Extracting Trust from Domain Analysis: a Case Study on Wikipedia Project

Pierpaolo Dondio[1], Stephen Barrett[1], Stefan Weber[1] and Jean-Marc Seigneur[2]

[1] School of Computer Science and Statistics, Distributed System Group,
Trinity College Dublin, College Green 2, Dublin
{dondiop, stephen.barrett, stefan.weber} @cs.tcd.ie
[2] University of Geneva, CUI,
24 Rue du General Dufour, Geneva 4, Switzerland
Jean-Marc.Seigneur@trustcomp.org

Abstract. The problem of identifying trustworthy information on the World Wide Web is becoming increasingly acute as new tools such as wikis and blogs simplify and democratize publications. Wikipedia is the most extraordinary example of this phenomenon and, although a few mechanisms have been put in place to improve contributions quality, trust in Wikipedia content quality has been seriously questioned. We thought that a deeper understanding of what in general defines high-standard and expertise in domains related to Wikipedia – i.e. *content quality* in a *collaborative environment* – mapped onto Wikipedia elements would lead to a complete set of mechanisms to sustain trust in Wikipedia context. Our evaluation conducted on about 8,000 articles, representing 65% of the overall Wikipedia editing activity, shows that the new trust evidence that we extracted from Wikipedia allows us to transparently and automatically compute trust values to isolate articles of great or low quality.

1 Introduction

In the famous 1996 article *Today's WWW, Tomorrow's MMM: The specter of multimedia mediocrity* [1] Cioleck predicted a seriously negative future for online content quality by describing the World Wide Web (WWW) as a nebulous, ever-changing multitude of computer sites that house continually changing chunks of multimedia information, the global sum of the uncoordinated activities of several hundreds of thousands of people who deal with the system as they please. Thus, the WWW may come to be known as the MMM (MultiMedia Mediocrity). Despite this vision, it is not hard to predict that the potential and the growth of the Web as a source of information and knowledge will increase rapidly. The Wikipedia project, started in January 2001, represents one of the most successful and discussed example of such phenomenon, an example of *collective knowledge*, a concept that is often lauded as the next step toward truth in online media.

On one hand, recent exceptional cases have brought to the attention the question of Wikipedia trustworthiness. In an article published on the 29th of November in USA Today [2], Seigenthaler, a former administrative assistant to Robert Kennedy, wrote about his anguish after learning about a false Wikipedia entry that listed him as hav-

ing been briefly suspected of involvement in the assassinations of both John Kennedy and Robert Kennedy. The 78-year-old Seigenthaler got Wikipedia founder Jimmy Wales to delete the defamatory information in October. Unfortunately, that was four months after the original posting. The news was further proof that Wikipedia has no accountability and no place in the world of serious information gathering [2].

On the other hand, Wikipedia is not only being negatively discussed. In December 2005, a detailed analysis carried out by the magazine *Nature* [3] compared the accuracy of Wikipedia against the Encyclopaedia Britannica. Nature identified a set of 42 articles, covering a broad range of scientific disciplines, and sent them to relevant experts for peer review. The results are encouraging: the investigation suggests that Britannica's advantage may not be great, at least when it comes to science entries. The difference in accuracy was not particularly great: the average science entry in Wikipedia contained around four inaccuracies; Britannica, about three. Reviewers also found many factual errors, omissions or misleading statements: 162 and 123 in Wikipedia and Britannica respectively. Moreover, Nature has stated that, among their scientific collaborators, 70% of them had heard of Wikipedia, 17% of those consult it on a weekly basis and about 10% help to update it.

