthy. We expect this effect to improve accuracy for all algorithms with our full algorithm taking the lead. We collected data from two fields of study, Computer Science and Geology. An expert from each of these fields compiled two lists of online Web resources: five trustworthy Web resources and five non-trustworthy Web resources. Participants were recruited from these two fields, 15 from each field for a total of 30 participants. We directed participants to a website showing the ten Web resources from their field and the ten resources from the other field, arranged randomly, and with controls to view and endorse each one. Participants were instructed to review the Web resources and endorse the ones that they considered trustworthy. The endorsement took the form of a unary decision. The UI featured a link next to each Web resource and participants clicked those that they choose to endorse as trustworthy. After each participant provided input, their endorsements were added to a cumulative set of endorsements across participants. We then sorted the Web resources by computed trust values and classified the first half to be trustworthy (partitioning value = 0.50) and the last half to be untrustworthy, preserving each classification result through time as the number of participants grew. Each classification result was compared to the experts classification, and the number of false positives/negatives and true positives/negatives was computed. These values were analyzed through time as a function of number of endorsements. The analysis was recomputed with another partitioning value (i.e. 0.20).

Participants signed a consent form prior to participation (see appendix B for IRB approval) and then were directed to a computer set up with the test website. They created an account on the system with randomly assigned IDs and provided demographic/academic information such as Title (e.g. Undergraduate Student, Professor), primary and secondary areas of knowledge (e.g. Computer Science, Geology), institution (e.g. UTEP) and position (e.g. Research Assistant, Associate Professor). No information that would identify a specific individual was collected.

Since we had ten trustworthy Web resources and ten non trustworthy (five of CS and

five of Geology), a partition value of 0.5 corresponds to a 100% recall (half of the resources are trustworthy if we take the upper half of the list we are selecting 5/5), a partition of 0.2 would correspond to a 40%.

5.1 Results

Tables 5.1, 5.2, 5.3, and 5.4 shows results obtained by applying the four algorithms (previously described) to the data entered by the participants of the evaluation study. The first two tables represent the results for the Computer Science use case and tables 5.3 and 5.3 list the results for the Geology use case. Tables 5.5 and 5.6 show the confusion matrices for both use cases. Figures 5.1 and 5.2 show the comparison between the accuracy of the two best performing algorithms, namely simple voting and our own algorithm.

Table 5.1: Precision with 0.5 cutoff (100% recall). In this table we see the precision for the Computer Science use case. The precision is seen as a function of the number of endorsements.

Algorithm Comparison				
# of Endorsements	Simple Algorithm	Algorithm w/o Credentials	Algorithm w/o dynamic credentials	Full Algorithm
20	1	0.4	0.2	0.2
40	1	0.8	0.6	0.8
60	1	0.8	0.8	1
80	1	0.8	0.8	1
100	1	0.8	0.8	1
120	1	0.8	0.8	1
140	1	0.8	0.8	1
160	1	0.8	0.8	1

Table 5.1 shows the results of applying the four tested algorithm to endorsements of Computer Science Web resources made by the study participants. The table shows the number of endorsements used to calculate the precision of the algorithm with a 100% recall. For example, with 20 endorsements, the Simple Algorithm has a precision of 1, which means that the Simple Algorithm accurately classified five (all) of the trustworthy resources as trustworthy based on user endorsements. For the same 20 endorsements, the Full Algorithm only classified one trustworthy resource as trustworthy. The 0.5 cutoff indicates that we consider the top five resources, that is those that have highest trust value (as produced by each particular algorithm) as trustworthy. Since we have ten resources of each field, a cutoff of 0.5 (50% of the resources) is equivalent to five resources, given the fact that the number of actual trustworthy resources is also five we have a 100% recall. This table shows that the Simple Algorithm has a 100% precision with any number of endorsements used in calculations. Our proposed algorithm reaches the 100% precision after 60 endorsements. The algorithm without credentials and the algorithm without dynamic credentials reach only 80% precision on 40 and 60 endorsements respectively.

Table 5.2: Precision with 0.2 cutoff (40% recall). In this table we see the precision for the Computer Science use case. The precision is seen as a function of the number of endorsements.

Algorithm Comparison				
# of Endorsements	Simple Algorithm	Algorithm w/o Credentials	Algorithm w/o dynamic credentials	Full Algorithm
20	1	0.5	0.5	0.5
40	1	0.5	0.5	1
60	1	0.5	0.5	1
80	1	0.5	0.5	1
100	1	0.5	0.5	1
120	1	0.5	0.5	1
140	1	0.5	0.5	1
160	1	0.5	0.5	1

Table 5.2 shows the results of applying the four tested algorithm to endorsements of Computer Science Web resources made by the study participants. The table shows the number of endorsements used to calculate the precision of the algorithm with a 40% recall. Using the same example, with 20 endorsements, the Simple Algorithm has a precision of 1, which means that the Simple Algorithm accurately classified two (all given the recall) of the trustworthy resources as trustworthy based on user endorsements. For the same 20 endorsements, the Full Algorithm only classified one trustworthy resource as trustworthy. The 0.2 cutoff indicates that we consider the top two resources, that is those that have highest trust value (as produced by each particular algorithm) as trustworthy. Since we have ten resources of each field, a cutoff of 0.2 (20% of the resources) is equivalent to two resources, given the fact that the number of actual trustworthy resources is five we have

a 40% recall (2/5=.4). This table shows that the Simple Algorithm has a 100% precision with any number of endorsements used in calculations. Our proposed algorithm reaches the 100% precision after 40 endorsements. The algorithm without credentials and the algorithm without dynamic credentials reach only 50% precision with any number of endorsements used in the calculations.

Table 5.3: Precision with 0.5 cutoff (100% recall). In this table we see the precision for the Geology use case. The precision is seen as a function of the number of endorsements.

Algorithm Comparison				
# of Endorsements	Simple Algorithm	Algorithm w/o Credentials	Algorithm w/o dynamic credentials	Full Algorithm
20	0.2	0	0.2	0.2
40	0.4	0.2	0.2	0.2
60	0.6	0.6	0.6	0.6
80	0.4	0.4	0.2	0.4
100	0.6	0.4	0.4	0.6
120	0.6	0.4	0.6	0.6
140	0.6	0.6	0.6	0.6
160	0.6	0.4	0.4	0.6
180	0.6	0.4	0.4	0.6

Table 5.3 shows the results of applying the four tested algorithm to endorsements of Geology Web resources made by the study participants. The table shows the number of endorsements used to calculate the precision of the algorithm with a 100% recall. This table shows that the Simple Algorithm achieves a 60% precision after 60 endorsements (albeit with a dip to 40% at the 80 endorsements mark). Our proposed algorithm reaches the same

precision at the same number of endorsements and presents the same dip in performance at the 80 endorsements mark. The simple algorithm does outperform our own with at low levels of endorsements, specifically, there is a 20% difference at the 40 endorsement mark. The algorithm without credentials and the algorithm without dynamic credentials reach only 60% precision both at the 60 endorsement mark but accuracy is unstable and degrades after that.

Table 5.4: Precision with 0.2 cutoff (40% recall). In this table we see the precision for the Geology use case. The precision is seen as a function of the number of endorsements.

Algorithm Comparison				
# of Endorsements	Simple Algorithm	Algorithm w/o Credentials	Algorithm w/o dynamic credentials	Full Algorithm
20	0	0	0	0
40	0.5	0	0	0.5
60	1	0.5	0.5	1
80	0.5	0.5	0.5	0.5
100	1	0.5	0.5	0.5
120	0.5	0.5	0.5	0.5
140	1	0.5	0.5	1
160	0.5	0.5	0.5	1
180	1	0.5	0.5	0.5

Table 5.4 shows the results of applying the four tested algorithm to endorsements of Geology Web resources made by the study participants. The table shows the number of endorsements used to calculate the precision of the algorithm with a 40% recall. This table shows that the Simple Algorithm has a precision that fluctuates between 50% and

100% for all number of endorsements used in calculations starting with 40. Our proposed algorithm also fluctuates starting with 40 endorsements although with more stability (the algorithm maintained a 50% from 80 to 120 endorsements and 10% from 140 to 160). The algorithm without credentials and the algorithm without dynamic credentials proved to be more stable in this test but only reached 50% precision both after 60 endorsements.

Table 5.5: In this table we see the confusion matrices for our algorithm in the Computer Science use case. Please note that since we are evaluating the accuracy of the algorithm n is the number of endorsements but the classified elements is always 10 (5 trustworthy and 5 non-trustworthy).

Confusion Matrices for Computer Science		
n=20 (100% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	1	4
Actual: Non-Trustworthy	4	1
n=40 (100% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	4	1
Actual: Non-Trustworthy	1	4
n=60+ (100% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	5	0
Actual: Non-Trustworthy	0	5
n=20 (40% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	1	4
Actual: Non-Trustworthy	1	4
n=40+ (40% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	2	3
Actual: Non-Trustworthy	0	5

In table 5.5 we can see the confusion matrices for the Computer Science use case. These

tables provides us with more detailed information on the classification of all resources by our algorithm at different number of endorsement and recall levels. For example, in the case of 20 endorsements and 100% recall. We know from table 5.1 that we had a 20% precision. This translates into one resource predicted as trustworthy that was actually trustworthy, four resources predicted as trustworthy that were actually non-trustworthy, and another five miss-classified resources (four trustworthy resources and one non-trustworthy). In the case of 40% recall we only classify the top two elements as trustworthy, if these two resources are indeed trustworthy we achieve a 100% precision and 50% if only one is correctly classified.In the same example of 20 endorsements we know from table 5.2 that our algorithm had a 50% precision, this translates to one correctly classified trustworthy element and one non-trustworthy incorrectly classified as trustworthy. For the elements predicted as non-trustworthy, three are classified as such by default due to our cutoff selection (we only allow 2 resources to be classified as trustworthy but in reality five are trustworthy), for a total of 4 predicted non-trustworthy (including the one actually miss-classified).

As with previous tables (5.1 and 5.2) notice that in the case of 100% recall our algorithm converges to full precision after 60 endorsements where there is no false positives or negatives. In the case of 40% recall (equivalent to two resources classified as trustworthy) the precision is 100% after 40 endorsements.

Table 5.6: In this table we see the confusion matrices for our algorithm in the Geology use case. Please note that since we are evaluating the accuracy of the algorithm n is the number of endorsements but the classified elements is always 10 (5 trustworthy and 5 non-trustworthy).

Confusion Matrices for Geology		
n=20-40 (100% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	1	4
Actual: Non-Trustworthy	4	1
n=80 (100% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	2	3
Actual: Non-Trustworthy	3	2
n=100+ (100% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	3	2
Actual: Non-Trustworthy	2	3
n=40, n=80-120, n=180 (40% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	1	4
Actual: Non-Trustworthy	1	4
n=60, n=140-160 (40% recall)	Predicted: Trustworthy	Predicted: Non-Trustworthy
Actual: Trustworthy	2	3
Actual: Non-Trustworthy	0	5

In table 5.6 we can see the confusion matrices for the Geology use case. Again notice that

in the case of 100% recall our algorithm converges to 60% precision after 100 endorsements, also notice that the proportions of true positives, false positives, true negatives and false negatives is unstable with a 40% recall.

