A Subjective Credibility Model for Participatory Media

Aaditeshwar Seth **Jie Zhang** **Robin Cohen**

School of Computer Science
University of Waterloo, ON, Canada

Abstract

We propose a method to determine the credibility of messages that are posted in participatory media (such as blogs), of use in recommender systems designed to provide users with messages that are considered to be the most credible to them. Our approach draws from theories developed in sociology, political science, and information science – this results in a method for evaluating the credibility of messages that is user-specific and sensitive to the social network in which the user resides. Our methodology rests on Bayesian learning, integrating new concepts of context and completeness of messages inspired by the strength of weak ties hypothesis from social network theory. We show that our credibility evaluation model can be used to significantly enhance the performance of collaborative filtering recommendation. Experimental validation is done using a dataset obtained from digg.com, a knowledge sharing website where users indicate their satisfaction with messages that are provided to them. Our results reinforce the value of using sociological insights in recommender system design.

Introduction

With the goal to provide personalization on the web, one topic of concern is how to assist users in processing vast collections of participatory media content that exist in this environment. A specific challenge is to design a personalized recommender system that will be able to propose messages of interest to users. For example, *citizen journalism* through participatory media content such as blogs and comments on news articles, has become a popular supplement to mass media (Gillmor 2006). It adds diversity of opinion to news topics, and provides an additional level of localization in news coverage, addressing issues that may have been skipped by national news agencies. However, these collections of participatory media content tend to be huge and dynamic; for example, on the order of 1.6 million blog posts are written each day (Sifry 2007). In an effort to help users to cope with this vast amount of news, personalized recommender systems that propose news articles of interest to users would be beneficial.

In this paper, we examine one particular concern in the processing of messages, the modeling of a message's credibility, and use this to construct a recommender system which provides to users those messages that are considered to be the most credible. We contend that credibility is an important component to judge the usefulness of participatory media content. This is because of the ease of publishing information on the Internet without any editorial checks by a formal agency: Anybody can publish "incorrect" information, or bad-mouth "correct" information. We depart from this conventional thinking on the polarity of information credibility, however, and instead develop a subjective credibility model suitable for the scenario of participatory media. Our model aims to capture the following principles:

- *D-1*: Different users may judge the credibilities of blogs differently according to their own contexts.
- *D-2*: Different users may have different propensities to accommodate views contrary to their own beliefs.
- *D-3*: The credibility of a blogger is topic specific; an expert in some area may not be an expert in another.
- *D-4*: A highly credible blogger can occasionally make mistakes and give inaccurate information. Analogously, useful blog-entries could be written by a blogger unknown so far.

We draw from research in media studies, information science, political science, and social networks to refine these design principles into specific criteria that can be used to judge the credibility of information. These criteria include, for example, the influence of public opinion, influence of close friends of people, and the extent to which different people may trust their own beliefs. We then use these criteria to build and learn a Bayesian network on a personalized basis for each user, to predict which messages the user may find to be credible. Our method makes extensive use of social network information to create the user model, and combines the link structure of social networks of users with information about authorship and ratings of messages by users. We test our method on a dataset obtained from a popular knowledge sharing website, digg.com. Experimental results show that our method outperforms other well-known methods such as Pagerank used to rank Internet web-pages in order of their importance (Brin and Page 2001), Eigentrust used in peer-to-peer (P2P) file-sharing systems to identify trustworthy peers that upload valid content (Kamvar, Scholsser, and Garcia-Molina 2003), and the beta-reputation system used in e-commerce to evaluate the trustworthiness of buyers and sellers (Whitby, Jøsang, and Indulska 2005).

Our method has important implications for the design of recommender systems for participatory media content. It serves to predict the probability of a user finding a message to be credible, and can hence be used as a pre- or post-filtering stage with existing recommendation algorithms. In this paper, we show that our method can be adapted to integrate closely with collaborative filtering (CF) (Adomavicius and Tuzhilin 2005); enhancing a CF algorithm with our credibility model can be shown to perform significantly better than the basic CF for a binary-classification of messages (i.e. $\{$*recommend*, *not recommend*$\}$) to a user. The validation is done using existing data about user ratings from the digg.com dataset.

In the sections that follow, we first use insights from sociology to determine the credibility judgement criteria used by people. We then describe the Bayesian network model in detail, and provide the evaluation of our method. We also show how our modeling technique can be adapted to improve the performance of collaborative filtering for recommendation of participatory media content. Finally, we present related work, a discussion, and future work.

Credibility Judgement Criteria

In this section, we build upon insights about credibility developed in different disciplines. We then use these insights to construct a Bayesian model for each user; the model parameters can be learned using $\pm$ve ratings given by a user to messages seen by her in the past, and can be used to predict whether the user will find a new unseen message to be credible.

Multi-dimensional construct Various researchers have proposed to model credibility as a multi-dimensional construct (Fogg and Tseng 1999; Sabater and Sierra 2001). Fogg and Tseng (1999) reason about credibility criteria used by people to judge the credibility of computerized devices and software, and identify the following different types of credibility:

- *Experienced*: This is based on *first-hand experience* of a user, and reflects her personal belief about a device or software.
- *Presumed*: This reflects personal biases of a user that give rise to *general assumptions* about certain categories of computing products; for example, presumptions based upon the company which developed the product, the cost of the product, the importance of the function performed by the product, etc.
- *Reputed*: This is based on *third-party reports* about different products.

A model with similar distinctions is developed in (Sabater and Sierra 2001) to evaluate the trustworthiness of agents in an e-commerce setting. Here, the authors distinguish *witness reputation* (i.e. general public opinion) from *direct reputation* (i.e. opinion from a user's own experience) and include as well *system reputation* (i.e. the reputation from the role of an agent, as buyer, seller or broker). We next consider relevant studies from sociology and political science for additional valuable insights.

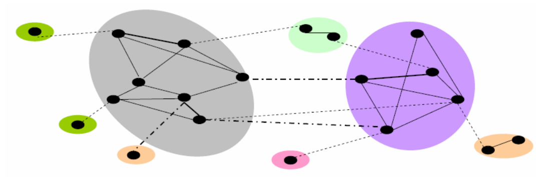

Figure 1: Strong and weak ties

Social networks People are embedded in real-world social networks of relationships as friends, acquaintances, family members, etc. The *strength-of-weak-ties* hypothesis in sociology (Granovetter 1973) states that such social networks consist of clusters of people with *strong* ties among members of each cluster, and *weak* ties linking people across clusters, shown in Fig. 1. Whereas strong ties are typically constituted of close friends, weak ties are constituted of remote acquaintances. The hypothesis claims that weak ties are useful for the diffusion of information and economic mobility, because they connect diverse people with each other. On the other hand, people strongly tied to each other in the same cluster may not be as diverse.