This paper seeks to face the problem of the trustworthiness of Wikipedia by using a computational trust approach: our goal is to set up an automatic and transparent mechanism able to estimate the trustworthiness of Wikipedia articles. The paper is organized as follows: in the next section 2 we review related work on trust and content quality issues; in section 3 we argue that, due to the fast changing nature of articles, it is difficult to apply the trust approaches proposed in related work. In section 4 this discussion will lead us to introduce our approach, that starts from an in-depth analysis of *content quality* and *collaborative editing* domains to give us a better understanding of what can support trust in these two Wikipedia related fields. In section 5 we map conclusions of the previous section onto elements extracted directly from Wikipedia in order to define a new set of sources of trust evidence. In section 6 we present our experiment and evaluation conducted on almost 8,000 Wikipedia articles selected among the most edited and visited, that by themselves represent 65% of the editing activity and 50% of visits of the overall encyclopedia. Section 7 will collect our conclusions and future work, where we anticipate our intention to organize the method used in this paper into a general trust methodology that can be applied to other application domains.

2 Related Work

There are many definitions of the human notion trust in a wide range of domains from sociology, psychology to political and business science, and these definitions may even change when the application domains change. For example, Romano's recent definition tries to encompass the previous work in all these domains: "trust is a subjective assessment of another's influence in terms of the extent of one's perceptions about the quality and significance of another's impact over one's outcomes in a given situation, such that one's expectation of, openness to, and inclination toward such influence provide a sense of control over the potential outcomes of the situation."[4]

However, the terms trust/trusted/trustworthy, which appear in the traditional computer security literature, are not grounded on social science and often correspond to an implicit element of trust. Blaze et al [5] first introduced “decentralized trust management" to separate trust management from applications. PolicyMaker [6] introduced the fundamental concepts of policy, credential, and trust relationship. Terzis et al. [7] have argued that the model of trust management [5,6] still relies on an implicit notion of trust because it only describes “a way of exploiting established trust relationships for distributed security policy management without determining how these relationships are formed”. Computational trust was first defined by S. Marsh [8], as a new technique able to make agents less vulnerable in their behavior in a computing world that appears to be malicious rather than cooperative, and thus to allow interaction and cooperation where previously there could be none. A computed trust value in an entity may be seen as the digital representation of the trustworthiness or level of trust in the entity under consideration. The EU project SECURE [9] represents an example of a trust engine that uses evidence to compute trust values in entities and corresponds to evidence-based trust management systems. Evidence encompasses outcome observations, recommendations and reputation. Depending on the application domain, a few types of evidence may be more weighted in the computation than other types. When recommendations are used, a social network can be reconstructed. Golbeck [10] studied the problem of propagating trust value in social networks, by proposing an extension of the FOAF vocabulary [11] and algorithms to propagate trust values estimated by users rather than computed based on a clear count of pieces of evidence. Recently, even new types of evidence have been proposed to compute trust values. For example, Ziegler and Golbeck [12] studied interesting correlation between similarity and trust among social network users: there is indication that similarity may be evidence of trust. In SECURE, evidence is used to select which trust profile should be given to an entity. Thus similar evidence should lead to similar profile selection. However, once again, as for human set trust value, it is difficult to clearly estimate people similarity based on a clear count of pieces of evidence. However, the whole SECURE framework may not be generic enough to be used with abstract or complex new types of trust evidence. In fact, in this paper, we extracted a few types of evidence present in Wikipedia (detailed in the next sections) that did not fit well with the SECURE framework and we had to build our own computational engine. We think that our approach to deeply study the domain of application and then extract the types of trust evidence from the domain is related to the approach done in expert systems where the knowledge engineer interacts with an expert in the domain to acquire the needed knowledge to build the expert system for the application domain. Lerch et al. [13] highlighted the impact of trust in expert systems advices. Ball et al. [14] proposed an expert system that has knowledge about the factors that are important in computing the trust in a certification authority used in a public key infrastructure. It shows that there are different application domains where our approach could be used and may indicate that to formalize a methodology based on our approach may be useful. However, in this paper, we focus on trust computation for content quality and Bucher [15] clearly motivates our contribution in this paper because he argues that on the Internet “we no longer have an expert system to which we can assign management of information quality”. We finish this section by two last computational projects related to content quality in a decentralised publishing system. Huang and Fox in [16] propose a