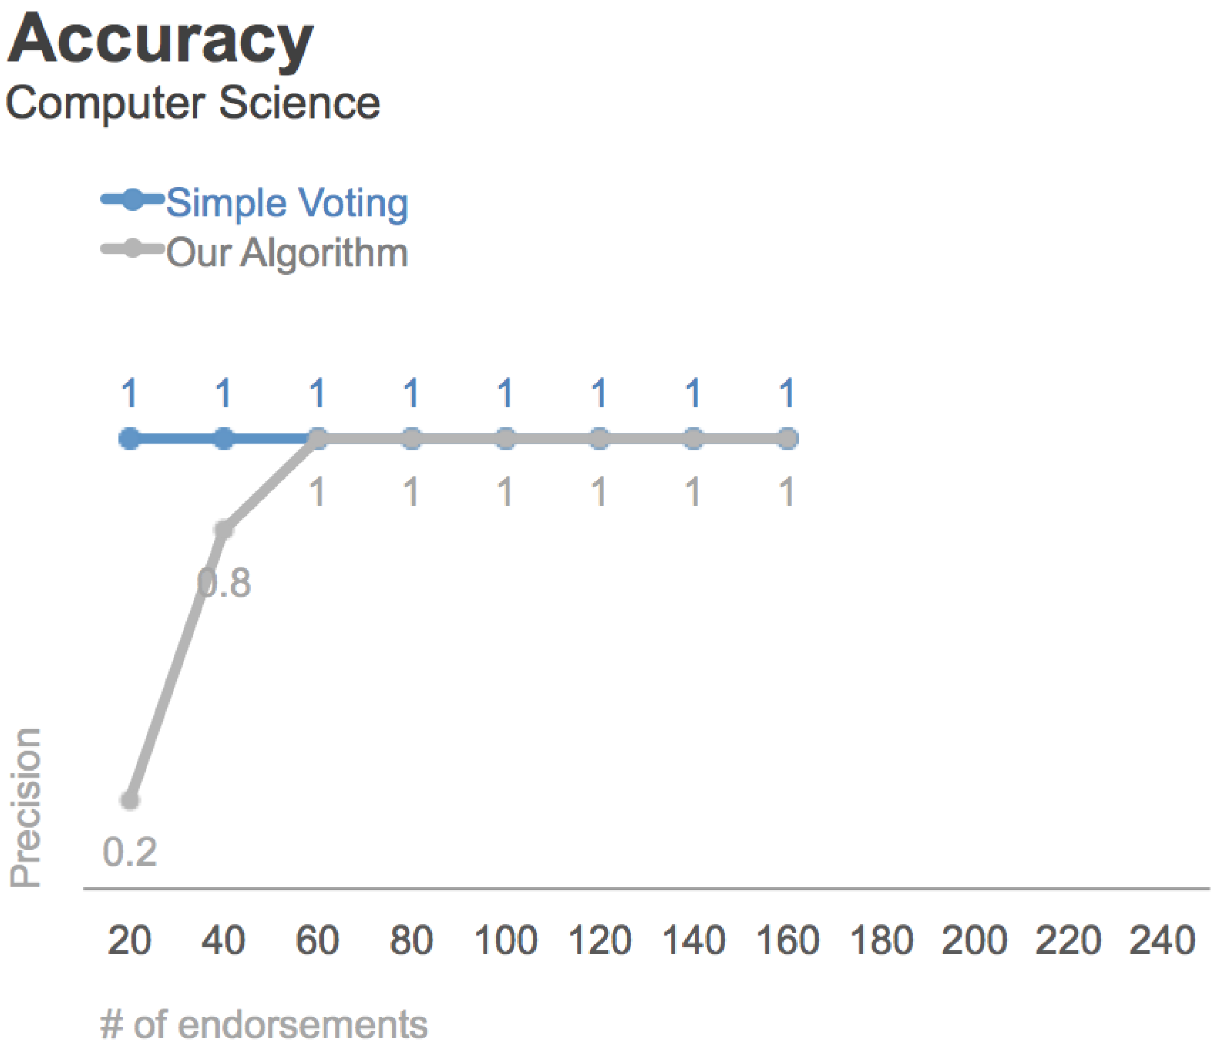

Figure 5.1: A graph comparing the simple algorithm vs our algorithm on the Computer Science use case (100% recall).

Figure 5.1 shows a comparison between the Simple Algorithm and our algorithm in the Computer Science use case (100% recall). You may notice that our algorithm takes 60

endorsements for its results to converge with Simple Algorithm's but both achieved a 100% precision.

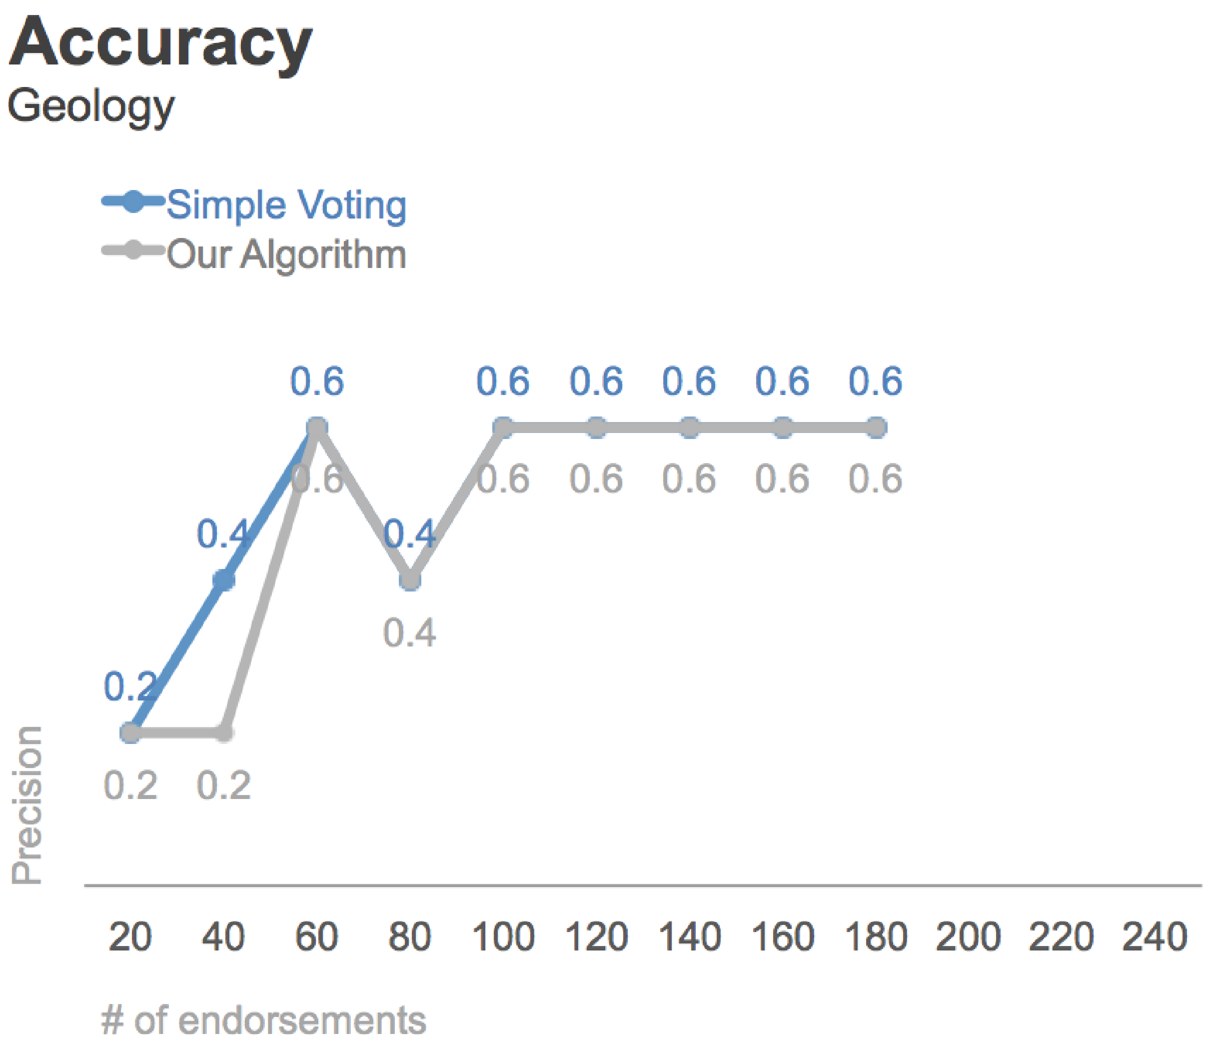

Figure 5.2: A graph comparing the simple algorithm vs our algorithm on the Geology use case.

Figure 5.2 shows a comparison between the Simple Algorithm and our algorithm in the Geology use case (100% recall). Again our algorithm requires a certain number of endorsements to match the accuracy of the Simple Algorithm. After 60 endorsements their

accuracy results are the same for both solutions maxing out at 60% precision.

Chapter 6

Related Work

Our algorithm sits at the intersection of two areas of study within the Computer Science field (below). According to [Jøsang et al., 2007] they can be described and differentiated as follows:

1. Reputation Systems - Systems that produce an entity's (public) reputation score as seen by the whole community.

2. Trust Systems - Systems that produce a score that reflects the relying party's subjective view of an entity's trustworthiness.

The differences between Reputation and Trust Systems are not always clear, in fact some approaches have elements of both; nevertheless each is given its own section to provide the reader a clearer and more straightforward way to traverse the existing literature.

6.1 Reputation Systems

The proposed algorithm for computing trust value of web-resources falls into the category of reputation systems (also known as karma systems). In reputations systems the past interactions of an agent are considered when making decisions in the present. According to [Resnick et al., 2000] reputation systems must have the following three properties:

1. Entities must be long lived, so that with every interaction there is always an expectation of future interactions.

2. Ratings about current interactions are captured and distributed.

3. Ratings about past interactions must guide decisions about current interactions.

Our system fulfills all three properties. Entities (i.e. users and Web resources) exist in the system forever after their creation or upload, respectively. Trust is propagated throughout the system upon endorsement and the trust values are meant to guide users to make decisions on Web resources.

Reputation systems can be divided into centralized and distributed [Jøsang et al., 2007] depending on how trust values are stored. Our system falls in the former category. Research in this area typically focuses on the e-commerce or professional domain while our solution focuses on an academic setting; nevertheless there are other algorithms that are similar in their general approach to the problem in both categories.

6.1.1 Centralized Reputation Systems

In this type of reputation system the performance of a given participant is obtained from other members in the community who have had direct experience with that participant. The central authority (reputation centre) collects the ratings and derives a reputation score for every participant, and makes all scores publicly available. Participants can then use these scores to decide with which participants to interact.

One of the most famous examples of reputation systems is EBay's feedback system [Resnick et al., 2006]. EBay is an electronic action house where buyers and sellers can leave feedback (similar to our own endorsement) on each other based on their transaction experience; this feedback can be positive, negative or neutral (1,-1,0). Trust in a seller is determined by the sum of positive feedback minus the sum of negative feedback. This differs from our work where we only deal with positive feedback.

The trust propagation in our system is graph-like in nature. A similar solution to our own is the work presented in [Golbeck and Hendler, 2004]. While their work also deals with trust in a graph-like manner, contrary to our solution where we aggregate values to obtain the trust of a node for the whole system, this work focuses on whether one particular node

can trust another by means of traversing intermediary nodes.

Another similar work is the one exemplified in [Guha et al., 2004]. In their solution they not only consider trust between users but also distrust, which is not considered in our approach.

A relevant related solution that combines a reputation system with provenance information encoded in PROV-O (http://www.w3.org/TR/prov-o/) is described in [Ceolin et al., 2012] which differs from our solution in their method to determine trustworthiness is by the number of times users attach the same tag to their content (videos) while our approach only requires each piece of content to be tagged once and determines trust by the endorsement process.

6.1.2 PageRank

As previously mentioned, our trust model is heavily inspired by a single algorithm PageRank [Page et al., 1999]. PageRank is Google's first and most famous algorithm. PageRank's main purpose is to measure the importance of pages in relation with the rest of the pages of the set (that set in the case of Google is the traversable World Wide Web).

PageRank's theory holds an imaginary surfer going through pages by clicking links and the PR is the probability of this user reaching a specific page. PageRank also take into consideration a damping factor which in the random surfer example would correspond to the probability that the random surfer getting bored from clicking links and would jump (by typing a new address in the URL) to a page at random.

In PageRank the assumption is that a webpage is of higher quality/has a higher ranking based on the number of inbound links coming from other pages and the quality of the websites with those links. PageRank achieves this ranking by the propagation of this "importance" value through links between pages. In this sense if page A has links to pages B and C part of the PageRank value of A will transfer to B and C, and if say page B has links to other pages, these pages will in turn receive part of B's PageRank.

Definition 11 The PageRank formula is as follows:

$$PR(A) = (1 - d) + d(PR(T1)/C(T1) + ... + PR(Tn)/C(Tn)) \tag{6.1}$$

Where PR is the PageRank of page A, d is the damping factor, and C the number of links going out of a page (count).

Similar to PageRank in our algorithm the assumption is that an entity (a user or Web resource) is more trustworthy depending on the number of endorsements it has and the quality of the same. However, our algorithm differs from PageRank in that PageRank predicted evaluation of each entity is one-dimensional (only 'importance' is measured) while our handles arbitrarily many predictions of trust depending on how many tags a resources has.

6.1.3 Distributed Reputation Systems

In this type of reputation system there is no central location that stores the reputation scores. Instead, there can be distributed stores where ratings can be submitted. Alternatively each participant keeps track of the experience with other parties, and provides this information on request to other parties. A party, who considers transacting with a given target party, must find the distributed stores, or attempt to obtain ratings from as many members as possible that have had direct experience with that target party.

A particularly famous example of Distributed Reputation Systems are Peer-to-Peer networks (P2P). In these systems tasks are divided among all participants that are both clients and server, sometimes called *servents*. Since all participants in the network have the same capabilities and responsibilities trust in the nodes is necessary for the integrity of the system.