One among many studies based on the *strength-of-weak-ties* hypothesis, (Baybeck and Huckfeldt 2002) traces the changes in political opinion of people before and after the 1996 presidential elections in USA, observed with respect to the social networks of people. It is shown that weak ties (identified as geographically dispersed ties of acquaintances) are primarily responsible for the diffusion of divergent political opinion into localized clusters of people having strong ties between themselves. As indicated by the *strength-of-weak-ties* hypothesis, this reflects that local community clusters of people are often homogeneous in opinion, and these opinions may be different from those of people belonging to other clusters. Furthermore, people have different propensities to respect opinions different from those of their immediate local community members. This reflects that the personal characteristics of people also influence the extent to which they would be comfortable in deviating from the beliefs of their immediate local cluster. These observations provide two insights:

- Reputed credibility has at least two sub-types: *cluster credibility* based on the opinions of people in the same cluster or local community, and *public credibility* based on the general opinions of everybody.
- Users have different personal characteristics to weigh the importance of different types of credibilities.

The first insight suggests refining *reputed* credibility to also consider reports from those in the same cluster. The second insight is reinforced by studies in information science (Rieh 2002), which argue that users have different preferences for different types of credibilities discussed so far. Inspired by these studies, we develop and operationalize a multi-dimensional subjective credibility model for participatory media as described next.

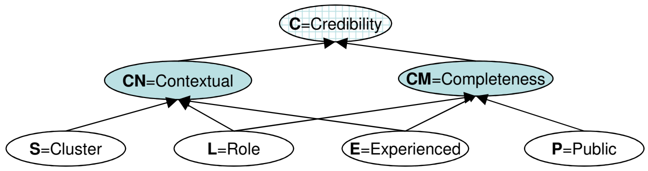

Figure 2: Credibility model

Bayesian User Model

Knowledge assumptions Suppose that we wish to predict whether a message m_k about a topic t and written by user u_j, will be considered credible by user u_i. We assume that we have the following prior knowledge:

- We consider a scenario where all older messages about topic t written in the past are labeled with the author of each message. In addition, a message may have also been assigned ratings by various recipient users, whenever users would have read the message, based on the credibility of the message for the recipient. The set of credibility ratings of any message are also assumed to be available [1].
- Users may declare a subset of other users as their "friends". We refer to an explicitly declared relationship between two users as a *link* between them, and assume to have knowledge of the social network graph formed by all users and the links between pairs of users.
- Users may also declare topics of interest to them. We use this information, and the social network graph, to derive the *topic specific social network graph* for topic t, as the induced subgraph of the overall social network graph consisting only of those users and edges between users who are interested in topic t.
- For each topic specific social network graph, community identification algorithms such as (Dongen 2000; Tantipathananandh, Berger-Wolf, and Kempe 2007) can identify dense clusters of users and links. We use the definition of *strong* and *weak* ties proposed by (Granovetter 1973), and refer to *strong* ties as links between users in the same cluster, and *weak* ties as links between users in different clusters. We use V_{it} to denote the local cluster of users strongly tied to user u_i with respect to topic t.

These assumptions are reasonable in contexts such as the website digg.com, which allows users to construct social networks by declaring some users as their friends. Information about message authorship and ratings given by users to messages is also available. We will show that we can use this knowledge to quantify different types of credibilities for each message with respect to each user. Then, based on ratings given by a particular user to older messages, we can use a Bayesian model to learn preferences of the user towards these different kinds of credibilities of messages. Finally, we can use this learned model to predict whether or not the new message m_k will be considered credible by user u_i.

Bayesian network The different types of credibilities that we choose to model are as follows:

- e_{ikt}= *experienced credibility*: Identical to the concept of experienced credibility discussed earlier, this is based only on ratings given by user u_i in the past, and denotes the credibility that u_i associates with the message m_k written by u_j.
- l_{ikt}= *role based credibility*: Similar to presumed credibility discussed earlier, this denotes the credibility that u_i associates with the message m_k written by users having the same role as that of u_j; for example, based on whether the messages' authors are students, or professors, or journalists, etc.
- s_{ikt} = *cluster credibility*: A sub-type of reputed credibility discussed earlier, this is based on the ratings given by other users in cluster V_{it}, that is, the cluster of user u_i. It denotes the credibility associated by the cluster or local community of u_i to the message m_k written by u_j.
- p_{kt} = *public credibility*: Another sub-type of reputed credibility, this is based on ratings by all the users, and reflects the general public opinion about the credibility for the message m_k written by u_j.

Each of these credibilities can be expressed as a real number $\in [0, 1]$, and we propose a Bayesian network to combine them into a single credibility score. The model is shown in Fig. 2. Our aim is to learn the distribution for $\mathrm{P}_{it}(\mathbf{C}|\mathbf{E,L,S,P})$ for each user and topic based on ratings given by various users to older messages; here, $\{\mathbf{E,L,S,P}\}$ are evidence variables for the four types of credibilities for a message, and **C** is a variable denoting the credibility that u_i associates with the message. Thus, for each topic t, a set of messages M about t will be used during the training phase with samples of (c_{ik}, e_{ik}, l_{ik}, s_{ik}, p_k) for different messages $m_k \in M$ to learn the topic specific credibility models for u_i. Assuming that a user's behavior with respect to preferences for different kinds of credibilities remains consistent over time, the learned model can now be used to predict c_{ix} for a new message m_x about topic t, that is, $\mathrm{P}_{it}(c_{ix}|e_{ix}, l_{ix}, s_{ix}, p_x)$.

Fig. 2 also shows two hidden variables as shaded ovals. The hidden variables help make the model more tractable to learn, and also capture an insight we developed in prior work (Seth 2007; Seth and Zhang 2008). We showed that a new message has two characteristics with respect to a recipient: it carries some contextual information for the recipient about the issue being discussed in the message, and some degree of completeness of information about the issue. Context and completeness are defined as follows:

- *Context* relates to the ease of understanding of the message, based on how well the message content explains the relationship of the message to its recipient. *Simplification* of the meaning of the message (Bryant and Zillman 2002), can be considered as an outcome of the amount of

[1] We also assume that we are beyond the cold-start stage so that the set of older messages have all received some ratings, and all users have provided at least some ratings.

context in the message. That is, messages that are more contextual for users, will be more simple for them to understand.

- *Completeness* denotes the depth and breadth of topics covered in the message. The *scope* of the message, or the *opinion diversity* expressed in the message (Bryant and Zillman 2002), can be considered as outcomes of the degree of *completeness* of the message. That is, messages that are more complete will carry more diverse opinions or more mention of relationships with other issues.

Examples about the concepts of context and completeness are given in the appendix. What is relevant to this paper is simply to understand that the Bayesian model is a hierarchical model: For each message, the model first estimates the credibilities of the contextual and complete information carried by the message, and then uses these two credibilities to generate the final estimate. We reason that cluster credibility will only influence contextual credibility, while public credibility will only influence completeness credibility. This is because general public opinion is by definition averaged over different contexts, and hence it will only add noise to any context specific credibility. Similarly, cluster credibility will double count the opinion of a specific cluster when judging the degree of completeness or diversity in a message. Other types of credibilities, experienced and role based, will influence both contextual and completeness credibility since they are based on the personal beliefs of the user.