metadata-based approach to determine the origin and validity of information on the Web. They provide a metadata language for expressing judgments on the trustworthiness of a proposition, its author or a specific field. Then, by maintaining information sources and dependencies, it is possible to compute the trustworthiness of a derived proposition based on the trustworthiness of those composed. In this work the key hypothesis is the validity of the basic assertion and their maintenance. Finally, Guha [17] presents a model to integrate user-driven ratings on top of e-services, as it has been done in Epinion, Amazon or Slashdot but not really for Wikipedia. In addition, we could not integrate our new types of trust evidence in Guha's model that merely focuses on recommendations and reputation propagated within the social network formed by the users.

3 The Problem of Wikipedia Articles Trustworthiness

Wikipedia shows intrinsic characteristics that make the utilization of trust solutions challenging. The main feature of Wikipedia, appointed as one of its strongest attribute, is the speed at which it can be updated. The most visited and edited articles reach an average editing rate of 50 modifications per day, while articles related to recent news can reach the number of hundreds of modifications. This aspect affects the validity of several trust techniques.

Human-based trust tools like feedback and recommendation systems require time to work properly, suffering from a well know ramp-up problem [7]. This is a hypothesis that clashes with Wikipedia, where pages change rapidly and recommendations could dramatically lose meaning. Moreover, the growing numbers of articles and their increasing fragmentation require an increasing number of ratings to keep recommendations significant. *Past-evidence trust paradigm* relies on the hypothesis that the trustor entity had enough past interactions with the trustee to collect a significant set of evidences. In Wikipedia the fact that past versions of a page are not relevant for assessing present trustworthiness and the changing nature of articles makes it difficult to compute trust values based on past evidences. In general, user past-experience with a Web site is only at 14th position among the criteria used to assess the quality of a Web site with an incidence of 4.6% [19]. We conclude that a mechanism to evaluate articles trustworthiness relying exclusively on their present state is required. We thought that such a method could be identified by a deeper understanding of the domains involved in Wikipedia, namely, the content quality domain and the collaborative editing domain. After understanding what brings trust in those domains, we mapped these set of evidences into Wikipedia elements that we previously isolated by defining a detailed model of the application. This resulting new set of evidence, extracted directly from Wikipedia, is source of trust, since it relies on proven domains' expertise. Through an evaluation phase, we exploited this set of evidences to support trust calculation and estimate the trustworthiness of articles.

4 Wikipedia Domain Analysis

We identified two relevant areas involved in Wikipedia: the content quality domain and collaborative working domain, in our case a *collaborative editing*. In this section, we analyze what can bring high quality in these two domains.

The quality of online content is a critical problem faced by many institutions. Alexander [20] underlines how information quality is a slippery subject, but it proposes hallmark of what is consistently good information. He identified three basic requirements: objectivity, completeness and pluralism. The first requirement guarantees that the information is unbiased, the second assesses that the information should not be incomplete, and the third stresses the importance of avoiding situations in which information is restricted to a particular viewpoint. University of Berkeley proposes a practical evaluation method [21] that stresses the importance of considering authorship, timeliness, accuracy, permanence and presentation. Authorship stresses the importance of collecting information on the authors of the information, accuracy deals with how the information can be considered good, reviewed, well referenced and if it is comparable to similar other Web content, in order to check if it is compliant to a standard. Timeliness considers how the information has changed during time: its date of creation, its currency and the rate of its update; permanence stresses how the information is transitory or stable. Finally, presentation concerns the layout of the text, the balance among its sections, the presence of images and the quality of the layout. In a study already cited [19], presentation resulted in the most important evaluation criterion with an incidence of 46%.