Since our application's structure and P2P network's resembles a graph of nodes, research on P2P reputation systems are relevant to our work.

Several works have delved into this subject, including [Aberer and Despotovic, 2001], in their work they consider trust as binary, in their work the untrustworthy case is consider the

exception, therefore all nodes are considered trustworthy until they perform a transaction incorrectly. Another effect of this assumption is that only untrustworthy records are kept and distributed since only dishonest interactions are relevant.

The work of [Cornelli et al., 2002] defines the P2PRep and the work of [Damiani et al., 2002] defines XRep; both are similar approaches to reputation where a servent client queries a list of servents that can fulfill their request. Instead of the standard approach of simply selecting the servent with the best reported quality (based on connection speed) the servent send another message to the network asking its peers on opinions on the servents that offered to service their requests. This approach has several benefits such as the inherent decentralization of reputation management (no need for central authority) characteristic of all p2p solutions but since in p2p networks nodes are anonymous and require a servent to use an id and since ids are declared rather than assigned malicious parties can impersonate trustworthy clients and continue doing harm. They solved this problem by using public-key cryptography, another issue with ids is that servlets can change it at any time, therefore malicious agents with bad reputation can reset their score fairly easily hence only positive reputation is relevant in real cases.

Mixed solutions to reputation exist, that is solutions with both elements of distributed and centralized systems. One such solution is [Fahrenholtz and Lamersdorf, 2002] where they define a Reputation Management System(RMS). In their solution there is a central authority to which the clients must be registered and authenticated with but the power and knowledge still resided on the clients unlike our approach where all knowledge and power resides in the central authority.

6.2 Trust Systems

On the side of trust systems there are works like that of [Carbone et al., 2003]. In their work they define a *trust module* that handles *principals* and *trust values* and where the later can be simple values like {trusted, distrusted} or structures like tuples representing

actions (e.g. access a file) and the trust values relative to that file.

Several approaches rely on the use and exchange of credentials, these types of approaches are often referred as *trust negotiation*. The concept and the credential exchange protocol was introduced by [Yu et al., 2001].

A work highly related to ours is that of LinkedIn [Work et al., 2013]. They also use the concept of endorsement, and deal with their network and their trust in a graph-like manner. The difference is their focus, while in our approach the endorsement creates trust from user to a Web resource over a field of study, in their system the endorsement provides social proof for a user's skills.

Chapter 7

Concluding Remarks

In this thesis we propose an algorithm to determine trust in a semantically-enabled collaborative environment. We also described the implementation of this algorithm in the VLC system. Our algorithm is domain-aware through the use of tags that represent a domain (currently academic domain) as well as credentials of the users. These tags separate the trust between domains for both Web resources and users. We have tested this algorithm and obtained good results (albeit below expectations, more on this in the next subsection). Upon analysis of our test results we found that given a relatively closely related populace the algorithm provides accuracy results on par with the best of the tested algorithms. It is our new hypothesis that given a more diverse population our algorithm will outperform all other algorithms we tested in this thesis. We have achieved the objectives of this research: defining an algorithm that calculates a value of trustworthiness of Web resources and compare the proposed algorithm with simpler algorithms that do not consider user preferences or consider them in a limited fashion (more on this on the testing chapter) to evaluate its performance.

7.1 Significance of the Result

The results reveal that across the board the worst performers are the algorithms that do not consider credentials and those that assign a constant value for them, while the simple voting and the full algorithm yielded more accurate results. Looking at the data we can see that does support our hypothesis that the algorithm yields accurate results (100% accuracy in the Computer Science use case) but it does not support our hypothesis that

our algorithm would outperform the others; the simple voting algorithm performed just as well. After consideration we believe that the results are to be expected given the test case and population selection. In our institution Geology and Computer Science regularly (though not always) work closely together, especially in academic environments such as labs and research groups from where most of our research participants were taken. The intention of the algorithm is to be used in the VLC or similar systems where users of any and all fields take part and collaborate. Thus as future work we intend to test the algorithm with a use case more closely resembling this average case.

7.2 Future Work

As future work we intend to test the algorithm with a third group of participants, in addition to the Geology and Computer Science groups. This third group would have miscellaneous credentials from fields of study more removed from the Web resource's research area (For example students from psychology, art, etc.). We believe this additional group will be representative of the population we can expect from systems such as the VLC. It is our hypothesis that a more heterogeneous group, which is expected to know less of the subject at hand and is therefore more prone to errors, would lead to more accurate results for our algorithm which naturally assigns more trust to users with high level credentials in the specific area of knowledge. The algorithm in theory will decrease the trust assigned to error-prone individuals (those without credentials in the area) and increase that of knowledgeable individuals (those with credentials in the area).

Results showed that none of the tested algorithms reached 100% precision in the Geology use. Therefore, we intend to further analyze the endorsements made by participants to check if the Web resources that were not correctly classified as trustworthy were consistent across the participants and if so identify the characteristics of these resources that prevented participants to classify them correctly.

In addition we will also perform extensive artificial testing. Our new hypothesis is

that our algorithm would outperform the other variations when the system has a more error-prone population. Therefore, simulating populations of varying accuracy would allow us to increase the support of our claim. This paradigm would not require extensive and time-consuming human testing. Human testing of the required scale is unfeasible until the system is deployed in a production environment and therefore used by a large population. The testing procedure would require the implementation of several additional modules:

1. User Creation Module

2. Controller Module

3. Reporting Module

The User Creation Module will generate an arbitrarily large population, whose diversity would also be modified by parameters both in credentials (e.g. 15% Professors, 40% undergraduates,etc.) and fields of study for said credentials (e.g. 40% Computer Science, 40% Geology, 20% other). These groups will also be able to be parameterized with varying levels of accuracy, which would take the form of correct and incorrect endorsements.

The Controller Module will create a large number of populations using the user creation module with small increments in the parameters from population to population. In this fashion the controller would start by creating a population of a couple dozen users, all with low credentials and all of them from the same subject area with and toward the end would create populations numbering in the millions of users with varying credential levels in different subject areas and with different inaccuracy levels for all groups.

The Reporting Module would be an extension and modification of functionality currently existing in the system. After each population is created and the algorithm is executed for said population, the reporting algorithm would apply the same testing methodology we have used in this work; it would sort the resources in descending order and compare to the ground truth (i.e. expert list) to calculate the accuracy of each algorithm with each particular population.

Based on its design we expect our algorithm to outperform the others when the population's accuracy is higher with same-field high credential individuals (the population we expect within the VLC). The results of this artificial population will be subsequently compared to the results of the human test with the third/miscellaneous group which we expect will provide more strength to our claim. Based on the human tests we have performed so far we also expect our algorithm will at least perform as well as the other algorithms with the rest of the populations in the artificial tests.

In addition to testing, we consider some variations of the algorithm are worthy of study. One variation we consider particularly promising is the introduction of what we denote as *inferred trust*. In such variation we would consider users with high credentials in certain fields to have varying degrees of initial trust on other, related fields. For example the opinion of a professor whose area of research is Semantic Web might be useful in the area of Machine Learning, given that both areas are under the Computer Science field. This relationship is not explicit, but it can be inferred and used by creating the association between these fields and modifying the algorithm to leverage it. Creating these relationships or mining them is an area of work unto itself; an initial version could be handmade, where the relationship between fields of study (tags) are manually defined. Another option would be to crowdsource the effort by providing controls to users to define these relationships. Suggestions could be made to the users based on the history of resource access or by traversing the content of resources and looking for text resembling other tags in the system. By far the most promising solution to us would be to take advantage of preexisting knowledge and resources, such as DBpedia [Auer et al., 2007] from which we could obtain these relationships by creating an ontology (or modifying our existing one) to match the structure of these sources. In this manner we could leverage the vast knowledge these external resources have on said relationships.

Efforts that facilitate Web users to assign values of trustworthiness to Web resources will contribute to the overall utility of the world wide web, it will leverage the knowledge of the world and will formalize said knowledge so that machine agents may use to our benefit.

Efforts in this area will enhance our ability to use one of mankind's greatest inventions and eventually take us closer to the realization of its evolution, the Semantic Web. It will do so by expanding on a currently underdeveloped but paramount piece of this evolution. We are proud to have made a modest contribution to this important area and to have had the opportunity to expand the field of knowledge.

References

[Aberer and Despotovic, 2001] Aberer, K. and Despotovic, Z. (2001). Managing trust in a peer-2-peer information system. In *Proceedings of the tenth international conference on Information and knowledge management*, pages 310–317. ACM.

[Aliprand et al., 2003] Aliprand, J. et al. (2003). *The unicode standard*. Addison-Wesley.

[Artz and Gil, 2007] Artz, D. and Gil, Y. (2007). A survey of trust in computer science and the semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web*, 5(2):58–71.

[Auer et al., 2007] Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., and Ives, Z. (2007). *Dbpedia: A nucleus for a web of open data*. Springer.

[Berners-Lee et al., 2001] Berners-Lee, T., Hendler, J., Lassila, O., et al. (2001). The semantic web. *Scientific american*, 284(5):28–37.

[Berners-Lee et al., 1994] Berners-Lee, T., Masinter, L., McCahill, M., et al. (1994). Uniform resource locators (url).

[Bray et al., 2000] Bray, T., Paoli, J., Sperberg-McQueen, C. M., Maler, E., and Yergeau, F. (2000). Extensible markup language 1.0.

[Brickley and Guha, 2003] Brickley, D. and Guha, R. V. (2003). Rdf vocabulary description language 1.0: Rdf schema. w3c working draft, 23 january 2003.

[Carbone et al., 2003] Carbone, M., Nielsen, M., and Sassone, V. (2003). A formal model for trust in dynamic networks. *BRICS Report Series*, 10(4).

[Ceolin et al., 2012] Ceolin, D., Groth, P. T., Van Hage, W. R., Nottamkandath, A., and Fokkink, W. (2012). Trust evaluation through user reputation and provenance analysis. *URSW*, 900:15–26.

[Cornelli et al., 2002] Cornelli, F., Damiani, E., di Vimercati, S. D. C., Paraboschi, S., and Samarati, P. (2002). Choosing reputable servents in a p2p network. In *Proceedings of the 11th international conference on World Wide Web*, pages 376–386. ACM.

[Cyganiak et al., 2014] Cyganiak, R., Wood, D., and Lanthaler, M. (2014). Rdf 1.1 concepts and abstract syntax. *W3C Recommendation*, 25:1–8.

[Damiani et al., 2002] Damiani, E., di Vimercati, D. C., Paraboschi, S., Samarati, P., and Violante, F. (2002). A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *Proceedings of the 9th ACM conference on Computer and communications security*, pages 207–216. ACM.

[Drupal.com, 2015] Drupal.com (2015). Drupal — a cms platform for great digital experiences — drupal.com.

[Fahrenholtz and Lamersdorf, 2002] Fahrenholtz, D. and Lamersdorf, W. (2002). Transactional security for a distributed reputation management system. In *E-Commerce and Web Technologies*, pages 214–223. Springer.

[Fallside and Walmsley, 2004] Fallside, D. C. and Walmsley, P. (2004). Xml schema part 0: primer second edition. *W3C recommendation*, 16.

[Gandara, 2013] Gandara, A. (2013). *A Semantic Web-based Methodology For Describing Scientific Research Efforts*. THE UNIVERSITY OF TEXAS AT EL PASO.

[GitHub, 2015] GitHub (2015). semsol/arc2.

[Golbeck and Hendler, 2004] Golbeck, J. and Hendler, J. (2004). Accuracy of metrics for inferring trust and reputation in semantic web-based social networks. In *Engineering knowledge in the age of the semantic web*, pages 116–131. Springer.

[Grandison and Sloman, 2000] Grandison, T. and Sloman, M. (2000). A survey of trust in internet applications. *Communications Surveys & Tutorials, IEEE*, 3(4):2–16.

[Guha et al., 2004] Guha, R., Kumar, R., Raghavan, P., and Tomkins, A. (2004). Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412. ACM.

[Horridge and Patel-Schneider, 2009] Horridge, M. and Patel-Schneider, P. F. (2009). Owl 2 web ontology language manchester syntax. *W3C Working Group Note*.

[Jøsang et al., 2007] Jøsang, A., Ismail, R., and Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision support systems*, 43(2):618–644.

[Kifer, 2008] Kifer, M. (2008). Rule interchange format: The framework. In *Web reasoning and rule systems*, pages 1–11. Springer.

[Lebo et al., 2013] Lebo, T., Sahoo, S., McGuinness, D., Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., and Zhao, J. (2013). Prov-o: The prov ontology. *W3C Recommendation*, 30.