Meeting the design principles

Our modeling method is able to satisfy three out of the four design principles listed in the "Introduction" section. (*D-1*) The model takes into account personal and contextual opinions of people that may influence their credibility judgements. (*D-2*) The model is learned in a personalized manner for each user, and allows to accommodate varying degrees of propensities of users to respect opinions of other users. (*D-3*) Different model instances are learned for different topics, making credibility judgements topic specific. (*D-4*) We will show in the next section that the fourth principle of allowing mistakes by credible users and useful messages by non-credible users can also be modeled in this framework.

Credibility Computation

In this section, we describe how the different types of credibilities can be computed based on social network information, ratings given by users to messages, and authorship information. We first list the axioms that are the basis for our formulation to quantify the various types of credibilities, and then give the actual computation process.

Axioms to calculate credibility

We use the information captured in the following relationships:

- *A-1*: A message is credible if it is rated highly by credible users.
- *A-2*: A user is credible if messages written by her are rated highly by other credible users.
- *A-3*: A user is also credible if ratings given by her are credible, that is, she gives high ratings to messages that appear to be credible to credible users, and low ratings to messages that appear to be non-credible.
- *A-4*: A user is also credible if she is linked to by other credible users in the social network.

There is clearly a recursive relationship between these axioms. We solve the recursion using fixed-point Eigenvector computations, as described next.

Calculation of evidence variables

We henceforth assume that we are operating within some topic t, and drop the subscript for simplicity. As stated in the knowledge assumptions earlier, we start with the following information that will be a part of our training set.

- **A[k,n]**: A matrix for k messages and n users, where $a_{ij} \in \{0, 1\}$ indicates whether message m_i was written by u_j
- **R[k,n]**: A ratings matrix for k messages and n users, where $r_{ij} \in \{0, 1\}$ [2] indicates the rating given to message m_i by user u_j
- **N[n,n]**: A social network matrix where $n_{ij} \in \{0, 1\}$ indicates the presence or absence of a link from user u_i to user u_j. We also assume that the clustering algorithm can identify clusters of strong ties among users, connected to other clusters through weak ties.

Our goal is to find a method to compute the evidence variables for the Bayesian model using the axioms given above. The evidence variables can be expressed as the matrices **E[n,k]**, **L[n,k]**, **S[n,k]**, and **P[k]**, containing the credibility values for messages. Here, p_k is the public credibility for message m_k authored by user u_j. e_{ij} and l_{ij} are the experienced and role based credibilities respectively for message m_k according to the self-beliefs of user u_i. Similarly, s_{ij} is the cluster credibility for message m_k according to the beliefs of the users in u_i's cluster V_i. Once these evidence variables are computed for older messages, they are used to learn the Bayesian model for each user. Subsequently, for a new message, the learned model for a user is used to predict the credibility of the new message for the user.

We begin with computation of the evidence variable matrix for public credibility **P**; we will explain later how other credibilities can be computed in a similar fashion.

1. Let **P′[n]** be a matrix containing the public credibilities of users, and consider the credibility of a message as the mean of the ratings for the message, weighted by the credibility of the raters (*A-1*):

   $$p_k = \sum_i r_{ki}.p'_i / |r_{ki} > 0|$$

   This is the same as writing **P=R**$_r$**.P′**, where **R**$_r$ is the row-stochastic form of **R**, ie. the sum of elements of each row = 1.
2. The credibility of users is calculated as follows:

[2] We assume in this paper that the ratings are binary. However, our method can be easily generalized to real-valued ratings as well.

2a. Consider the credibility of a user as the mean of the credibilities of the messages written by her (*A-2*):

$$p'_i = \sum_k p_k / |p_k|$$

This is the same as writing $\mathbf{P'}$=$\mathbf{A}_c^T$.$\mathbf{P}$, where $\mathbf{A}_c$ is the column-stochastic form of $\mathbf{A}$; and $\mathbf{A}_c^T$ is the transpose of $\mathbf{A}_c$.

2b. The above formulation indicates a fixed point computation:

$$\mathbf{P'}\text{=}\mathbf{A}_c^T.\mathbf{R}_r.\mathbf{P'} \quad (1)$$

Thus, $\mathbf{P'}$ can be computed as the dominant Eigenvector of $\mathbf{A}_c^T.\mathbf{R}_r$. This formulation models the first two axioms, but not yet the ratings-based credibility (*A-3*) and social network structure of the users (*A-4*). This is done as explained next.

2c. Perform a fixed-point computation to infer the credibilities **G[n]** acquired by users from the social network (*A-4*):

$$\mathbf{G}\text{=}(\beta.\mathbf{N}_r^T + (1\text{-}\beta).\mathbf{Z}_c.\mathbf{1}^T).\mathbf{G} \quad (2)$$

Here, $\beta \in (0, 1)$ denotes a weighting factor to combine the social network matrix **N** with the matrix **Z** that carries information about ratings given to messages by users. We generate **Z** by computing z_i as the mean similarity in credibility ratings of user u_i with all other users. The ratings similarity between a pair of users is computed as the Jacquard's coefficient of common ratings between the users. Thus, z_i will be high for users who give credible ratings, that is, their ratings agree with the ratings of other users (*A-3*). In this way, combining the social-network matrix with ratings-based credibility helps to model the two remaining axioms as well. Note that $\mathbf{Z}_c$**[n]** is a column stochastic matrix and **1[n]** is a unit column matrix; augmenting **N** with $\mathbf{Z}_c.\mathbf{1}^T$ provides an additional benefit of converting **N** into an irreducible matrix so that its Eigenvector can be computed [3].

2d. The ratings and social network based scores are then combined together as:

$$\mathbf{P'}\text{=}(\alpha.\mathbf{A}_c^T.\mathbf{R}_r + (1\text{-}\alpha).\mathbf{G}_c.\mathbf{1}^T).\mathbf{P'} \quad (3)$$

Here again **1** is a unit column matrix, and $\alpha \in (0, 1)$ is a weighting factor. The matrix $\mathbf{P'}$ can now be computed as the dominant Eigenvector using the power method.

3. Once $\mathbf{P'}$ is obtained, **P** is calculated in a straightforward manner as $\mathbf{P}$=$\mathbf{R}_r.\mathbf{P'}$.

Note that the above method is only one way of combining the different pieces of information we have. Our objective in presenting this method is to show that information about social networks, ratings, and authorship can be combined together and to then examine the performance of this method compared to competing approaches.

