On Modeling Trust in Social Media using Link Polarity[1]

Anubhav Kale, Pranam Kolari, Akshay Java, Tim Finin and Anupam Joshi

University of Maryland, Baltimore County, Baltimore MD 21250, USA

Technical Report EB06, May 2007

ABSTRACT

There is a growing interest in exploring the role of social networks to understand how communities and individuals spread influence. In a densely connected online world, social media and networks have a great potential in influencing our thoughts and actions. We describe techniques to model trust in social media and present experimental results on finding "like minded" blogs based on blog-to-blog link sentiment for a particular domain. Using simple sentiment detection techniques, we identify the polarity (positive, negative or neutral) of the text surrounding links that point from one blog post to another. We use a trust propagation model to spread this sentiment from a subset of connected blogs to other blogs and deduce like-minded blogs in the blog graph. We then extend the same technique to label main stream news sources as left- or right-leaning based on the links between blogs and news sources. Our results confirm that the simple heuristics to analyze the text surrounding links and our trust propagation model are highly applicable for domains having weak link structure. These techniques demonstrate the potential of using polar links for more generic problems such as detecting trustworthy nodes in web graphs.

1. INTRODUCTION

Social media is a dynamic and growing area that includes collection of blogs, wikis, forums, photos and videos sharing sites. According to wikipedia[1] "social media describes the online tools and platforms that people use to share opinions, insights, experiences, and perspectives with each other". What makes social media really interesting is the level of user participation and conversations around different topics. This leads to formation of communities around topics that can vary from politics, technology, arts to knitting or photography and public relations.

Blogs have become a means by which new ideas and information spreads rapidly on the web. Blogs in general contain mostly user generated content, as do other sites like delicious[2], flickr[3] etc. Bloggers link to interesting posts or might even comment on someone else's blog and these links tend to be the basis of conversations. As communities in social media like blogs emerge, there are thought leaders and individuals who have a significant share of the community's attention. *Influential* nodes in a social network can be responsible for starting a buzz or getting the community to notice a new trend or product. Monitoring and tracking both influential nodes and their opinions on the blogosphere, can thus have a significant number of applications in the realm of product marketing.

In this paper, we address the problem of modeling trust in the social media in general and blogosphere in particular. Our approach uses link structure of the blog graph to associate sentiments with the links connecting two blogs. (By "link" we mean the url that blogger ***a*** uses in his blog post to refer to blogger *b*'s post). We call this sentiment as *link polarity* and the sign and magnitude of this value is based on the sentiment of text surrounding the link. These polar edges indicate the *bias/trust/distrust* between the respective blogs. In order to associate a given blog *foo* to the set of its like-minded blogs, we *create* new *polar links* between all pairs of blogs using initial *polar links*. We use trust propagation models to "spread" the initial polarity values to all possible pairs of nodes. Finally, we compute the trust/distrust score for *foo* from the seed set of *influential* blogs (discussed later) to determine its community. In order to demonstrate that our technique of modeling trust is domain independent, we apply our methods to the main stream media news sources and label them as left or right leaning.

Our approach uses simple shallow natural language processing to determine *link polarity*, yet preliminary results [18] indicate that our approach has the potential to aid conventional community detection techniques based on path distance and reachability metrics. Since, we do not process entire blog-post text for sentiment detection between two blogs and use shallow NLP techniques, we speculate that the approach should scale well for real-time applications (e.g., analyzing blogs for dynamic situations like elections) than traditional off-line and computation intense approaches. This paper presents some of our results in the domain of blogosphere, however a long-term goal of our work is to deduce trustworthy nodes for a given node in any web-graph. We believe that *directed polar links* have a tremendous potential for addressing this hard problem.

1.1 Background

Bloggers typically discuss views about varied topics and are based on personal experiences. Such views are expressed almost instantaneously as soon as any new event occurs. The blogosphere has matured a lot since its inception and hence, when an event occurs, the first reaction is to turn to the blogosphere to see what people are saying about it. For example during the London bombings, people were interested in finding first hand reports, pictures, emotions and experiences of Londoners. As time progressed, people might have looked for more information about the event - what happened? Why? How many people were killed? Have there been any arrests? Which group(s) has claimed responsibility for this act? etc.