[Manola et al., 2004] Manola, F., Miller, E., McBride, B., et al. (2004). Rdf primer. *W3C recommendation*, 10(1-107):6.

[Moreau and Missier, 2013] Moreau, L. and Missier, P. (2013). Prov-dm: The prov data model.

[Mui et al., 2002] Mui, L., Mohtashemi, M., and Halberstadt, A. (2002). A computational model of trust and reputation. In *System Sciences, 2002. HICSS. Proceedings of the 35th Annual Hawaii International Conference on*, pages 2431–2439. IEEE.

[Page et al., 1999] Page, L., Brin, S., Motwani, R., and Winograd, T. (1999). The pagerank citation ranking: bringing order to the web.

[Pennington et al., 2012] Pennington, D., Gandara, A., and Gris, I. (2012). The virtual learning commons: Supporting the fuzzy front end of scientific research with emerging technologies. In *AGU Fall Meeting Abstracts*, volume 1, page 04.

[PrudHommeaux et al., 2008] PrudHommeaux, E., Seaborne, A., et al. (2008). Sparql query language for rdf. *W3C recommendation*, 15.

[Resnick et al., 2000] Resnick, P., Kuwabara, K., Zeckhauser, R., and Friedman, E. (2000). Reputation systems. *Communications of the ACM*, 43(12):45–48.

[Resnick et al., 2006] Resnick, P., Zeckhauser, R., Swanson, J., and Lockwood, K. (2006). The value of reputation on ebay: A controlled experiment. *Experimental Economics*, 9(2):79–101.

[Work et al., 2013] Work, J. D., Blue, A., and Hoffman, R. (2013). Determining reputations of users based on endorsements. US Patent App. 13/851,998.

[Yu et al., 2001] Yu, T., Winslett, M., and Seamons, K. E. (2001). Interoperable strategies in automated trust negotiation. In *Proceedings of the 8th ACM conference on Computer and Communications Security*, pages 146–155. ACM.

Appendix A

Ontology

The following is the used ontology.

```
Prefix: : <http://ontology.cybershare.utep.edu/CARP/carp-ns.owl#>
Prefix: cc: <http://web.resource.org/cc/>
Prefix: dc: <http://purl.org/dc/elements/1.1/>
Prefix: ns: <http://creativecommons.org/ns#>
Prefix: owl: <http://www.w3.org/2002/07/owl#>
Prefix: rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
Prefix: xml: <http://www.w3.org/XML/1998/namespace>
Prefix: xsd: <http://www.w3.org/2001/XMLSchema#>
Prefix: carp: <http://ontology.cybershare.utep.edu/CARP/carp.owl#>
Prefix: dcam: <http://purl.org/dc/dcam/>
Prefix: foaf: <http://xmlns.com/foaf/0.1/>
Prefix: prov: <http://www.w3.org/ns/prov#>
Prefix: rdfs: <http://www.w3.org/2000/01/rdf-schema#>
Prefix: sioc: <http://rdfs.org/sioc/ns#>
Prefix: skos: <http://www.w3.org/2004/02/skos/core#>
Prefix: terms: <http://purl.org/dc/terms/>
Prefix: vocab: <http://www.w3.org/1999/xhtml/vocab#>
Prefix: aboutdcmi: <http://purl.org/dc/aboutdcmi#>
Prefix: wgs84_pos: <http://www.w3.org/2003/01/geo/wgs84_pos#>
Prefix: XMLSchema-datatypes: <http://www.w3.org/2001/XMLSchema-datatypes#>
```

```
Ontology: <http://ontology.cybershare.utep.edu/CARP/carp.owl>


Annotations:
    owl:versionInfo "Recommendation version 2013-04-30"@en,
   rdfs:comment "The namespace that supports the Collect Annotate Refine Publish methodology
    rdfs:label "CARP Namespace",
   rdfs:comment "This document is published by the Provenance Working Group (http://www.w3.c

If you wish to make comments regarding this document, please send them to public-prov-comment
    rdfs:seeAlso <http://www.w3.org/TR/prov-o/>,
    rdfs:seeAlso <http://www.w3.org/ns/prov>,
    rdfs:label "W3C PROVenance Interchange Ontology (PROV-O)"@en

AnnotationProperty: prov:aq

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubPropertyOf:
        rdfs:seeAlso
```

```
AnnotationProperty: prov:wasRevisionOf

    Annotations:
        prov:category "expanded",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "A revision is a derivation that revises an entity into a revised version
        prov:qualifiedForm <http://www.w3.org/ns/prov#Revision>,
        prov:component "derivations",
        prov:inverse "hadRevision",
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedRevision>,
        rdfs:label "wasRevisionOf"

    Annotations:
      rdfs:comment "A revision is a derivation that revises an entity into a revised version
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "wasRevisionOf",
        prov:category "expanded",
        prov:component "derivations",
        prov:inverse "hadRevision",
        prov:qualifiedForm <http://www.w3.org/ns/prov#Revision>,
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedRevision>


AnnotationProperty: rdfs:isDefinedBy


AnnotationProperty: prov:editorialNote
```

```
    Annotations:

        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,

      rdfs:comment "A note by the OWL development team about how this term expresses the PRO


    Annotations:

      rdfs:comment "A note by the OWL development team about how this term expresses the PRO

        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>



AnnotationProperty: prov:sharesDefinitionWith


    Annotations:

        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>


    Annotations:

        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>


    SubPropertyOf:

        rdfs:seeAlso



AnnotationProperty: rdfs:seeAlso


    Annotations:

        rdfs:comment ""@en


    Annotations:

        rdfs:comment ""@en
```

```
AnnotationProperty: owl:versionInfo


AnnotationProperty: prov:category

    Annotations:
      rdfs:comment "Classify prov-o terms into three categories, including 'starting-point'
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    Annotations:
      rdfs:comment "Classify prov-o terms into three categories, including 'starting-point'
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>


AnnotationProperty: prov:editorsDefinition

    Annotations:
      rdfs:comment "When the prov-o term does not have a definition drawn from prov-dm, and
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    Annotations:
      rdfs:comment "When the prov-o term does not have a definition drawn from prov-dm, and
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubPropertyOf:
        prov:definition
```

```
AnnotationProperty: prov:unqualifiedForm

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "Classes and properties used to qualify relationships are annotated with

    Annotations:
      rdfs:comment "Classes and properties used to qualify relationships are annotated with
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubPropertyOf:
        rdfs:seeAlso


AnnotationProperty: prov:order

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "The position that this OWL term should be listed within documentation. T

    Annotations:
      rdfs:comment "The position that this OWL term should be listed within documentation. T
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>


AnnotationProperty: rdfs:comment
```

```
    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:comment ""@en

    Annotations:
        rdfs:comment ""@en,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>


AnnotationProperty: prov:specializationOf

    Annotations:
      prov:definition "An entity that is a specialization of another shares all aspects of t
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
        prov:category "expanded",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-specialization"^^xsd:a
        prov:component "alternate",
        rdfs:label "specializationOf",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:inverse "generalizationOf",
        rdfs:seeAlso <http://www.w3.org/ns/prov#alternateOf>,
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-specialization"^^x

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "specializationOf",
        rdfs:seeAlso <http://www.w3.org/ns/prov#alternateOf>,
```

```
        prov:category "expanded",
        prov:component "alternate",
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
      prov:definition "An entity that is a specialization of another shares all aspects of t
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-specialization"^^xsd:a
        prov:inverse "generalizationOf",
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-specialization"^^x


AnnotationProperty: prov:definition


    Annotations:
      rdfs:comment "A definition quoted from PROV-DM or PROV-CONSTRAINTS that describes the
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>


    Annotations:
      rdfs:comment "A definition quoted from PROV-DM or PROV-CONSTRAINTS that describes the
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>



AnnotationProperty: prov:qualifiedForm


    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "This annotation property links a subproperty of prov:wasInfluencedBy wi


Example annotation:
```

```
prov:wasGeneratedBy prov:qualifiedForm prov:qualifiedGeneration, prov:Generation .
```

Then this unqualified assertion:

```
:entity1 prov:wasGeneratedBy :activity1 .
```

can be qualified by adding:

```
:entity1 prov:qualifiedGeneration :entity1Gen .
:entity1Gen
    a prov:Generation, prov:Influence;
    prov:activity :activity1;
    :customValue 1337 .
```

Note how the value of the unqualified influence (prov:wasGeneratedBy :activity1) is mirrored

Annotations:
rdfs:comment "This annotation property links a subproperty of prov:wasInfluencedBy wi

Example annotation:

```
prov:wasGeneratedBy prov:qualifiedForm prov:qualifiedGeneration, prov:Generation .
```

Then this unqualified assertion:

```
:entity1 prov:wasGeneratedBy :activity1 .
```

can be qualified by adding:

```
   :entity1 prov:qualifiedGeneration :entity1Gen .
   :entity1Gen
       a prov:Generation, prov:Influence;
       prov:activity :activity1;
       :customValue 1337 .
```