The above process is to compute the public credibilities **P[k]** of messages. The processes to compute cluster **S[n,k]**, experienced **E[n,k]**, and role based **L[n,k]** credibilities are identical, except that different cluster credibilities are calculated with respect to each cluster in the social network, and different experienced and role based credibilities are calculated with respect to each user. This is why cluster and experienced credibility matrices are 2-dimensional, while the public credibility is only 1-dimensional. For example, considering a message m_3 and a recipient user u_1, **P[3]** is the public credibility of message m_3; **E[1,3]** is the experienced credibility of message m_3 according to the self-belief of recipient u_1; **L[1,3]** is the role based credibility of message m_3 also according to the self-belief of recipient u_1; and **S[1,3]** is the cluster credibility of message m_3 according to the beliefs of users in cluster V_1 of recipient u_1. The processing steps for computing these quantities are outlined in Algorithm-1; a description is below.

- The cluster credibilities **S[n,k]** are computed in the same manner as the public credibilities, but after modifying the ratings matrix **R** to contain only the ratings of members of the same cluster. Thus, the above process is repeated for each cluster, modifying **R** in every case. For each users u_i belonging to cluster V_i, s_{ik} is then equal to the cluster credibility value for message m_k with respect to u_i.

  The matrix **Z** in the computation on the social network matrix is also modified. When computing the cluster credibilities for cluster V_i, element z_j of **Z** is calculated as the mean similarity of user u_j with users in cluster V_i. Thus, z_j will be high for users who are regarded credible by members of cluster V_i because their ratings agree with the ratings of the cluster members.
- The experienced credibilities **E[n,k]** are computed in the same manner as well, but this time for each user by modifying the ratings matrix **R** to contain only the ratings given by the user. The matrix **Z** is also modified each time by considering z_j as the similarity between users u_i and u_j, when calculating the experienced credibilities for u_i.
- Role based credibility is computed as the mean experienced credibilities of users having the same role. However, we do not use role based credibility in our evaluation because sufficient user profile information was not available in the digg dataset used by us. Henceforth, we ignore **L[n,k]** in our computations.

Model learning

Once various types of credibilities for messages are calculated with respect to different users, this data is used to learn the Bayesian model for each user and topic of interest using the Expectation-Maximization (EM) algorithm (Russel and Norvig 2003). Model parameters are learned to predict for user u_i interested in topic t, the probability $\mathrm{P}_{it}(c_{ix}|e_{ix}, s_{ix}, p_x)$ that u_i will find a new message m_x to be credible.

[3] This step is similar to the Pagerank or HITS computations for the importance of Internet web pages (Brin and Page 2001; Kleinberg 1998). The matrix **N** can be considered as the link matrix of web-pages, and the matrix **Z** as the pagerank personalization matrix. The output matrix **G** then essentially ranks the web-pages in order of their importance, after taking personalization into account.

Algorithm 1: Training set preparation

```
Input: A[k,n], R[k,n], N[n,n]
Output: P[k], E[n,k], S[n,k], P'[k], E'[n,n], S'[n,n]

1. Compute similarity matrix Y[n,n]
forall i ∈ 1..n, j ∈ 1..n, i ≠ j do
    forall m ∈ 1..k do
        if R[m,i] = R[m,j] then
            Y[i,j] ← Y[i,j]+1/k

2. Compute public credibilities P[k], P'[n]
Z[n] ← 0
forall i ∈ 1..n do
    forall j ∈ 1..n do
        Z[i] ← Z[i] + Y[j,i]
Solve for G[n]: G=(β.N_r^T+(1-β).Z_c.1^T).G
Solve for P'[n]: P'=(α.A_c^T.R_r + (1-α).G_c.1^T).P'
P ← R_r.P'

3. Compute cluster credibilities S[n,k], S'[n,n]
forall Cluster V_c ∈ clusters in social network do
    Z[n] ← 0
    G[n] ← 0, P[n] ← 0, P'[n] ← 0, R[k,n] ← 0
    forall j ∈ users in V_c do
        forall i ∈ 1..n do
            Z[i] ← Z[i] + Y[j,i]
        forall m ∈ 1..k do
            R[m,j] ← R[m,j]
    Solve for G[n]: G=(β.N_r^T+(1-β).Z_c.1^T).G
    Solve for P'[n]: P'=(α.A_c^T.R_r+(1-α).G_c.1^T).P'
    P=R_r.P'
    forall j ∈ users in V_c do
        forall m ∈ 1..k, u ∈ 1..n do
            S'[j,u] ← P'[u]; S[j,m] ← P[m]

4. Compute experienced credibilities E[n,k], E'[n,n]
forall User i ∈ 1..n do
    Z[n] ← 0
    G[n] ← 0, P[n] ← 0, P'[n] ← 0, R[k,n] ← 0
    forall j ∈ 1..n do
        Z[j] ← Y[j,i]
    forall m ∈ 1..k do
        R[m,i] ← R[m,i]
    Solve for G[n]: G=(β.N_r^T+(1-β).Z_c.1^T).G
    Solve for P'[n]: P'=(α.A_c^T.R_r+(1-α).G_c.1^T).P'
    P ← R_r.P'
    forall m ∈ 1..k, u ∈ 1..n do
        E'[i,u] ← P'[u]; E[i,m] ← P[m]
```

Algorithm 2: Inference phase (ratings based)

```
Input: User i, Cluster V_i of user i, Message m;
       Ratings R[n,m] given by other users to m;
       Learned model for user i
Output: P(user i will find m to be credible | R[k])

p_m ← mean(R[j,m].P'[j])_{j∈1..n}
s_im ← mean(R[j,m].S'[i,j])_{j∈1..n}
e_im ← mean(R[j,m].E'[i,j])_{j∈1..n}
P(c_im | p_im, s_im, e_im) ← MCMC on learned model for i
```

Inference

Now, for a new message m_x, the evidence variables are calculated with respect to a recipient user u_i in one of two ways as described next, and the learned model is used to produce a probabilistic prediction of whether u_i would find m_x to be credible.

- *Authorship*: The four types of credibilities of the message are considered to be the same as the corresponding four types of credibilities of its author with respect to u_i.
- *Ratings*: The cluster and public credibilities are calculated as the weighted mean of ratings for the message given by other users and the credibilities of these users with respect to u_i. The experienced and role based credibilities are the same as the corresponding credibilities of the message author with respect to u_i.

As we will show in the evaluation, the ratings method performs better than the authorship method. This also meets the fourth design principle (*D-4*) listed in the "Introduction" section. Since credibility is evaluated through ratings given to the message by various users, it allows new users to popularize useful messages written by them because their own credibility does not play a role in the computations. It also allows credible users to make mistakes because the credibility of the author is not taken into account.

Given the evidence variables for the new message, and the learned Bayesian model, the probability of u_i finding the message to be credible is computed using standard belief propagation methods such as Markov-Chain-Monte-Carlo (MCMC) (Russel and Norvig 2003). The outline is given in Algorithm-2.