The Persuasive Technology Lab has been running the Stanford Web Credibility Research since 1997 to identify which are the sources of credibility and expertise in Web content. Among the most well-known results are the ten guidelines for Web credibility [22], compiled to summarize what brings credibility and trust in a Web site. The guidelines confirm what we described so far and again they emphasize the importance of the non anonymity of the authors, the presence of references, the importance of the layout, the constant updating and they underline how typographical errors and broken links, no matter how small they could be, strongly decrease trust and represent evidence of lack of accuracy.

Beside content quality domain, Wikipedia cannot be understood if we do not take into consideration that it is done entirely in a collaborative way. Researches in *collaborative working* [23] help us to define a particular behaviour strongly involved in Wikipedia dynamics, the balance in the editing process. A collaborative environment is more effective when there is a kind of emerging leadership among the group; the leadership is able to give a direction to the editing process and avoid fragmentation of the information provided. Anyway, this leadership should not be represented by one or two single users to avoid the risk of lack of pluralism and the loss of collaborative benefits like merging different expertise and points of view. We summarize our analysis with the prepositions shown in table 1: in the first column are theoretical propositions affecting trust, second column lists the domains from which each preposition was taken.

Table 1. A Trust domain-compatible theory. CQ is Content Quality domain and CE is Collaborative Editing domain

	Prepositions about Trustworthiness of an article (T). T increases if the article…	**Domain**
1	was written by expert and identifiable authors	CQ
2	has similar features or it is compliant to a standard in its category	CQ
3	there is a clear leadership/direction in the group directing the editing process and acting like a reference	CE
4	there is not a dictatorship effect, where most of the editing reflect one person's view	CQ/CE
5	the fragmentation of the contributions is limited: there is more cohesion than dissonance among authors	CE
6	has good balance among its sections, the right degree of details, it contains images if needed, it has a varied sentence structure, rhythm and length	CQ
7	is constantly visited and reviewed by authors	CQ
8	the article is stable	CQ
9	use a neutral point of view	CQ
10	the article is well referenced	CQ

Preposition 1 covers the authorship problem. Preposition 2 derives from the accuracy issues. Preposition 3, 4 and 5 underline the importance that the article should have a sense of unity, even if written by more than one author. Preposition 7 underlines the fact that a good article is constantly controlled and reviewed by a reasonable high number of authors. Preposition 8 stresses the stability of the article: a stable text means that it is well accepted, it reached a consensus among the authors and its content is almost complete. Preposition 9 emphasizes the risk, especially for historical or political issues, that different authors may express personal opinions instead of facts, leading to a subjective article or controversial disputes among users. In order to have meaning, these prepositions need to be considered together with their interrelationships along with some conditions. For example, the length of an article needs to be evaluated in relation to the popularity and importance of its subjects, to understand if the article is too short, superficial or too detailed; the stability of an article has no meaning if the article is rarely edited, since it could be stable because it is not taken in consideration rather than because it is complete.

5 Mapping Domains Theories Onto Wikipedia Elements.

We first need a model of Wikipedia in order to extract elements useful for our purpose. Wikipedia has been designed so that any past modification, along with information about the editor, is accessible. This transparency, that by itself gives an implicit sense of trust, allows us to collect all the information and elements needed.

Our Wikipedia model is composed of two principal objects (*Wiki Article* and *Wiki User*) and a number of supporting objects, as depicted in fig. 1. Since each user has a personal page, user can be treated as an article with some editing methods like *creat-*

ing, modifying and deleting article or uploading images. An article contains the main text page (class *wiki page*) and the *talk page*, where a user can add comments and judgments on the article. *Wiki pages* include properties such as its length, a count of the number of sections, images, external links, notes, and references. Each page has a history page associated, containing a complete list of all modifications. A modification contains information on User, date and time and article text version.