Note how the value of the unqualified influence (prov:wasGeneratedBy :activity1) is mirrored

```
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubPropertyOf:
        rdfs:seeAlso



AnnotationProperty: prov:todo



AnnotationProperty: prov:inverse

    Annotations:
       rdfs:seeAlso <http://www.w3.org/TR/prov-o/#names-of-inverse-properties>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
       rdfs:comment "PROV-O does not define all property inverses. The directionalities defi

    Annotations:
       rdfs:comment "PROV-O does not define all property inverses. The directionalities defi
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
       rdfs:seeAlso <http://www.w3.org/TR/prov-o/#names-of-inverse-properties>
```

```
AnnotationProperty: rdfs:label

    Annotations:
        rdfs:comment ""@en,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    Annotations:
        rdfs:comment ""@en,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>


AnnotationProperty: prov:constraints

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "A reference to the principal section of the PROV-CONSTRAINTS document th

    Annotations:
      rdfs:comment "A reference to the principal section of the PROV-CONSTRAINTS document th
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubPropertyOf:
        rdfs:seeAlso


AnnotationProperty: prov:dm
```

```
    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "A reference to the principal section of the PROV-DM document that descri

    Annotations:
      rdfs:comment "A reference to the principal section of the PROV-DM document that descri
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubPropertyOf:
        rdfs:seeAlso


AnnotationProperty: prov:component

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "Classify prov-o terms into six components according to prov-dm, includir

    Annotations:
      rdfs:comment "Classify prov-o terms into six components according to prov-dm, includir
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>


AnnotationProperty: prov:n

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
```

```
    rdfs:comment "A reference to the principal section of the PROV-DM document that descri

  Annotations:
    rdfs:comment "A reference to the principal section of the PROV-DM document that descri
      rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

  SubPropertyOf:
      rdfs:seeAlso


Datatype: rdfs:Literal


Datatype: xsd:dateTime


Datatype: xsd:string


Datatype: rdf:PlainLiteral


Datatype: xsd:double


Datatype: xsd:anyURI
```

```
ObjectProperty: carp:hasEndorsee

    Domain:
        carp:Endorsement

    Range:
        carp:Resource


ObjectProperty: prov:qualifiedUsage

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "qualifiedUsage",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#used>,
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Usage>,
        prov:category "qualified",
      rdfs:comment "If this Activity prov:used Entity :e, then it can qualify how it used it
        prov:inverse "qualifiedUsingActivity",
        prov:component "entities-activities"

    Domain:
        prov:Activity

    Range:
        prov:Usage

    SubPropertyOf:
```

```
        prov:qualifiedInfluence


ObjectProperty: carp:hasFieldSite

    Domain:
        carp:Project

    Range:
        carp:FieldSite

    InverseOf:
        carp:isFieldSiteOf


ObjectProperty: prov:qualifiedInvalidation

    Annotations:
        prov:inverse "qualifiedInvalidationOf",
      rdfs:comment "If this Entity prov:wasInvalidatedBy Activity :a, then it can qualify ho
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasInvalidatedBy>,
        prov:category "qualified",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Invalidation>,
        rdfs:label "qualifiedInvalidation",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "entities-activities"

    Domain:
```

```
        prov:Entity

    Range:
        prov:Invalidation

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: prov:qualifiedPrimarySource

    Annotations:
      rdfs:comment "If this Entity prov:hadPrimarySource Entity :e, then it can qualify how
        prov:category "qualified",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#hadPrimarySource>,
        rdfs:label "qualifiedPrimarySource",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#PrimarySource>,
        prov:component "derivations",
        prov:inverse "qualifiedSourceOf"

    Domain:
        prov:Entity

    Range:
        prov:PrimarySource

    SubPropertyOf:
```

```
        prov:qualifiedInfluence


ObjectProperty: carp:isEndorserOf

    Domain:
        carp:Person

    Range:
        carp:Endorsement


ObjectProperty: carp:isFundingAgencyOf

    InverseOf:
        carp:hasFundingAgency


ObjectProperty: carp:hasOwner

    Annotations:
        rdfs:label "hasOwner"@en

    Domain:
        carp:Collection

    Range:
        carp:Person
```

```
ObjectProperty: carp:hasFundingAgency

    InverseOf:
        carp:isFundingAgencyOf


ObjectProperty: carp:hasProjectInvestigator

    Annotations:
        rdfs:label "has project investigator"

    Domain:
        carp:Project

    Range:
        carp:ProjectInvestigator

    InverseOf:
        carp:isProjectInvestigatorOf


ObjectProperty: prov:alternateOf

    Annotations:
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-alternate"^^xsd:ar
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
```

```
        prov:component "alternate",
      prov:definition "Two alternate entities present aspects of the same thing. These aspec
        prov:inverse "alternateOf",
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
        prov:category "expanded",
        rdfs:label "alternateOf",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-alternate"^^xsd:anyURI
        rdfs:seeAlso <http://www.w3.org/ns/prov#specializationOf>

    Domain:
        prov:Entity

    Range:
        prov:Entity


ObjectProperty: prov:wasGeneratedBy

    Annotations:
        prov:component "entities-activities",
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedGeneration>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "wasGeneratedBy",
        prov:inverse "generated",
        prov:qualifiedForm <http://www.w3.org/ns/prov#Generation>,
        prov:category "starting-point"

    Domain:
```

```
        prov:Entity

    Range:
        prov:Activity

    SubPropertyOf:
        prov:wasInfluencedBy

    InverseOf:
        prov:generated

    SubPropertyChain:
        prov:qualifiedGeneration o prov:activity


ObjectProperty: prov:hadGeneration

    Annotations:
        prov:category "qualified",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Generation>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "hadGeneration",
        prov:inverse "generatedAsDerivation",
        prov:component "derivations",
       rdfs:comment "The _optional_ Generation involved in an Entity's Derivation."@en

    Domain:
        prov:Derivation
```

```
    Range:
        prov:Generation


ObjectProperty: carp:hasAreaOfKnowledge

    EquivalentTo:
        carp:hasKeyword,
        carp:hasTag


ObjectProperty: prov:influencer

    Annotations:
        prov:category "qualified",
        prov:inverse "hadInfluence",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-influence"^^xsd:anyURI
      prov:editorsDefinition "This property is used as part of the qualified influence patte
      prov:editorialNote "This property and its subproperties are used in the same way as th
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "influencer",
      rdfs:comment "Subproperties of prov:influencer are used to cite the object of an unqua

    Domain:
        prov:Influence

    Range:
```

```
        owl:Thing


ObjectProperty: prov:invalidated

    Annotations:
        rdfs:label "invalidated",
        prov:category "expanded",
      prov:editorialNote "prov:invalidated is one of few inverse property defined, to allow
        prov:component "entities-activities",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Invalidation>,
        prov:inverse "wasInvalidatedBy"

    Domain:
        prov:Activity

    Range:
        prov:Entity

    SubPropertyOf:
        prov:influenced

    InverseOf:
        prov:wasInvalidatedBy


ObjectProperty: prov:wasQuotedFrom
```

```
    Annotations:
        prov:category "expanded",
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedQuotation>,
        prov:qualifiedForm <http://www.w3.org/ns/prov#Quotation>,
      rdfs:comment "An entity is derived from an original entity by copying, or 'quoting', s
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "derivations",
        prov:inverse "quotedAs",
        rdfs:label "wasQuotedFrom"

    Domain:
        prov:Entity

    Range:
        prov:Entity

    SubPropertyOf:

          Annotations: rdfs:comment "Quotation is a particular case of derivation (see http:
        prov:wasDerivedFrom

    SubPropertyChain:
        prov:qualifiedQuotation o prov:entity


ObjectProperty: prov:hadMember
```

```
    Annotations:
        rdfs:label "hadMember",
        prov:inverse "wasMemberOf",
        prov:category "expanded",
        prov:component "expanded",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Collection>

    Domain:
        prov:Collection

    Range:

          Annotations: rdfs:comment "A collection is an entity that provides a structure to s
                     prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-collecti
        prov:Entity

    SubPropertyOf:
        prov:wasInfluencedBy


ObjectProperty: prov:qualifiedStart

    Annotations:
      rdfs:comment "If this Activity prov:wasStartedBy Entity :e1, then it can qualify how i
        prov:category "qualified",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasStartedBy>,
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Start>,
```

```
        prov:inverse "qualifiedStartOf",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "entities-activities",
        rdfs:label "qualifiedStart"

    Domain:
        prov:Activity

    Range:
        prov:Start

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: carp:hasPart

    Annotations:
        rdfs:label "hasPart"@en

    InverseOf:
        carp:isPartOf


ObjectProperty: carp:hasUploader

    Domain:
        carp:Resource
```

```
    Range:
        carp:Person


ObjectProperty: owl:topObjectProperty


ObjectProperty: carp:hasInstitution

    InverseOf:
        carp:isInstitutionOf


ObjectProperty: prov:qualifiedAttribution

    Annotations:
        rdfs:label "qualifiedAttribution",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasAttributedTo>,
        prov:component "agents-responsibility",
        prov:inverse "qualifiedAttributionOf",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:category "qualified",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Attribution>,
       rdfs:comment "If this Entity prov:wasAttributedTo Agent :ag, then it can qualify how i

    Domain:
        prov:Entity
```

```
    Range:
        prov:Attribution

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: prov:qualifiedDerivation

    Annotations:
        prov:component ”derivations”,
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Derivation>,
      rdfs:comment ”If this Entity prov:wasDerivedFrom Entity :e, then it can qualify how it
        prov:category ”qualified”,
        prov:inverse ”qualifiedDerivationOf”,
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasDerivedFrom>,
        rdfs:label ”qualifiedDerivation”,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    Domain:
        prov:Entity

    Range:
        prov:Derivation

    SubPropertyOf:
        prov:qualifiedInfluence
```

```
ObjectProperty: prov:activity

    Annotations:
        prov:category "qualified",
        rdfs:label "activity",
      prov:editorialNote "This property behaves in spirit like rdf:object; it references the
        prov:inverse "activityOfInfluence",
      prov:editorsDefinition "The prov:activity property references an prov:Activity which
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    Domain:
        prov:ActivityInfluence

    Range:
        prov:Activity

    SubPropertyOf:
        prov:influencer


ObjectProperty: prov:wasAssociatedWith

    Annotations:
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedAssociation>,
        prov:inverse "wasAssociateFor",
        prov:category "starting-point",
```

```
        prov:qualifiedForm <http://www.w3.org/ns/prov#Association>,
        rdfs:label "wasAssociatedWith",
      rdfs:comment "An prov:Agent that had some (unspecified) responsibility for the occurre
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "agents-responsibility"

    Domain:
        prov:Activity

    Range:
        prov:Agent

    SubPropertyOf:
        prov:wasInfluencedBy

    SubPropertyChain:
        prov:qualifiedAssociation o prov:agent


ObjectProperty: prov:qualifiedRevision

    Annotations:
        prov:component "derivations",
        rdfs:label "qualifiedRevision",
        prov:inverse "revisedEntity",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasRevisionOf>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "If this Entity prov:wasRevisionOf Entity :e, then it can qualify how it
```

```
        prov:category "qualified",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Revision>

    Domain:
        prov:Entity

    Range:
        prov:Revision

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: prov:wasInfluencedBy

    Annotations:
        rdfs:label "wasInfluencedBy",
       prov:editorialNote "The sub-properties of prov:wasInfluencedBy can be elaborated in mo

Subproperties of prov:wasInfluencedBy may also be asserted directly without being qualified.

prov:wasInfluencedBy should not be used without also using one of its subproperties.
"@en,
        prov:category "qualified",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Influence>,
        prov:component "agents-responsibility",
       rdfs:comment "This property has multiple RDFS domains to suit multiple OWL Profiles. S
        prov:qualifiedForm <http://www.w3.org/ns/prov#Influence>,
```

```
        prov:inverse "influenced",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
       rdfs:comment "Because prov:wasInfluencedBy is a broad relation, its more specific subp
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedInfluence>

    Domain:

          Annotations: prov:definition "influencee: an identifier (o2) for an entity, activi
                       prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-influenc
        prov:Activity or prov:Agent or prov:Entity

    Range:

          Annotations: prov:definition "influencer: an identifier (o1) for an ancestor entit
                       prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-influenc
        prov:Activity or prov:Agent or prov:Entity

    InverseOf:
        prov:influenced


ObjectProperty: prov:influenced

    Annotations:
        prov:inverse "wasInfluencedBy",
        prov:category "expanded",
        prov:component "agents-responsibility",
        rdfs:label "influenced",
```

```
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Influence>

    InverseOf:
        prov:wasInfluencedBy


ObjectProperty: carp:hasResource

    Range:
        carp:Resource


ObjectProperty: prov:hadActivity

    Annotations:
        prov:component "derivations",
        rdfs:label "hadActivity",
        prov:inverse "wasActivityOfInfluence",
      rdfs:comment "The _optional_ Activity of an Influence, which used, generated, invalida
      prov:editorialNote "The multiple rdfs:domain assertions are intended. One is simpler a
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "This property has multiple RDFS domains to suit multiple OWL Profiles. S
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Activity>,
        prov:category "qualified"

    Domain:
        prov:Influence,
```

```
        prov:Delegation or prov:Derivation or prov:End or prov:Start

    Range:
        prov:Activity


ObjectProperty: prov:hadRole

    Annotations:
        prov:inverse "wasRoleIn",
        prov:category "qualified",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "This property has multiple RDFS domains to suit multiple OWL Profiles. S
      rdfs:comment "The _optional_ Role that an Entity assumed in the context of an Activity
        prov:component "agents-responsibility",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Role>,
        rdfs:label "hadRole",
      prov:editorsDefinition "prov:hadRole references the Role (i.e. the function of an enti

    Domain:
        prov:Influence,
        prov:Association or prov:InstantaneousEvent

    Range:
        prov:Role


ObjectProperty: carp:isPartOf
```

```
    Annotations:
        rdfs:label "isPartOf"@en

    Domain:
        carp:Resource

    Range:
        carp:Collection

    Characteristics:
        Transitive

    InverseOf:
        carp:hasPart


ObjectProperty: carp:isFieldSiteOf

    InverseOf:
        carp:hasFieldSite



ObjectProperty: prov:wasInformedBy

    Annotations:
        prov:category "starting-point",
        prov:inverse "informed",
```

```
      rdfs:comment "An activity a2 is dependent on or informed by another activity a1, by wa
       prov:component "entities-activities",
      prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedCommunication>,
       rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
       rdfs:label "wasInformedBy",
       prov:qualifiedForm <http://www.w3.org/ns/prov#Communication>

    Domain:
       prov:Activity

    Range:
       prov:Activity

    SubPropertyOf:
       prov:wasInfluencedBy

    SubPropertyChain:
       prov:qualifiedCommunication o prov:activity


ObjectProperty: prov:wasRevisionOf

    Annotations:
       prov:category "expanded",
       rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "A revision is a derivation that revises an entity into a revised version
       prov:qualifiedForm <http://www.w3.org/ns/prov#Revision>,
       prov:component "derivations",
```

```
        prov:inverse "hadRevision",
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedRevision>,
        rdfs:label "wasRevisionOf"

    Domain:
        prov:Entity

    Range:
        prov:Entity

    SubPropertyOf:

          Annotations: rdfs:comment "Revision is a derivation (see http://www.w3.org/TR/prov
http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#term-Revision 23 April 2012 'wasRe
        prov:wasDerivedFrom

    SubPropertyChain:
        prov:qualifiedRevision o prov:entity


ObjectProperty: prov:generated

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "entities-activities",
      prov:editorialNote "prov:generated is one of few inverse property defined, to allow Ac
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Generation>,
        rdfs:label "generated",
```

```
        prov:inverse "wasGeneratedBy",
        prov:category "expanded"

    Domain:
        prov:Activity

    Range:
        prov:Entity

    SubPropertyOf:
        prov:influenced

    InverseOf:
        prov:wasGeneratedBy


ObjectProperty: prov:qualifiedDelegation

    Annotations:
       rdfs:comment "If this Agent prov:actedOnBehalfOf Agent :ag, then it can qualify how wi
        prov:inverse "qualifiedDelegationOf",
        prov:category "qualified",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Delegation>,
        prov:unqualifiedForm <http://www.w3.org/ns/prov#actedOnBehalfOf>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "agents-responsibility",
        rdfs:label "qualifiedDelegation"
```

```
    Domain:
        prov:Agent

    Range:
        prov:Delegation

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: prov:qualifiedEnd

    Annotations:
      rdfs:comment "If this Activity prov:wasEndedBy Entity :e1, then it can qualify how it
        rdfs:label "qualifiedEnd",
        prov:category "qualified",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasEndedBy>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "entities-activities",
        prov:inverse "qualifiedEndOf",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#End>