Evaluation

We evaluate our method over a dataset of ratings by real users obtained from a popular knowledge sharing website, digg.com (Lerman 2007). The website allows users to submit links to news articles or blogs, which are called *stories* in the terminology used by the website. Other users can vote for these stories; this is known as *digging* the stories. Stories that are *dugg* by a large number of users are promoted to the front-page of the website. In addition, users are allowed to link to other users in the social network. Thus, the dataset provides us with all the information we need:

- Social network of users: We use this information to construct the social network link matrix between users **N[n,n]**. The social network is clustered using MCL, a flow-stochastic graph clustering algorithm (Dongen 2000), to produce classifications of ties as strong or weak (Seth 2007). The cluster of users strongly connected to user u_i is referred to as V_i.
- Stories submitted by various users: This is used to construct the authorship matrix **A[k,n]**. Since all the stories in the dataset were related to technology, we consider them as belonging to a single topic.
- Stories dugg by various users: We use this information to construct the ratings matrix **R[k,n]**. We consider a vote of 1 as an evidence for credibility of the story, and a vote of 0 as an evidence of non-credibility.

Although the dataset is quite large with over 200 stories, we are able to use only 85 stories which have a sufficiently large number of ratings by a common set of users. This is because we require the same users to rate many stories so that we have enough data to construct training and test datasets for these users. Eventually, we assemble a dataset of 85 stories with ratings by 27 users. We do not include users who rate more than 65 stories as all credible or all non-credible, because a good predictor for such users would trivially be to always return 1 or 0, and besides, such user behavior may amount to attacks on the system which we consider as future work. A few assumptions we make about the validity of the dataset for our experiments are as follows:

- The original submission of a story to Digg may not have been made by the author of the story. However, we regard the submitting user as the message author because it distinguishes this user from other users who only provide ratings to the messages.
- The ratings provided on the Digg website may not reflect credibility ratings, but rather usefulness ratings given to messages by users. We however consider them to be equivalent to credibility because of the smaller dataset size we use. We argue that since the users in the dataset vote for at least 20 stories out of 85 (25% of the total number of stories), they are likely to be interested in the topic and all the stories; therefore, the only reason for their not voting for a story would be its credibility.

We use an open-source package, OpenBayes, to program the Bayesian network. We simplify the model by discretizing the evidence variables **E,S,P** into 3 states, and a binary classification for the hidden variables **N, M**, and the credibility variable **C**. The discretization of the evidence variables into 3 states is performed by observing the Cumulative Distribution Frequency (CDF) and Complementary Cumulative Distribution Frequency (CCDF) of each variable with respect to the credibility rating of users. The lower cutoff is chosen such that the product of the CDF for rating=0 and CCDF for rating=1 is maximum, and the upper cutoff is chosen such that the CCDF for rating=0 and CDF for rating=1 is maximum. This gives a high discrimination ability to the classifier because the cutoffs are selected to maximize the pair-wise correlation of each evidence variable with the credibility rating.

Choice of parameters

The first set of experiments shown here find good values of α (eqn. 3) and β (eqn. 2), and compare ratings with authorship based evidence variable computation (the "Inference" section). We evaluate the performance of the model for each user by dividing the 85 stories into a training set of 67 stories and a test set of 17 stories (80% and 20% of the dataset respectively). We then repeat the process 20 times with different random selections of stories to get confidence bounds for the cross validation. For each evaluation, we use two kinds of performance metrics (Davis and Goadrich 2006):

- *Matthew's correlation coefficient (MCC)*:

$$\text{MCC} = \frac{(t_p.t_n - f_p.f_n)}{\sqrt{(t_p+f_p)(t_p+f_n)(t_n+f_p)(t_n+f_n)}}$$

  Here, f_p = false positives, t_p = true positives, f_n = false negatives, t_n = true negatives. The MCC is a convenient measure because it gives a single metric for the quality of binary classifications.

- *TPR Vs FPR*: This plots on an XY-scale the true positive rate (TPR) with the false positive rate (FPR) of a binary classification. Maximum accuracy implies TPR=1.0 and FPR=0.0, while TPR=FPR is the random baseline. Therefore, points above the random baseline are considered to be good.

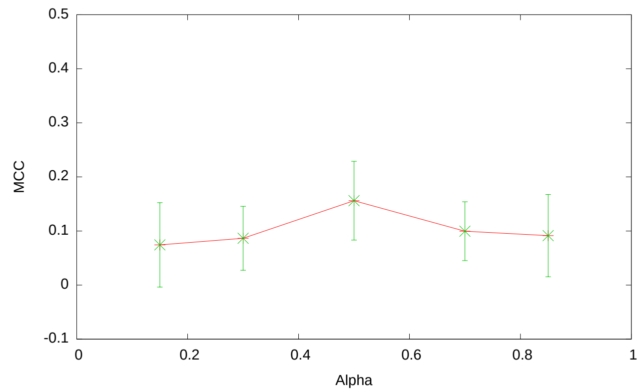

Figure 3: Performance with different parameters

Fig. 3 shows the mean MCC across all users for different values of α (eqn. 3) to combine the ratings and social network matrices. The best performance happens at $\alpha = 0.5$, conveying our message that all of authorship, ratings, and social networks provide valuable credibility information. All the experiments are done using ratings-based inference with $\beta = 0.85$ (eqn. 2). Larger or smaller values of β both give poorer results.

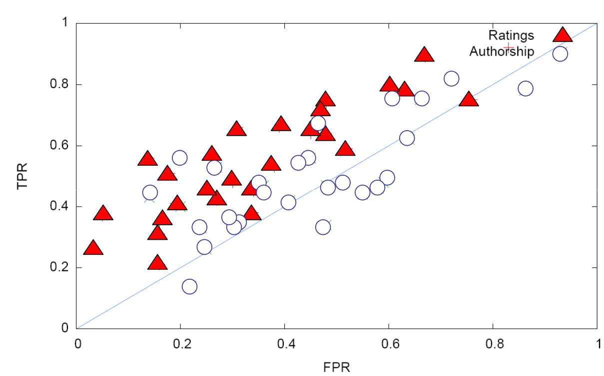

Figure 4: Performance of Bayesian credibility model

Inference methods

Fig. 4 shows the TPR-FPR plot for ratings and authorship based evidence variable computation when $\alpha = 0.5$ and $\beta = 0.85$. As can be seen visually, the ratings-based method performs better than the authorship-based method. The former gives MCC = 0.156 (σ=0.073), while the latter gives MCC = 0.116 (σ=0.068). However, the authorship performance is still successful for a majority, which is encouraging. This indicates that authorship information may be used to solve the problem of cold-start for new messages that have not acquired a sufficient number of ratings. Similarly, ratings may be used to solve cold-start for new users who have not acquired sufficient credibility.