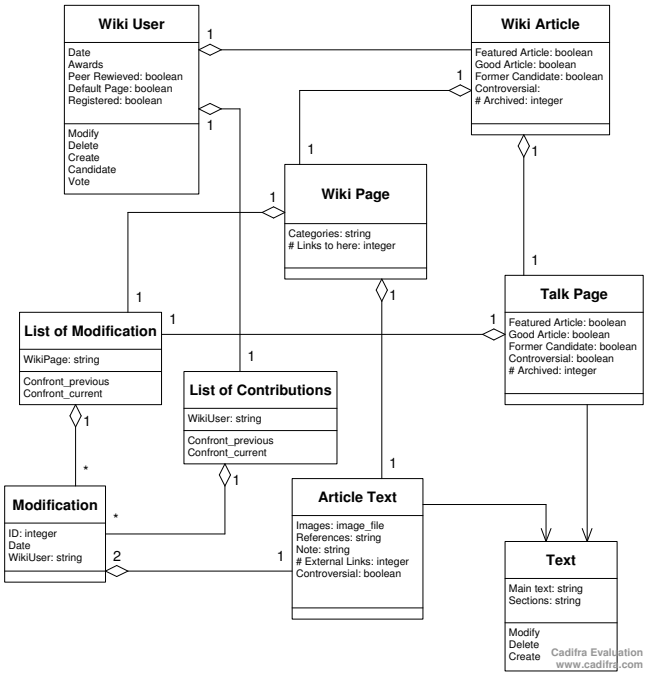

Fig. 1. The Wikipedia UML model.

The community of users can modify articles or adding discussion on article's topic (the *talk page* for that article). We are now ready to map the proposition listed in table 1 onto elements of our Wikipedia model. We outlined 8 macro-areas and we identified about 50 trust factors. As an example, we described how we model the two macro-areas *User's Distribution* and *Stability* (a complete list is available here [24]).

Users' distribution/Leadership (Propositions 3,5,9)

We model this issue with following formulas. We define *E(u)* as the number of edits for user u for article *w*. We define:

$$T(w) : w \rightarrow \Im \tag{1}$$

$$P(n) : [0..1] \rightarrow \Im \ , \ P(n) = \sum_{U_a} E(u) \tag{2}$$

Where U_a is the set of n% most active users in *U(w)*. T represent the total number of edits for article *w* while *P(n)*, given a normalized percentage *n*, returns the number of edits done by the top n% most active users among the set *U(w)*. Similar to *P(n)* is

$$Pe(n):\Im \to \Im \ , \ Pe(n)=\sum_{U_n} E(u) \ , \ U_n=\{u \in U \mid E(u)>n\} \quad (3)$$

that, given a number of edits *n*, represent the number of edits done by users with more than n edits. The different between *P* and *Pe* is that *P* considers the most active users in relation to the set of users contributing to the article, while *Pe* considers the most active users in relation to an absolute number of edit n. Table 3 explains our considerations for each formula:

Table 2. User Distribution Factors

Trust Factors	Comments
Average of *E*	Average number of edits per user.
Standard Deviation of *E*	Standard deviation of edits
$\frac{P(n)}{T}$	% of edits produced by the most active users
$\frac{Pe(t)}{T}$	% of edits produced by users with more than n edit for that article
Number of discussions (talk edit)	It represents how much an article is discussed

We explain the meaning of the functions defined: P(n)/T tells us how much of the article has been done by a subset of users. If we pose n=5 and we obtain:

$$\frac{P(5)}{T}=0.45$$

this means that the 45% of the edits have been done by the top 5% of the most active users. If the value is low the leadership for that article is low, if it is high it means that a relatively small group of users is responsible for most of the editing process.
We have introduced the function *Pe(n)/T* to evaluate leadership from a complementary point of view. *Pe(n)/T* is the percentage of edits done by users that did more than n edits for the article. If we pose n=3 and we obtain

$$\frac{Pe(3)}{T}=0.78$$

it means that 78% of the edits where done by users with more than 3 edits and only 22% by users that did 1,2 or 3 edits. Thus, 1-Pe(n)/T with n small (typically 3) indicates how much of the editing's process was done by occasional users, with a few edits. Thus, it can represent a measurement of the fragmentation of the editing process.
The average and standard deviation of the function *E(u)* (total edits per user u) reinforces the leadership as well: average close to 1 means high fragmentation, high standard deviation means high leadership. Finally, we model the dictatorship effect (Propositions 3,9) by using the function *P(n)/T* keeping n very small (<1) to verify if a very small group of users (typically no more than 5) did a great percentage of edits. We also consider the standard deviation of *E(u)*: very high values can represent a strong predominance of few users.