    Domain:
        prov:Activity

    Range:
        prov:End
```

```
    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: carp:isInstitutionOf

    SubPropertyOf:
        owl:topObjectProperty

    InverseOf:
        carp:hasInstitution


ObjectProperty: prov:specializationOf

    Annotations:
      prov:definition "An entity that is a specialization of another shares all aspects of t
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
        prov:category "expanded",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-specialization"^^xsd:a
        prov:component "alternate",
        rdfs:label "specializationOf",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:inverse "generalizationOf",
        rdfs:seeAlso <http://www.w3.org/ns/prov#alternateOf>,
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-specialization"^^x

    Domain:
```

```
        prov:Entity

    Range:
        prov:Entity

    SubPropertyOf:
        prov:alternateOf


ObjectProperty: prov:used

    Annotations:
        prov:qualifiedForm <http://www.w3.org/ns/prov#Usage>,
        prov:category "starting-point",
        prov:inverse "wasUsedBy",
      rdfs:comment "A prov:Entity that was used by this prov:Activity. For example, :baking
        prov:component "entities-activities",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "used",
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedUsage>

    Domain:
        prov:Activity

    Range:
        prov:Entity

    SubPropertyOf:
```

```
        prov:wasInfluencedBy

    SubPropertyChain:
        prov:qualifiedUsage o prov:entity


ObjectProperty: prov:qualifiedInfluence

    Annotations:
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Influence>,
        rdfs:label "qualifiedInfluence",
        prov:category "qualified",
      rdfs:comment "Because prov:qualifiedInfluence is a broad relation, the more specific
        prov:component "derivations",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:inverse "qualifiedInfluenceOf",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasInfluencedBy>

    Domain:
        prov:Activity or prov:Agent or prov:Entity

    Range:
        prov:Influence


ObjectProperty: prov:wasEndedBy

    Annotations:
```

```
        prov:component "entities-activities",
        prov:qualifiedForm <http://www.w3.org/ns/prov#End>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:inverse "ended",
        rdfs:label "wasEndedBy",
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedEnd>,
       rdfs:comment "End is when an activity is deemed to have ended. An end may refer to an e
        prov:category "expanded"

    Domain:
        prov:Activity

    Range:
        prov:Entity

    SubPropertyOf:
        prov:wasInfluencedBy

    SubPropertyChain:
        prov:qualifiedEnd o prov:entity


ObjectProperty: prov:qualifiedGeneration

    Annotations:
        prov:component "entities-activities",
        prov:inverse "qualifiedGenerationOf",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasGeneratedBy>,
```

```
        prov:category "qualified",
       rdfs:comment "If this Activity prov:generated Entity :e, then it can qualify how it pe
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Generation>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "qualifiedGeneration"

    Domain:
        prov:Entity

    Range:
        prov:Generation

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: prov:qualifiedCommunication

    Annotations:
        prov:qualifiedForm <http://www.w3.org/ns/prov#Communication>,
        rdfs:label "qualifiedCommunication",
        prov:category "qualified",
        prov:inverse "qualifiedCommunicationOf",
        prov:component "entities-activities",
       rdfs:comment "If this Activity prov:wasInformedBy Activity :a, then it can qualify how
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Communication>
```

```
    Domain:
        prov:Activity

    Range:
        prov:Communication

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: prov:wasDerivedFrom

    Annotations:
        prov:qualifiedForm <http://www.w3.org/ns/prov#Derivation>,
      prov:definition "A derivation is a transformation of an entity into another, an update
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:inverse "hadDerivation",
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedDerivation>,
        rdfs:label "wasDerivedFrom",
      rdfs:comment "The more specific subproperties of prov:wasDerivedFrom (i.e., prov:wasQ
        prov:component "derivations",
        prov:category "starting-point"

    Domain:
        prov:Entity

    Range:
        prov:Entity
```

```
    SubPropertyOf:

          Annotations: rdfs:comment "Derivation is a particular case of trace (see http://www
        prov:wasInfluencedBy

    SubPropertyChain:
        prov:qualifiedDerivation o prov:entity


ObjectProperty: prov:actedOnBehalfOf

    Annotations:
        prov:qualifiedForm <http://www.w3.org/ns/prov#Delegation>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "actedOnBehalfOf",
        prov:component "agents-responsibility",
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedDelegation>,
        prov:inverse "hadDelegate",
        prov:category "starting-point",
      rdfs:comment "An object property to express the accountability of an agent towards anc

    Domain:
        prov:Agent

    Range:
        prov:Agent
```

```
    SubPropertyOf:
        prov:wasInfluencedBy

    SubPropertyChain:
        prov:qualifiedDelegation o prov:agent


ObjectProperty: prov:wasAttributedTo

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:category "starting-point",
        prov:qualifiedForm <http://www.w3.org/ns/prov#Attribution>,
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedAttribution>,
      prov:definition "Attribution is the ascribing of an entity to an agent."@en,
        rdfs:label "wasAttributedTo",
        prov:inverse "contributed",
      rdfs:comment "Attribution is the ascribing of an entity to an agent."@en,
        prov:component "agents-responsibility"

    Domain:
        prov:Entity

    Range:
        prov:Agent

    SubPropertyOf:
```

```
      Annotations: rdfs:comment "Attribution is a particular case of trace (see http://w
                 prov:definition "IF wasAttributedTo(e2,ag1,aAttr) holds, THEN wasInflue
        prov:wasInfluencedBy

    SubPropertyChain:
        prov:qualifiedAttribution o prov:agent


ObjectProperty: prov:wasInvalidatedBy

    Annotations:
        prov:qualifiedForm <http://www.w3.org/ns/prov#Invalidation>,
        prov:component "entities-activities",
        prov:category "expanded",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedInvalidation>,
        prov:inverse "invalidated",
        rdfs:label "wasInvalidatedBy"

    Domain:
        prov:Entity

    Range:
        prov:Activity

    SubPropertyOf:
        prov:wasInfluencedBy
```

```
    InverseOf:
        prov:invalidated

    SubPropertyChain:
        prov:qualifiedInvalidation o prov:activity


ObjectProperty: prov:hadUsage

    Annotations:
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Usage>,
        rdfs:label "hadUsage",
        prov:inverse "wasUsedInDerivation",
        prov:category "qualified",
      rdfs:comment "The _optional_ Usage involved in an Entity's Derivation."@en,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "derivations"

    Domain:
        prov:Derivation

    Range:
        prov:Usage


ObjectProperty: carp:hasTag

    EquivalentTo:
```

```
        carp:hasAreaOfKnowledge


ObjectProperty: prov:qualifiedQuotation

    Annotations:
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasQuotedFrom>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:inverse "qualifiedQuotationOf",
      rdfs:comment "If this Entity prov:wasQuotedFrom Entity :e, then it can qualify how usi
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Quotation>,
        prov:category "qualified",
        rdfs:label "qualifiedQuotation",
        prov:component "derivations"

    Domain:
        prov:Entity

    Range:
        prov:Quotation

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: carp:isProjectInvestigatorOf

    Characteristics:
```

```
        Transitive

    InverseOf:
        carp:hasProjectInvestigator


ObjectProperty: prov:entity

    Annotations:
        prov:inverse "entityOfInfluence",
      prov:editorsDefinition "The prov:entity property references an prov:Entity which infl
        rdfs:label "entity",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:category "qualified",
      prov:editorialNote "This property behaves in spirit like rdf:object; it references the

    Domain:
        prov:EntityInfluence

    Range:
        prov:Entity

    SubPropertyOf:
        prov:influencer


ObjectProperty: prov:qualifiedAssociation
```

```
    Annotations:
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasAssociatedWith>,
        prov:category "qualified",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "agents-responsibility",
        rdfs:label "qualifiedAssociation",
        prov:inverse "qualifiedAssociationOf",
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Association>,
       rdfs:comment "If this Activity prov:wasAssociatedWith Agent :ag, then it can qualify t

    Domain:
        prov:Activity

    Range:
        prov:Association

    SubPropertyOf:
        prov:qualifiedInfluence


ObjectProperty: prov:atLocation

    Annotations:
       rdfs:comment "This property has multiple RDFS domains to suit multiple OWL Profiles. S
        prov:category "expanded",
       prov:editorialNote "This property is not functional because the many values could be a
        prov:inverse "locationOf",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
```

```
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Location>,
        rdfs:comment "The Location of any resource."@en,
      prov:editorialNote "The naming of prov:atLocation parallels prov:atTime, and is not na
        rdfs:label "atLocation"

    Domain:
        prov:Activity or prov:Agent or prov:Entity or prov:InstantaneousEvent

    Range:
        prov:Location


ObjectProperty: prov:wasStartedBy

    Annotations:
        prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedStart>,
      rdfs:comment "Start is when an activity is deemed to have started. A start may refer to
        rdfs:label "wasStartedBy",
        prov:inverse "started",
        prov:category "expanded",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "entities-activities",
        prov:qualifiedForm <http://www.w3.org/ns/prov#Start>

    Domain:
        prov:Activity

    Range:
```

```
        prov:Entity

    SubPropertyOf:
        prov:wasInfluencedBy

    SubPropertyChain:
        prov:qualifiedStart o prov:entity


ObjectProperty: prov:hadPlan

    Annotations:
        prov:category "qualified",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "agents-responsibility",
        rdfs:label "hadPlan",
        prov:inverse "wasPlanOf",
      rdfs:comment "The _optional_ Plan adopted by an Agent in Association with some Activit
        prov:sharesDefinitionWith <http://www.w3.org/ns/prov#Plan>

    Domain:
        prov:Association

    Range:
        prov:Plan


ObjectProperty: carp:hasKeyword
```

```
    Annotations:
        rdfs:label "hasKeyword"@en

    EquivalentTo:
        carp:hasAreaOfKnowledge


ObjectProperty: prov:agent

    Annotations:
        prov:category "qualified",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      prov:editorsDefinition "The prov:agent property references an prov:Agent which influe
      prov:editorialNote "This property behaves in spirit like rdf:object; it references the
        rdfs:label "agent",
        prov:inverse "agentOfInfluence"

    Domain:
        prov:AgentInfluence

    Range:
        prov:Agent

    SubPropertyOf:
        prov:influencer
```

```
ObjectProperty: prov:hadPrimarySource

    Annotations:
       prov:qualifiedForm <http://www.w3.org/ns/prov#qualifiedPrimarySource>,
        rdfs:label "hadPrimarySource",
        prov:category "expanded",
        prov:component "derivations",
        prov:qualifiedForm <http://www.w3.org/ns/prov#PrimarySource>,
        prov:inverse "wasPrimarySourceOf",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    Domain:
        prov:Entity

    Range:
        prov:Entity

    SubPropertyOf:

          Annotations: rdfs:comment "hadPrimarySource property is a particular case of wasDe
        prov:wasDerivedFrom

    SubPropertyChain:
        prov:qualifiedPrimarySource o prov:entity


DataProperty: prov:value
```

```
  Annotations:
    prov:definition "Provides a value that is a direct representation of an entity."@en,
      prov:category "expanded",
    prov:editorialNote "The editor's definition comes from http://www.w3.org/TR/rdf-prime
    prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-attribute-value"^^xsd:
      prov:component "entities-activities",
      rdfs:label "value",
      rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
    prov:editorialNote "This property serves the same purpose as rdf:value, but has been r

  Domain:
      prov:Entity


DataProperty: carp:hasQualityRating

  Annotations:
    rdfs:comment "This may be the trust rating that Joaquin is working on. If not, provide
      rdfs:label "has quality rating"@en

  Range:
      xsd:double


DataProperty: carp:hasKeyword

  Annotations:
      rdfs:label "hasKeyword"@en
```

```
    Range:
        xsd:string

    SubPropertyOf:
        owl:topDataProperty


DataProperty: carp:hasValue

    Range:
        xsd:double


DataProperty: prov:startedAtTime

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:qualifiedForm <http://www.w3.org/ns/prov#Start>,
      rdfs:comment "The time at which an activity started. See also prov:endedAtTime."@en,
      prov:editorialNote "It is the intent that the property chain holds: (prov:qualifiedSta
        prov:component "entities-activities",
        prov:qualifiedForm <http://www.w3.org/ns/prov#atTime>,
        prov:category "starting-point",
        rdfs:label "startedAtTime"

    Domain:
        prov:Activity
```

```
    Range:
        xsd:dateTime


DataProperty: prov:atTime

    Annotations:
        prov:unqualifiedForm <http://www.w3.org/ns/prov#invalidatedAtTime>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "The time at which an InstantaneousEvent occurred, in the form of xsd:dat
        prov:unqualifiedForm <http://www.w3.org/ns/prov#generatedAtTime>,
        prov:category "qualified",
      prov:sharesDefinitionWith <http://www.w3.org/ns/prov#InstantaneousEvent>,
        prov:unqualifiedForm <http://www.w3.org/ns/prov#startedAtTime>,
        rdfs:label "atTime",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#endedAtTime>,
        prov:component "entities-activities"

    Domain:
        prov:InstantaneousEvent

    Range:
        xsd:dateTime


DataProperty: carp:hasName
```

```
    Annotations:
        rdfs:label "hasName"@en

    Range:
        xsd:string


DataProperty: owl:topDataProperty


DataProperty: carp:hasRelevanceRating

    Annotations:
        rdfs:label "hasRelevanceRaiting"@en

    Range:
        xsd:double


DataProperty: carp:hasLatitude

    Range:
        xsd:double


DataProperty: carp:hasStartDate

    Range:
```

```
        xsd:dateTime


DataProperty: prov:endedAtTime

    Annotations:
      rdfs:comment "The time at which an activity ended. See also prov:startedAtTime."@en,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "entities-activities",
        rdfs:label "endedAtTime",
        prov:qualifiedForm <http://www.w3.org/ns/prov#End>,
      prov:editorialNote "It is the intent that the property chain holds: (prov:qualifiedEnd
        prov:qualifiedForm <http://www.w3.org/ns/prov#atTime>,
        prov:category "starting-point"

    Domain:
        prov:Activity

    Range:
        xsd:dateTime


DataProperty: prov:invalidatedAtTime

    Annotations:
        prov:qualifiedForm <http://www.w3.org/ns/prov#atTime>,
        prov:category "expanded",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
```

```
      prov:qualifiedForm <http://www.w3.org/ns/prov#Invalidation>,
     rdfs:comment "The time at which an entity was invalidated (i.e., no longer usable)."@e
     prov:editorialNote "It is the intent that the property chain holds: (prov:qualifiedInv
      rdfs:label "invalidatedAtTime",
      prov:component "entities-activities"

    Domain:
        prov:Entity

    Range:
        xsd:dateTime



DataProperty: carp:hasLongitude

    Range:
        xsd:double



DataProperty: carp:hasEndDate

    Range:
        xsd:dateTime



DataProperty: carp:hasWebsite
```

```
DataProperty: prov:generatedAtTime

    Annotations:
        prov:component "entities-activities",
      rdfs:comment "The time at which an entity was completely created and is available for 
        prov:qualifiedForm <http://www.w3.org/ns/prov#Generation>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:category "expanded",
        rdfs:label "generatedAtTime",
      prov:editorialNote "It is the intent that the property chain holds: (prov:qualifiedGe
        prov:qualifiedForm <http://www.w3.org/ns/prov#atTime>

    Domain:
        prov:Entity

    Range:
        xsd:dateTime


Class: carp:Resource

    Annotations:
        rdfs:label "Resource"@en

    SubClassOf:
        carp:hasStartDate some xsd:dateTime,
        carp:hasEndDate some xsd:dateTime,
        carp:hasQualityRating some xsd:double,
```

```
        carp:hasKeyword some xsd:string,
        prov:Entity,
        carp:hasRelevanceRating some rdfs:Literal,
        carp:hasName some xsd:string


Class: carp:FieldSite

    SubClassOf:
        carp:hasName some xsd:string,
        carp:Place



Class: owl:Thing



Class: prov:Bundle

    Annotations:
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-bundle-declaration
        rdfs:label "Bundle",
        prov:category "expanded",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-bundle-entity"^^xsd:an
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      prov:definition "A bundle is a named set of provenance descriptions, and is itself an 
      rdfs:comment "Note that there are kinds of bundles (e.g. handwritten letters, audio re

    SubClassOf:
```

```
        prov:Entity


Class: prov:Usage

    Annotations:
        prov:category "qualified",
      prov:definition "Usage is the beginning of utilizing an entity by an activity. Before
        prov:unqualifiedForm <http://www.w3.org/ns/prov#used>,
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
      rdfs:comment "An instance of prov:Usage provides additional descriptions about the bi
        prov:component "entities-activities",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Usage"^^xsd:anyURI,
        rdfs:label "Usage",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Usage"^^xsd:anyURI

    SubClassOf:
        prov:EntityInfluence,
        prov:InstantaneousEvent


Class: carp:EducationResource

    Annotations:
        rdfs:label "Education Resource"@en

    SubClassOf:
```

```
        carp:Resource


Class: carp:Person

    Annotations:
      rdfs:comment "A person involved or referred to by the resources within a Resource Coll
        rdfs:label "Person"

    SubClassOf:
        prov:Person



Class: prov:Plan

    Annotations:
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Association"^^xsd:
        prov:component "agents-responsibility",
        rdfs:label "Plan",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:comment "There exist no prescriptive requirement on the nature of plans, their re
        prov:category "expanded",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Association"^^xsd:anyU
        prov:category "qualified",
      prov:definition "A plan is an entity that represents a set of actions or steps intende

    SubClassOf:
        prov:Entity
```

```
Class: prov:Delegation

    Annotations:
      prov:definition "Delegation is the assignment of authority and responsibility to an ag

For example, a student acted on behalf of his supervisor, who acted on behalf of the departme
        prov:unqualifiedForm <http://www.w3.org/ns/prov#actedOnBehalfOf>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "Delegation",