We notice that the classifier performs very well for some users, but close to random for some other users. We therefore investigate various characteristics that may prove useful to determine for which users our method may work well and when it may not.

- We compute the variance of cluster and experienced credibility scores for different users. We then compare the variances by good performing users ($\frac{TPR}{FPR} > 1.5$) with the variances by the remaining users. We find that for both cluster and experienced credibilities, the variances by good performing users are more than twice the variances by poorly performing users.

  This shows the more the discrimination produced in the cluster and experienced credibility scores by a user, the better the performance of the user, because greater discrimination ability implies higher entropy in the information theoretic sense.
- We find that on an average, 85% of users in the same cluster are likely to be all good performing or all poorly performing. This is an interesting result because we also find that users in the same cluster are four times more similar to each other in their credibility ratings than to users in other clusters. Although the similarity of ratings explains why the majority of users also perform similarly, an open question is whether the performance of a user goes up or down because of the cluster in which she is a member, or simply because the ratings given by her are too inconsistent to be captured by the Bayesian model.

As part of future work, we will try to identify more features to classify ratings, authorship, and social network matrices in terms of their characteristics to yield good or bad performance for users.

Comparison with other methods

We next compare our method with other well known methods for trust and reputation computation meant for different applications. All these methods perform very close to random, even with personalization. We believe this to be due to a fundamental drawback of these methods: they try to form an objective assessment of credibility for users and messages, which is not appropriate for participatory media content.

- An Eigenvector computation on $\mathbf{A}_c^T.\mathbf{R}_r$ by leaving out the social network part (eqn. 1), is identical to the Eigentrust algorithm (Kamvar, Scholsser, and Garcia-Molina 2003). The best choice of parameters could only give a performance of MCC = -0.015 (σ = 0.062). Eigentrust has primarily been shown to work in P2P file sharing scenarios to detect malicious users that inject viruses or corrupted data into the network. However, the P2P context requires an objective assessment of the trustworthiness of a user, and does not allow for subjective differences, as desired for participatory media.
- An Eigenvector computation on the social network matrix (eqn. 2), personalized for each user, is identical to the Pagerank algorithm used to rank Internet web pages (Brin and Page 2001). However, this too performs poorly with an MCC = 0.007 (σ = 0.017). This suggests that users are influenced not only by their own experiences, but also by the judgement of other users in their cluster, and by public opinion. Methods ignoring these factors may not perform well.
- The beta-reputation system (Whitby, Jøsang, and Indulska 2005) is used in e-commerce environments to detect good or bad buying and selling agents. It estimates the credibility of agents in an objective manner using a probabilistic model based on the beta probability density function. Only the public opinion is considered; ratings are filtered out if they are not in the majority amongst other ratings. It too does not perform well in the context of participatory media, giving an MCC = 0.064 (σ = 0.062).

Our conclusion is that approaches which subjectively model credibility, allowing users to be influenced in different ways by different sources, perform better than objective modeling approaches.

Use in Recommender Systems

As mentioned earlier, our method for credibility computation can be used in two ways to improve recommender systems: (i) Since our method serves to predict the probability of a user finding a message to be credible or non-credible, it can be used as a pre- or post-filtering stage with existing recommendation algorithms. (ii) As shown in this section, our proposed model can be adapted to integrate closely with recommendation algorithms; we show how to do this with collaborative filtering (CF) (Adomavicius and Tuzhilin 2005).

A basic CF algorithm works in two steps. First, similarity coefficients are computed between all pairs of users, based on the similarity of message ratings given by each pair. Second, to make a decision whether or not to recommend a new message to a user, the mean of the message ratings given by other similar users is computed, weighted on the coefficients of similarity to these users. If the mean is greater than a threshold, the message is recommended; else it is rejected.

The drawback of the CF method is that it only learns the average user behavior. However, as we have argued, user behavior can be different in different circumstances. We therefore develop an adaptation of our method. Rather than computing a single similarity coefficient between each pair of users, we compute four similarity coefficients based upon whether messages are believed to be highly contextual by both users, or highly complete by both users, or contextual by the first user and complete by the second user, or vice versa. Essentially, we break down the average user behavior into four components based upon the context and completeness of messages to users, as follows:

1. For each user, we run the EM algorithm on the training set to learn the model.
2. We use the learned model to infer the probabilities of the hidden variables of context and completeness for each story in the training set: $P_i(\mathbf{CN}|\mathbf{E,S,P,C})$ and $P_i(\mathbf{CM}|\mathbf{E,S,P,C})$ shown in Fig. 2. That is, for each story m_j, we infer P(cn_{ji}=0,1$|e_{ji}, s_{ji}, p_{ji}, c_{ji}$) and P($cm_{ji}$=0,1$|e_{ji}, s_{ji}, p_{ji}, c_{ji}$).
3. We then discretize the probabilities for **CN** and **CM** in same way as we did earlier, by finding cutoffs that maximized the product of the CDF for c_{ji}=0 and CCDF for c_{ji}=1. This gives us samples of ($c_{ji} \in \{0,1\}$, $cn_{ji} \in \{0,1\}$, $cm_{ji} \in \{0,1\}$), that is, which stories appear contextual or complete to a user, and the rating given by the user to these stories.
4. For every pair of users, their samples are then compared to produce four similarity coefficients on how similar the users are in their contextual opinion, completeness opinion, and cross opinions between messages that appear contextual to one user and complete to the other, or vice versa.
5. Finally, when evaluating the decision to recommend a test message to a user, the mean of the message ratings is computed over all the four coefficients of similarity, rather than over a single coefficient as in the basic CF algorithm.

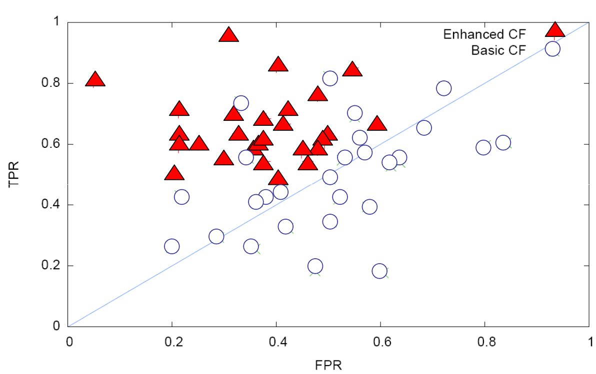

Figure 5: Enhancement of collaborative filtering

Fig. 5 shows the performance of the basic CF scheme and our enhanced version. The basic scheme performs worse than random for many users, but when enhanced with breaking up the average user behavior into contextual and completeness components, the performance improves considerably. The mean MCC for the basic scheme is 0.017 (σ = 0.086), and for the enhanced scheme is 0.278 (σ = 0.077), a sixteen-fold improvement. We consider this to be a huge improvement over the existing methodologies for trust, reputation, and recommendation algorithms, especially to build applications related to participatory media. Our results reinforce the value of using sociological insights in recommender system design.