Stability (Proposition 8)

We define the function

$$N(t) : t \rightarrow \Im \tag{4}$$

That gives the number of edits done at time t. Then we define:

$$Et(t) = \sum_{t}^{P} N(t) \tag{5}$$

that, given time t it gives the number of edits done from time t to the present time *P*. We then define

$$Txt(t) : t- > \Im \tag{6}$$

that gives the number of words that are different form the version at time *t* and the current one. We define *U* as the largest period of time for that article, i.e. its age. We define L as the number of words in the current version.

Table 3. Article's Stability Factors

Trust Factors	Comments
$\frac{Et(t)}{U}$	Rate of edits from time t
$\frac{Txt(t)}{L}$	Percentage of text different from current version and version at time t

We evaluate the stability of an article looking at the values of these two functions. If an article is stable it means that *Et*, from a certain point of time *t*, should decrease or be almost a constant that means that the number of editing is stable or decreasing: the article is not being to be modified. The meaning of *Txt(t)* is an estimation of how different was the version at time *t* compared to the current version. When *t* is close to the current time point, *Txt* goes to 0, and it is obviously 0 when t is the current time. An article is stable if *Txt*, from a certain point of time *t* not very close to the current time is almost a constant value. This means that the text is almost the same in that period of time. As mentioned above, an article can be stable because it is rarely edited, but this may mean it is not taken in consideration rather than it is complete. To avoid this, the degree of activity of the article and its text quality are used as a logic condition for stability: only active and articles with good text can be considered stable.

6 Evaluation

We selected a case study and developed a working prototype in C able to calculate the trustworthiness of a Wikipedia article. A diagram of the prototype is depicted in figure 2. The system, using the *factors updater* module, is continuously fed by Wikipe-

dia, and it stores the results in the factor DB. The Wikipedia database is completely available for download here [25].

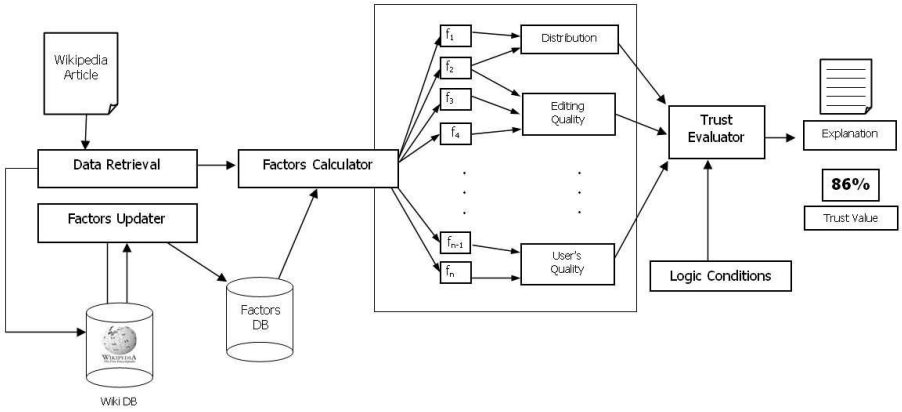

Fig. 2. The Wikipedia Trust Calculator.