        prov:component "agents-responsibility",
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-delegation"^^xsd:a
        prov:category "qualified",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-delegation"^^xsd:anyUR
      rdfs:comment "An instance of prov:Delegation provides additional descriptions about t

    SubClassOf:
        prov:AgentInfluence


Class: carp:Project

    Annotations:
        rdfs:label "Project"@en

    SubClassOf:
        carp:hasProjectInvestigator some carp:ProjectInvestigator,
```

```
        carp:hasEndDate some xsd:dateTime,
        prov:Entity,
        carp:hasStartDate some xsd:dateTime,
        carp:hasName some xsd:string,
        carp:hasResource some carp:Resource


Class: prov:SoftwareAgent

    Annotations:
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-types"^^xsd:anyURI
        prov:category "expanded",
        rdfs:label "SoftwareAgent",
        prov:component "agents-responsibility",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:definition "A software agent is running software."@en,
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-agent"^^xsd:anyURI

    SubClassOf:
        prov:Agent


Class: carp:ResearchResource

    Annotations:
        rdfs:label "Research Resource"@en

    SubClassOf:
```

```
        carp:Resource


Class: carp:Collection

    Annotations:
        rdfs:label "VLC Collection",
        rdfs:comment "A semantic description of a collection of Web resources"

    SubClassOf:
        carp:hasKeyword some xsd:string,
        prov:Collection,
        carp:hasResource some carp:Resource,
        carp:hasName some xsd:string,
        carp:hasOwner some carp:Person


Class: prov:Agent

    Annotations:
        prov:category "starting-point",
       prov:definition "An agent is something that bears some form of responsibility for an a
       prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Agent"^^xsd:anyURI
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "agents-responsibility",
        rdfs:label "Agent",
       prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-agent"^^xsd:anyURI
```

```
    DisjointWith:
        prov:InstantaneousEvent


Class: carp:Endorsement

    SubClassOf:
        prov:Activity


Class: prov:Communication

    Annotations:
       prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Communication"^^xsd:ar
       rdfs:comment "An instance of prov:Communication provides additional descriptions abou
        rdfs:label "Communication",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
       prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-wasInformedBy"^^xs
        prov:component "entities-activities",
       prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasInformedBy>,
       prov:definition "Communication is the exchange of an entity by two activities, one act
        prov:category "qualified"

    SubClassOf:
        prov:ActivityInfluence
```

```
Class: prov:Revision

    Annotations:
      rdfs:comment "An instance of prov:Revision provides additional descriptions about the
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Revision"^^xsd:any
      prov:definition "A revision is a derivation for which the resulting entity is a revise
        rdfs:label "Revision",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasRevisionOf>,
        prov:category "qualified",
        prov:component "derivations",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-revision"^^xsd:anyURI

    SubClassOf:
        prov:Derivation


Class: prov:Activity

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Activity"^^xsd:any
        prov:component "entities-activities",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Activity"^^xsd:anyURI,
        prov:category "starting-point",
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
        rdfs:label "Activity",
      prov:definition "An activity is something that occurs over a period of time and acts u
```

```
    DisjointWith:
        prov:Entity


Class: prov:Association

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "agents-responsibility",
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Association"^^xsd:
      prov:definition "An activity association is an assignment of responsibility to an agen
        prov:category "qualified",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Association"^^xsd:anyU
      rdfs:comment "An instance of prov:Association provides additional descriptions about
        rdfs:label "Association",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasAssociatedWith>

    SubClassOf:
        prov:AgentInfluence


Class: prov:Quotation

    Annotations:
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-quotation"^^xsd:an
        prov:component "derivations",
      prov:definition "A quotation is the repeat of (some or all of) an entity, such as text
```

```
        prov:category "qualified",
        rdfs:label "Quotation",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasQuotedFrom>,
       rdfs:comment "An instance of prov:Quotation provides additional descriptions about th
       prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-quotation"^^xsd:anyURI
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubClassOf:
        prov:Derivation


Class: prov:Attribution

    Annotations:
       prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-attribution"^^xsd:anyU
        prov:component "agents-responsibility",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasAttributedTo>,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
       prov:definition "Attribution is the ascribing of an entity to an agent.

When an entity e is attributed to agent ag, entity e was generated by some unspecified activi
       prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
       prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-attribution"^^xsd:
       rdfs:comment "An instance of prov:Attribution provides additional descriptions about
        prov:category "qualified",
        rdfs:label "Attribution"

    SubClassOf:
```

```
        prov:AgentInfluence



Class: prov:Invalidation

    Annotations:
        prov:component "entities-activities",
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Invalidation"^^xsd
        rdfs:label "Invalidation",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Invalidation"^^xsd:any
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasInvalidatedBy>,
      rdfs:comment "An instance of prov:Invalidation provides additional descriptions about
        prov:category "qualified",
      prov:definition "Invalidation is the start of the destruction, cessation, or expiry of
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubClassOf:
        prov:InstantaneousEvent,
        prov:ActivityInfluence



Class: carp:OutreachResource

    Annotations:
        rdfs:label "Outreach Resource"@en

    SubClassOf:
```

```
        carp:Resource



Class: carp:Institution


    SubClassOf:
        carp:Place



Class: prov:AgentInfluence


    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      prov:editorsDefinition "AgentInfluence is the capacity of an agent to have an effect o
        rdfs:seeAlso <http://www.w3.org/ns/prov#agent>,
        prov:category "qualified",
      rdfs:comment "AgentInfluence provides additional descriptions of an Agent's binary in
        rdfs:label "AgentInfluence",
      rdfs:comment "It is not recommended that the type AgentInfluence be asserted without a

    SubClassOf:
        prov:Influence



Class: prov:End


    Annotations:
      prov:definition "End is when an activity is deemed to have been ended by an entity, kno
```

```
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "End",
        prov:component "entities-activities",
        prov:category "qualified",
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-End"^^xsd:anyURI,
      rdfs:comment "An instance of prov:End provides additional descriptions about the binar
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-End"^^xsd:anyURI,
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasEndedBy>,
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c

    SubClassOf:
        prov:InstantaneousEvent,
        prov:EntityInfluence


Class: carp:ProjectInvestigator

    SubClassOf:
        carp:Person,
        carp:hasInstitution some carp:Institution


Class: prov:EntityInfluence

    Annotations:
      rdfs:comment "EntityInfluence provides additional descriptions of an Entity's binary
        rdfs:seeAlso <http://www.w3.org/ns/prov#entity>,
      prov:editorsDefinition "EntityInfluence is the capacity of an entity to have an effect
```

```
    rdfs:comment "It is not recommended that the type EntityInfluence be asserted without
      rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      rdfs:label "EntityInfluence",
      prov:category "qualified"

    SubClassOf:
        prov:Influence

    DisjointWith:
        prov:ActivityInfluence



Class: carp:TrainingResource

    Annotations:
        rdfs:label "Training Resource "@en

    SubClassOf:
        carp:Resource



Class: prov:Organization

    Annotations:
        rdfs:label "Organization",
      prov:definition "An organization is a social or legal institution such as a company, s
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-types"^^xsd:anyURI
        prov:component "agents-responsibility",
```

```
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-agent"^^xsd:anyURI,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:category "expanded"

    SubClassOf:
        prov:Agent


Class: carp:Place

    SubClassOf:
        carp:hasLongitude some xsd:double,
        carp:hasLatitude some xsd:double


Class: prov:Collection

    Annotations:
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-collection"^^xsd:anyUR
        rdfs:label "Collection",
        prov:component "collections",
        prov:category "expanded",
      prov:definition "A collection is an entity that provides a structure to some constitue
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubClassOf:
        prov:Entity
```

```
Class: prov:Entity

    Annotations:
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Entity"^^xsd:anyUR
        prov:category "starting-point",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "Entity",
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
        prov:component "entities-activities",
      prov:definition "An entity is a physical, digital, conceptual, or other kind of thing
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-entity"^^xsd:anyURI

    DisjointWith:
        prov:InstantaneousEvent, prov:Activity


Class: prov:Start

    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "entities-activities",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Start"^^xsd:anyURI,
      rdfs:comment "An instance of prov:Start provides additional descriptions about the bin
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasStartedBy>,
      prov:definition "Start is when an activity is deemed to have been started by an entity
        prov:category "qualified",
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
```

```
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Start"^^xsd:anyURI
        rdfs:label "Start"

    SubClassOf:
        prov:EntityInfluence,
        prov:InstantaneousEvent


Class: prov:Location

    Annotations:
        prov:category "expanded",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-attribute-location"^^x
        rdfs:seeAlso <http://www.w3.org/ns/prov#atLocation>,
      prov:definition "A location can be an identifiable geographic place (ISO 19112), but i
        rdfs:label "Location",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-attribute"^^xsd:an


Class: prov:Person

    Annotations:
        rdfs:label "Person",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-agent"^^xsd:anyURI,
        prov:category "expanded",
        prov:definition "Person agents are people."@en,
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
```

```
        prov:component "agents-responsibility",
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-types"^^xsd:anyURI

    SubClassOf:
        prov:Agent


Class: prov:Role

    Annotations:
        prov:component "agents-responsibility",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "Role",
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-attribute"^^xsd:an
        rdfs:seeAlso <http://www.w3.org/ns/prov#hadRole>,
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-attribute-role"^^xsd:a
      prov:definition "A role is the function of an entity or agent with respect to an activi
        prov:category "qualified"


Class: prov:EmptyCollection

    Annotations:
        prov:component "collections",
      prov:definition "An empty collection is a collection without members."@en,
        prov:category "expanded",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "EmptyCollection"@en
```

```
    SubClassOf:
        prov:Collection


Class: prov:Generation

    Annotations:
        prov:component "entities-activities",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "Generation",
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-Generation"^^xsd:a
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasGeneratedBy>,
        prov:category "qualified",
      prov:definition "Generation is the completion of production of a new entity by an acti
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Generation"^^xsd:anyUR
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
      rdfs:comment "An instance of prov:Generation provides additional descriptions about t

    SubClassOf:
        prov:ActivityInfluence,
        prov:InstantaneousEvent


Class: prov:Influence

    Annotations:
      rdfs:comment "An instance of prov:Influence provides additional descriptions about th
```

```
    prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-influence"^^xsd:anyURI
    rdfs:comment "Because prov:Influence is a broad relation, its most specific subclasses
      rdfs:label "Influence",
      prov:category "qualified",
      prov:unqualifiedForm <http://www.w3.org/ns/prov#wasInfluencedBy>,
    prov:definition "Influence is the capacity of an entity, activity, or agent to have an
      prov:component "derivations",
      rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
    prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-influence"^^xsd:an


Class: prov:InstantaneousEvent

  Annotations:
    prov:definition "The PROV data model is implicitly based on a notion of instantaneous
      rdfs:label "InstantaneousEvent",
    rdfs:comment "An instantaneous event, or event for short, happens in the world and mar
      prov:component "entities-activities",
      rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
    prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#dfn-event
      prov:category "qualified"

  DisjointWith:
      prov:Entity, prov:Agent


Class: prov:Derivation
```

```
    Annotations:
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        prov:component "derivations",
      rdfs:comment "An instance of prov:Derivation provides additional descriptions about t
      rdfs:comment "The more specific forms of prov:Derivation (i.e., prov:Revision, prov:Q
        rdfs:label "Derivation",
        prov:unqualifiedForm <http://www.w3.org/ns/prov#wasDerivedFrom>,
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#Derivation-Relation"^^xsd:any
      prov:definition "A derivation is a transformation of an entity into another, an update
      prov:constraints "http://www.w3.org/TR/2013/REC-prov-constraints-20130430/#prov-dm-c
        prov:category "qualified",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-Derivation"^^xsd:anyUR

    SubClassOf:
        prov:EntityInfluence


Class: prov:PrimarySource

    Annotations:
      prov:n "http://www.w3.org/TR/2013/REC-prov-n-20130430/#expression-original-source"^^
      prov:definition "A primary source for a topic refers to something produced by some age

Because of the directness of primary sources, they 'speak for themselves' in ways that cannot

A primary source relation is a particular case of derivation of secondary materials from thei
      rdfs:comment "An instance of prov:PrimarySource provides additional descriptions abou
        prov:category "qualified",
```

```
        prov:unqualifiedForm <http://www.w3.org/ns/prov#hadPrimarySource>,
        rdfs:label "PrimarySource",
      prov:dm "http://www.w3.org/TR/2013/REC-prov-dm-20130430/#term-primary-source"^^xsd:a
        prov:component "derivations",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubClassOf:
        prov:Derivation



Class: carp:FundingAgency

    SubClassOf:
        carp:Institution



Class: prov:ActivityInfluence

    Annotations:
        rdfs:label "ActivityInfluence",
        rdfs:seeAlso <http://www.w3.org/ns/prov#activity>,
        prov:category "qualified",
      rdfs:comment "It is not recommended that the type ActivityInfluence be asserted withou
      prov:editorsDefinition "ActivitiyInfluence is the capacity of an activity to have an e
      rdfs:comment "ActivityInfluence provides additional descriptions of an Activity's bin
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>

    SubClassOf:
```

```
        prov:Influence,
        prov:hadActivity max 0 owl:Thing

    DisjointWith:
        prov:EntityInfluence


Individual: carp:Joaquin

    Types:
        carp:Person


Individual: carp:NVR

    Annotations:
        rdfs:label "Natalia Villanueva Rosales"@en

    Types:
        carp:Person


Individual: carp:SemanticWebWorkingGuide

    Types:
        carp:TrainingResource
```

```
Individual: carp:FromProgramSynthesisToOptimalProgramSynthesis

    Types:
        carp:ResearchResource


Individual: <http://www.w3.org/ns/prov-o-20130312>


Individual: carp:IntroToSemanticWeb

    Types:
        carp:EducationResource


Individual: carp:NSF

    Types:
        carp:FundingAgency


Individual: carp:NVR-SemanticWeb

    Annotations:
        rdfs:label "Semantic Web"@en,
        rdfs:comment "A collection about semantic web resources."

    Types:
```

```
        carp:Collection

    Facts:
     carp:hasResource  carp:IntroToSemanticWeb,
     carp:hasResource  carp:SemanticWebWorkingGuide,
     carp:hasOwner  carp:NVR,
     carp:hasKeyword  "data integration",
     carp:hasKeyword  "semantic web",
     carp:hasName  "SemanticWeb"


Individual: carp:JoaquinProject

    Types:
        carp:Project

    Facts:
     carp:hasFieldSite  carp:WhiteSands,
     carp:hasFundingAgency  carp:NSF,
     carp:hasProjectInvestigator  carp:Joaquin,
     carp:hasName  "Joaquin Project",
     carp:hasLongitude  "-77.037852"^^xsd:double,
     carp:hasStartDate  "2001-10-26T21:32:52"^^xsd:dateTime,
     carp:hasLatitude  "38.898556"^^xsd:double,
     carp:hasWebsite  "scidesign-test.utep.edu/JoaquinProject",
     carp:hasEndDate  "2001-10-26T21:32:52+02:00"^^xsd:dateTime
```

```
Individual: carp:IntroToCS

    Types:
        carp:EducationResource


Individual: prov:EmptyCollection

    Annotations:
        prov:component "collections",
      prov:definition "An empty collection is a collection without members."@en,
        prov:category "expanded",
        rdfs:isDefinedBy <http://www.w3.org/ns/prov-o#>,
        rdfs:label "EmptyCollection"@en


Individual: carp:Miproyecto

    Annotations:
        rdfs:label "Miproyecto"@en

    Types:
        carp:Project

    Facts:
     carp:hasPart  carp:NVR-SemanticWeb
```

```
Individual: <http://www.w3.org/ns/prov-o>


Individual: carp:WhiteSands

    Types:
        carp:FieldSite


Individual: <http://www.w3.org/ns/prov-o#>

    Annotations:
        prov:specializationOf <http://www.w3.org/ns/prov-o>,
        prov:wasRevisionOf <http://www.w3.org/ns/prov-o-20130312>


Individual: carp:JOaquin-collection

    Annotations:
        rdfs:label "Joaquin-collection"@en

    Types:
        carp:Collection

    Facts:
     carp:hasResource  carp:IntroToCS,
     carp:hasOwner  carp:Joaquin,
     carp:hasPart  carp:Miproyecto,
```

```
carp:hasResource  carp:FromProgramSynthesisToOptimalProgramSynthesis
```