Related Work

Our credibility model allows the credibility of messages to be evaluated, makes use of information about social network of users and ratings of messages, and learns for each user a Bayesian model to combine different types of credibilities. In this section, we provide a brief summary of some existing research and point out how they are different from our approach.

Various researchers in the P2P community have focused on Eigenvector based methods to compute the reputation of peers in sharing reliable content (Kamvar, Scholsser, and Garcia-Molina 2003). The ratio of successful to unsuccessful content exchanges is computed for each pair of peers who have interacted in the past, and these values are propagated in a distributed manner assuming a transitive trust relationship between peers. However, this is used to only compute the peer reputations (i.e. evaluating users) and not the reliability of content that is shared by the peers. A similar approach of Eigenvector propagation was also used in (Pujol, Sanguesa, and Delgado 2002) to compute reputation scores in a blog network, but the reputation of individual blog-entries was not computed. In our approach, we make use of message ratings and compute the credibility of each message.

For P2P networks, a method was proposed in (Walsh and Sirer 2006) where the object reputation is directly calculated to determine whether or not to accept a file being shared on a peer-to-peer network. Transaction history is used to assign edge weights between pairs of peers based on the similarity of ratings given by them to common objects rated in the past. Instead of using Eigenvector propagation to compute an absolute reputation score, a small set of shortest paths is found for each pair of peers, and the relative trust between the peers is computed as the mean of the product of edge weights along the paths. In our approach, we offer a richer multi-dimensional representation, integrating concepts of cluster, experienced and public credibility.

Researchers in the AI community have examined trust models for multi-agent based electronic marketplaces. For example, (Zhang and Cohen 2006) and (Whitby, Jøsang, and Indulska 2005) offer systems that determine the trustworthiness of an agent (i.e. a user). In addition, the use of an extensive trust model is promoted in (Sabater and Sierra 2001), to include features of contextual, role-based and experienced trust. We also have a multi-dimensional model, but we place great emphasis on representing and making use of the social network of a user, in order to learn a user-specific credibility rating for messages.

Discussions and Future Work

Confidence bounds: Methods for combining trust and confidence have been proposed by researchers such as (Kuter and Golbeck 2007) and (Huynh, Jennings, and Shadbolt 2004). For future work, it may be valuable to explore how to incorporate the concept of confidence into our model, for example as a way of placing bounds on the statistical hypotheses that are formed at each step of our algorithm.

Dataset size: Given the limited size of our dataset, we have not been able to form significant insights about the size of the training data required for our model to perform well. We will work with larger datasets in the future to understand this aspect in a better way.

Model extensions: We view our proposed method more as an extensible framework that can be extended to incorporate new insights or information. For example, we could explore the concept of *expert credibility* in the future, for which we would repeat the Eigenvector computations by considering ratings only by a specific set of users categorized as expert users by expert identification algorithms (Kolari et al. 2007). Another piece of information that is typically available in participatory media content, although it is not available in the digg dataset that we used, is the *message link matrix* based on hyperlinks between messages. An axiom that credible messages link to other credible messages can be modeled through pagerank or HITS, and included as an additional weighting factor in the Eigenvector computations. Alternatively, the polarity between links can be derived by sentiment analysis of the anchor text (Kale et al. 2007), and distrust propagation methods can be used to produce credibility scores based on the message link matrix (Guha et al. 2004).

Robustness to attacks: It would be desirable to have our model be robust in the face of attacks by malicious users. This may include scenarios where attackers could add noise to the ratings matrix by giving random ratings to various messages, or attackers could pollute the social network matrix by inviting unsuspecting users to link to them as friends, or even more sophisticated scenarios where attackers could collude with each other. In future work, we would like to examine the robustness of our model against such types of attacks. We also believe that attack analysis could give important insights about the implicit interactions between various pieces of information that are modeled together; such insights are likely to help improve performance.

Optimized computation: The proposed credibility model may be computationally intensive. However, Eigenvector optimization schemes are available that can decompose a large matrix into smaller matrices, and then combine the components together in an approximate fashion (Kamvar et al. 2003). We will experiment with such schemes in future work.

Recommender systems: In this paper, we showed how our model can be applied to collaborative filtering. We plan to apply the model to other recommendation algorithms as well, such as a model based algorithm we developed in prior work (Seth and Zhang 2008).

Conclusions

In this paper, we made use of insights from sociology, political and information science, and HCI, to propose a subjective credibility model for participatory media content. We formulated the model as a Bayesian network that can be learned in a personalized manner for each user, making use of information about the social network of users and ratings given by the users. We showed that our method works better than existing methods on trust and reputation computation. In addition, an adaptation of our method to recommendation algorithms such as collaborative filtering (CF) improves the performance of CF. This encourages the use of sociological insights in recommender system research.

Acknowledgements

We would like to express our sincerest thanks to Prof. S. Keshav for invaluable discussions about the nature of credibility, and to Prof. Kristina Lerman for providing us with the Digg dataset used in her experiments.

References

Adomavicius, G., and Tuzhilin, A. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. Knowledge and Data Engineering* 17(6).

Baybeck, B., and Huckfeldt, R. 2002. Urban contexts, spatially dispersed networks, and the diffusion of political information. *Political Geography* 21.

Brin, S., and Page, L. 2001. The pagerank citation ranking: Bringing order to the web. Technical Report http://dbpubs.stanford.edu:8090/pub/1999-66, Technical Report.

Bryant, J., and Zillman, D. 2002. *Media Effects: Advances in Theory and Research.* Lawrence Erlbaum Associates.

Davis, J., and Goadrich, M. 2006. The relationship between precision-recall and roc curves. In *Proceedings of ICML,*.

Dongen, S. 2000. *MCL: A Cluster Algorithm for Graphs*. PhD Thesis, University of Utrecht.

Fogg, B., and Tseng, H. 1999. The elements of computer credibility. In *Proceedings of SIGCHI*.

Gillmor, D. 2006. We the media: Grassroots journalism by the people, for the people. *O'Reilly Media*.

Granovetter, M. 1973. The strength of weak ties. *American Journal of Sociology* 78(6).

Guha, R.; Kumar, R.; Raghavan, P.; and Tomkins, A. 2004. Propagation of trust and distrust. In *Proceedings of WWW*.

Huynh, T. D.; Jennings, N. R.; and Shadbolt, N. 2004. FIRE: An integrated trust and reputation model for open multi-agent systems. In *Proceedings of 16th European Conference on Artificial Intelligence*, 18–22.

Kale, A.; Karandikar, A.; Kolari, P.; Java, A.; Joshi, A.; and Finin, T. 2007. Modeling trust and influence in the blogosphere using link polarity. *ICWSM*.

Kamvar, S.; Haveliwala, T.; Manning, C.; and Golub, G. 2003. Exploiting the block structure of the web for computing pagerank. Technical Report http://www-nlp.stanford.edu/pubs/blockrank.pdf, Tech Report.