When we want to estimate the trustworthiness of an article, the Data Retrieval module query the Wikipedia DB (it could retrieve information directly from the web site as well), and it collects the needed data: article page, talk page, modification lists, user's list, article's category and old versions. Then, the factors calculator module calculates each of the trust factors, merging them into the macro-areas defined. Using the values contained in the Factors DB of pages of the same category or comparable it computes a ranking of the page for each macro-area. Finally, the trust evaluator module estimates a numeric trust values and a natural language explanation of the value. The output is achieved by merging the partial trust value of each macro areas using constraints taken from the Logic Conditions module. This contains logic conditions that control the meaning of each trust factor in relationship to the others:

- *IF leadership is high AND dictatorship is high THEN warning*
- *IF length is high AND importance is low THEN warning*
- *IF stability is high AND (length is short OR edit is low OR importance is low) THEN warning*

By looking at the page rank in each macro area and considering the warnings coming from the logic condition module, explanations like the following can be provided:
We evaluated our model an extensive calculation over almost 8,000 articles. The experiment was conducted on the 17th of March 2006 on 7 718 Wikipedia articles. These articles included all 846 featured articles plus the most visited pages with at least 25 edits. These articles represent the 65% of the editing activity of Wikipedia and the high majority of its access, thus it can be considered a significant set. The results are summarized in graph 1. The graph represents the distribution of the articles on the base of their trust values. We have isolated the featured articles (grey line) from standard articles (black line). Results obtained are positive and encouraging: graph 1 clearly shows the difference between standard articles distribution, mainly around a trust value of 50%, and the featured articles one, around 75%.

The 77.8% of the featured articles are distributed in the region with trust values > 70%, meaning that they are all considered good articles, while only 13% of standard articles are considered good. Furthermore, 42,3% of standard articles are distributed in the region with trust values < 50%, where there are no featured articles, demonstrating the selection operated by the computation. Only 23 standard articles are in the region >85%, where there are 93 featured ones. The experiment, covering articles from different categories, was conducted on an absolute scale, and it shows a minimal imprecision if compared with a previous experiment conducted on a set of 200 articles from the same category "*nations*" [24], where we relied on relative comparisons of similar articles. This shows that the method has a promising general validity.

Fig. 3. Articles and Featured Articles distribution by Trust Value.

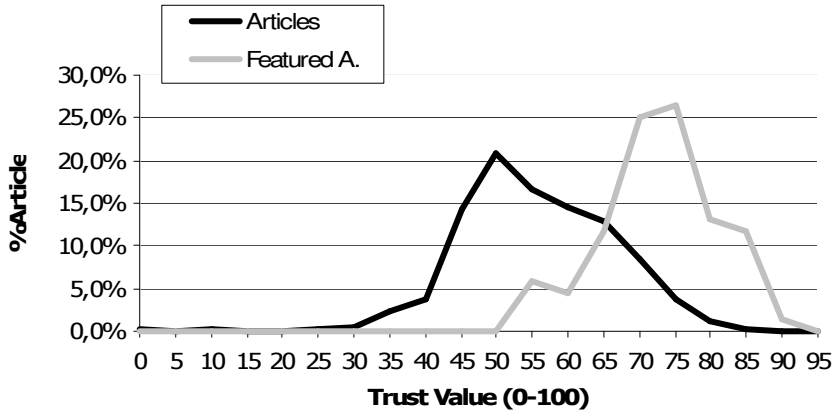

7 Conclusion and Future Work

In this paper we have proposed a transparent, non invasive and automatic method to evaluate the trustworthiness of Wikipedia articles. The method was able to estimate the trustworthiness of articles relying only on their present state, a characteristic needed to cope with the changing nature of Wikipedia. After having analyzed what brings credibility and expertise in the domains composing Wikipedia, i.e. content quality and collaborative working, we identified a set of new trust sources, trust evidence, to support our trust computation. The experimental evidence that we collected from almost 8,000 pages covering the majority of the encyclopedia activity lead to promising results. This suggests a role for such a method in the identification of trustworthy material on the Web.