Appendix B

IRB documents

The follwing is the IRB approval for this project.

THE UNIVERSITY OF TEXAS AT EL PASO
Office of the Vice President for Research and Sponsored Projects
Institutional Review Board
El Paso, Texas 79968-0587
phone: 915 747-8841 fax: 915 747-5931

FWA No: 00001224

DATE:	November 11, 2015
TO:	Joaquin Reyna
FROM:	University of Texas at El Paso IRB
STUDY TITLE:	[815967-1] A Semantically-Enabled Trust Model for Collaborative Environments
IRB REFERENCE #:	College of Engineering
SUBMISSION TYPE:	New Project
ACTION:	APPROVED
APPROVAL DATE:	November 11, 2015
EXPIRATION DATE:	November 10, 2016
REVIEW TYPE:	Expedited Review

Thank you for your submission of New Project materials for this research study. University of Texas at El Paso IRB has APPROVED your submission. This approval is based on an appropriate risk/benefit ratio and a study design wherein the risks have been minimized. All research must be conducted in accordance with this approved submission.

This study has received Expedited Review based on the applicable federal regulation.

Please remember that informed consent is a process beginning with a description of the study and insurance of participant understanding followed by a signed consent form. Informed consent must continue throughout the study via a dialogue between the researcher and research participant. Federal regulations require each participant receive a copy of the signed consent document.

Please note that any revision to previously approved materials must be approved by this office prior to initiation. Please use the appropriate revision forms for this procedure.

All SERIOUS and UNEXPECTED adverse events must be reported to this office. Please use the appropriate adverse event forms for this procedure. All FDA and sponsor reporting requirements should also be followed.

Please report all NON-COMPLIANCE issues or COMPLAINTS regarding this study to this office.

Please note that all research records must be retained for a minimum of three years after termination of the project.

Based on the risks, this project requires Continuing Review by this office on an annual basis. Please use the appropriate renewal forms for this procedure.

If you have any questions, please contact the IRB Office at (915) 747-8841 or irb.orsp@utep.edu. Please include your study title and reference number in all correspondence with this office.

cc:

Curriculum Vitae

Joaquin Andres Reyna was born on January 28, 1988. The first son of Joaquin Alberto Reyna and Blanca Cruz, he graduated from El Chamizal High School, Ciudad Juarez, Chihuahua, Mexico, in the spring of 2006. He entered El Paso Community College University in the fall of 2006, and, in the spring of 1982, The University of Texas at El Paso. While pursuing his bachelor's degree in Computer Science he worked as a Technical Assistant at Lockheed Martin Storefront under The Facilities Development and Operations (FDOC) NASA contract. He received his bachelor's degree in Computer Science in spring of 2012.

In the fall of 2012, he entered the Graduate School of The University of Texas at El Paso. While pursuing a master's degree in Computer Science he worked as a Research Assistant in the Cyber-Share center.

Permanent address: 352 N Carolina Dr 222
El Paso, Texas 79915