Kamvar, S.; Scholsser, M.; and Garcia-Molina, H. 2003. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of WWW*.

Kleinberg, J. 1998. Authoritative sources in a hyperlinked environment. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms,*.

Kolari, P.; Finin, T.; Lyons, K.; Yesha, Y.; Yesha, Y.; Perelgut, S.; and Hawkins, J. 2007. On the structure, properties, and utility of internal corporate blogs. In *Proceedings of ICWSM*.

Kuter, U., and Golbeck, J. 2007. Sunny: A new algorithm for trust inference in social networks using probabilistic confidence models. In *Proceedings of AAAI*.

Lerman, K. 2007. Social information processing in news aggregation. *IEEE Internet Computing* 11(6).

Pujol, J. M.; Sanguesa, R.; and Delgado, J. 2002. Extracting reputation in multi agent systems by means of social network topology. In *Proceedings of AAMAS,*.

Rieh, S. 2002. Judgement of information quality and cognitive authority on the web. *Information Science and Technology* 53(2).

Russel, S., and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*. Pearson Education.

Sabater, J., and Sierra, C. 2001. Regret: A reputation model for gregarious societies. In *Proceedings of the Fifth International Conference on Autonomous Agents Workshop on Deception, Fraud and Trust in Agent Societies*, 61–69.

Seth, A., and Zhang, J. 2008. A social network based approach to personalized recommendation of participatory media content. In *Proceedings of ICWSM*.

Seth, A. 2007. Understanding participatory media using social networks. Technical Report CS-2007-47, University of Waterloo.

Sifry, D. 2007. The state of the live web. *http://www.sifry.com/alerts/archives/000493.html*.

Tantipathananandh, C.; Berger-Wolf, T.; and Kempe, D. 2007. A framework for community identification in dynamic social networks. In *Proceedings of SIGKDD*.

Walsh, K., and Sirer, E. G. 2006. Experience with an object reputation system for peer-to-peer filesharing. In *Proceedings of USENIX NSDI,*.

Whitby, A.; Jøsang, A.; and Indulska, J. 2005. Filtering out unfair ratings in bayesian reputation systems. *The Icfain Journal of Management Research* 48–64.

Zhang, J., and Cohen, R. 2006. A personalized approach to address unfair ratings in multiagent reputation systems. In *Proceedings of AAMAS Workshop on Trust in Agent Societies*.

Appendix: Clarifying context and completeness

Participatory messages such as blog entries and online discussions are not static: they *evolve* with participation from users when users write comments, or generate trackbacks to blog entries. These comments affect the usefulness of messages in different ways. Consider the following examples from two popular news websites.

Example 1: BBC News: An article was published on November 4^{th} 2007 about the Emergency declared in Pakistan. The article described some aspects of the event, such as President Musharraf's justification of his decision, condemnation by other political leaders of the country, and reactions of the judiciary [4]. Following are two comments on the same article.

- "I recently graduated in electrical engineering from Comsats Islamabad and got a job after a long struggle in one of the telecom companies here in Islamabad. I am hired on the basis that they are starting a new project in NWFP and FATA areas. After this emergency declaration company is now thinking to cancel the project in that area for which I was hired for, as NWFP and FATA areas are prime hiding places for Taliban... Now my job is in jeopardy and don't know what my future holds for me..."
- "I have family in Karachi and we are leading normal lives going about our daily work, parties, schools and all, a few changes like more uniformed men and barriers not a big problem, in fact most of us are glad that Musharraf took this action, he should have done this earlier... If any Pakistani leader is to be trusted with leadership it is Musharraf, not traitors and looters..."

We believe that the first comment may have been useful for other people in similar circumstances as the message author, and could have spurred corrective actions on their part. The second comment seems to have instead increased the diversity of opinions expressed about the event. Therefore, both these comments improved the usefulness of the original article in different ways, by exploring aspects of the event that had not been considered earlier.

Example 2: Economist: The Economist published an article titled *Malaria and how to beat it* on January 31^{st} 2008, about a study in Kenya which concluded that malaria nets distributed for free produced better results than when they were sold for nominal prices [5]. The study was meant to counter the popular notion that people do not attach significant importance to goods unless they pay for the goods. Consider two comments on the article.

- "It is a very timely article and subject. Brazil is having a yellow fever scare, which is also transmitted by mosquitoes, and I have not seen any of the measures The Economist mentions in the articles published by Brazilian newspapers, just vaccination, which can be dangerous for people with some illnesses."
- "The Acumen fund took a different approach to this same solution, with the added benefit of capitalism. The science of fighting malaria with an insect barrier is good and effective. Agreed. But remove the aspect of just giving

[4] http://news.bbc.co.uk/2/hi/south_asia/7077310.stm

[5] http://www.economist.com/daily/news/displaystory.cfm?story_id=10610398

the poor some charity; and replace it with support for the establishment of a local business solution; and you solve the health problem, make progress on the economic situation, and allow people the dignity of helping themselves locally instead of just receiving largesse... Always be cautious about just giving some product en masse to a population. You may inadvertently be putting an important local economy out of business."

As in the first example, both these comments explored aspects of the topic that had not been examined in the original article, and improved its usefulness. The first comment extended the implications of the study to a similar disease, but in a different geographical and cultural setting. The second comment raised an issue which could have implications in the formulation of appropriate policies by governments and health agencies.

More examples are given in (Seth 2007), and it is evident that participatory messages evolve with time and gain usefulness with more participation. However, it is questionable whether each of the comments given above will be useful for every reader of the article. It is possible that the first comment in the BBC article will be particularly useful for batchmates of the recent graduate, and the first comment in the Economist article will be useful for health workers in Brazil, because both of the comments would help these respective groups of people to better *understand* the relevance of the articles for them. The second comments in both the examples may also be useful, but for a different reason of increasing the *scope* of the articles. What is more interesting, however, is that different people may have different preferences for these features of *understandability* and *scope*, based on their circumstances and degree of interest with reference to the message topic.

It is this simplification and increase of scope in participatory messages that we refer to as *context* and *completeness* respectively. Context helps to "situate" a message better with reference to the circumstances of a recipient, and leads to simplification of the message. Completeness helps to "associate" a message with other issues, or other viewpoints of the same issue, and conveys deeper and broader information to the recipient.

Since message authors and recipients are embedded in an underlying social network of friendships and acquaintances, we use insights from the *strength-of-weak-ties* hypothesis in social network theory (Granovetter 1973) to explain how context and completeness of messages may arise based on the implicit relationships between authors and recipients (Seth 2007). A participatory message written by a strong tie of a recipient, or having a large number of comments written by users strongly tied to the recipient, is likely to provide context to the recipient. Similarly, a message having participation from users weakly tied to a recipient, is likely to provide completeness.