Our future work will be focused on the generalization of the scheme proposed in the paper into a formal evidence-based trust methodology called Domain ANalysis/Trust Extraction (DANTE). DANTE will be centered on the idea that in a particular domain humans have developed theories and expertise that have been proved to be effective to guarantee gold-standard and high quality in that domain. We refer to this set of expert-knowledge, proven theories, modus operandi as *domain related theories*. We

trust an entity that can show evidence related to these domain-related theories. The more the entity proves to be associated with such evidence, the more it is trustworthy. Theories define what is trustworthy in that domain, entity evidence related to these theories allow us to compute the degree of trustworthiness.

References

1. Ciolek, T.: M. Today's WWW, Tomorrow's MMM: The specter of multi-media mediocrity, IEEE COMPUTER, January 1996, Vol 29(1) pp. 106-108
2. How much do you trust wikipedia?, article from the website http://news.com.com/2009-1025_3-5984535.html, visited March 2006
3. Gales, J.: Encyclopedias goes head a head, Nature Magazine, issue N. 438, 15 December 2005
4. Romano, D. M.: The Nature of Trust: Conceptual and Operational Clarification, Louisiana State University. PhD Thesis (2003)
5. Blaze, M., Feigenbaum, J. and Lacy, J.: Decentralized Trust Management, Proceedings of IEEE Conference on Security and Privacy (1996)
6. Chu, Y.: Trust Management for the World Wide Web, Master Thesis, MIT (1997)
7. Terzis, S., W. Wagealla: The SECURE Collaboration Model, SECURE Deliverables D2.1, D.2.2 and D2.3, 2004
8. Marsh, S.: Formalizing Trust as a Computational Concept. PhD thesis, University of Stirling, Department of Computer Science and Mathematics (1994)
9. V. Cahill, et al. Using Trust for Secure Collaboration in Uncertain Environments. IEEE Pervasive Computing Magazine, (July-September 2003)
10. Golbeck, J., Hendler, J., and Parsia, B.: Trust Networks on the Semantic Web, University of Maryland, College Park, (2002)
11. www.foaf-project.com, FOAF project website
12. Ziegler, C., Golbeck, J.: Investigating Correlations of Trust and Interest Similarity - Do Birds of a Feather Really Flock Together", Decision Support Services, (2005)
13. Lerch, J., M. Prietula & C. Kulik. (1997). The Turing Effect: The nature of trust in expert systems advice. In P. Feltovich.
14. Ball, E., Chadwick, D., Basden, A.: The Implementation of a System for Evaluating Trust in a {PKI} Environment, Trust in the Network Economy, SpringerWein (2003) 263-279.
15. Bucher, H.: Crisis Communication and the Internet: Risk and Trust in a Global Media, First Monday, Volume 7, Number 4 (2002)
16. Huang, J., Fox, S.: Uncertainty in knowledge provenance, Proceedings of the first European Semantic Web Symposium, Heraklio, Greece (May 2004)
17. Guha, R. Open Rating Systems. Technical Report, Stanford University (2003)
18. Burke, R. Knowledge-based Recommender Systems. In Encyclopedia of Library and Information Systems.Vol. 69, Supplement 32, New York (2000)
19. Fogg, B. J.: How Do Users Evaluate The Credibility of Web Sites? In Proceedings of the 2003 conference on Designing for user experiences, ACM Press, NY, USA (2003)
20. Alexander, J., Tate, M.: Web Wisdom: How to Evaluate and Create Information Quality on the Web, Lawrence Eribaum Associates Inc, New Jersey, USA (1999)
21. Cassel. R.: Selection Criteria for Internet Resources, College and Research Library News, N. 56, pag. 92-93, (February 1995)
22. Stanford Web Credibility Guidelines web site, http://credibility.stanford.edu/guidelines
23. Roberts, T.: Online Collaborative Learning, Theory and Practice, Idea Group Pub, USA (2004)
24. Dondio, P., Barrett, S., Weber, S.: Calculating the Trustworthiness of a Wikipedia Article Using Dante Methodology, IADIS eSociety conference, Dublin (2006)
25. http://download.wikimedia.org, download site of the Wikipedia project