Suppose that your goal was to market a new kind of mp3 player which would compete with ipod. One of the starting points would be to use advertising products such as Google's popular Adsense network. Using content of the webpage, this service matches relevant web pages to advertisements that relate to the topic of the page. While this gives a wide coverage and a

[1] Partial support was provided by NSF awards ITR-IIS-0326460 and ITR-IDM-0219649 and funding from I.B.M. Contact the authors as {akale1,kolari1,aks1,finin,joshi}@cs.umbc.edu

1 http://en.wikipedia.org/
2 http://del.icio.us
3 http://www.flickr.com

significant audience, there is very little the advertiser can do to actively promote the product to the *right set of individuals*. Using a blog search engine one can find a ranked list of relevant blog posts for different generic query terms. However, most blog search engines use link based ranking schemes that measure popularity as opposed to influence. While a number of popular blogs may talk about ipods in general, if the marketing division of your company can target the set of blogs that has a negative bias about ipod then chances of spreading good word about the new mp3 player is considerably high than targeting the blogs having a strong positive bias about ipod already. Thus having an insight into the communities in social media can aid in accurately targeting key personnel for marketing new products.

Temporal analysis of the swing in trends among communities has interesting applications for scenarios such as elections where a study of cause and effect phenomena has tremendous potential to gain an insight into change in voters' (or bloggers') bias during the election campaign events. This further implies that a community detection system capable of performing highly efficient real-time analysis of streaming data from social media can play a vital role for analyzing the effects of a candidate's meetings, speeches etc during election time.

There has been considerable amount of work in cluster formation and community detection on web graphs, however to our knowledge; none of the prior work involves using polarity of links as a parameter for the problem of community detection. Also, most of the well-known clustering algorithms like [1] are based on the analysis of link structure and do not work well for sparsely connected graphs. Our work is an initial step to address this problem. The remainder of the paper proceeds as follows. Section 2 covers related work. Section 3 describes the details of our approach, heuristic and data modeling. Section 4 covers the experiments and we discuss conclusions and future work in section 5 and 6.

2. RELATED WORK

Sentiment analysis can be defined as determining the overall *polarity* of a given document. Researchers have focused on many interesting challenges in this area such as predicting correct polarity irrespective of references to different objects in the same text corpus, modeling the context of text for topic categorization, analyzing language specific nuances such as negated words, n-grams, metaphors and subtle expressions; to name a few. Turney [2] propose a simple unsupervised learning algorithm for classifying reviews on the web as "thumbs up" or "thumbs down". Turney's work is focused on using the "semantic orientation" of phrases which is calculated as the difference between the mutual information gain between a given phrase and "excellent" and the mutual information gain between the same phrase and "poor". This work provides a simple, yet effective way of handling the complex natural language processing problem of sentiment classification. Pang et al. [3] provide a detailed analysis of various machine learning algorithms for the movie review classification problem. Their analysis of the "thwarted expectations" in the domain of movie reviews indicates an interesting challenge in the domain of sentiment classification. Church et al. [4] present work on "word association norms" (classifying words based on the co-occurrence with other words in corpus). Their approach uses information theoretic models of mutual information for estimating word association norms and they claim that models based on information theory are more effective than the traditional and costly way of testing few thousands of subjects on few hundred words to determine word associations.

Hearst [5] uses cognitive linguistics to determine the *directionality* of a sentence. This approach is a loose-case effort for applications that do not have sufficient resources to engage into complex NLP techniques; however the approach is useful only if the cost of building and executing the proposed methods does not compromise the quality of results. The work is independent of any domain-specific ontologies and uses isolated text interpretation in the realm of a generic metamorphic model adopted from [6].

People on Web 1.0 and software agents on Web 2.0 have to interact with unknown entities (*strangers*) to accomplish a variety of online tasks. Most of the commercial e-commerce websites in Web 1.0 (e.g. Amazon[4]) rely heavily on models for representing trust based on ranking schemes. Since it is not practical for every entity (people or software agents) on the web to have explicit knowledge of trust about every other entity, it is required to predict the trust score for a stranger from the knowledge of trust scores for known (or trusted) entities. Researchers have focused on this problem of formally modeling the notions of trust, distrust, influence and techniques for deriving trust scores for unknown entities.

Huang and Fox [7] provide a formal framework for representing trust and study the transitivity of trust. They classify trust as "trust in belief" and "trust in performance" and prove the transitivity of "trust in belief". They define the concept of trust as *Trust is the psychological state comprising (1) expectancy: the truster expects a specific behavior of the trustee such as providing valid information or effectively performing cooperative actions; (2) belief: the truster believes that expectancy is true, based on evidence of the trustee's competence and goodwill; (3) willingness to be vulnerable: the truster is willing to be vulnerable to that belief in a specific context where the information is used or the actions are applied*. To the best of our knowledge, this is the most precise and complete definition of trust since it provides a domain-independent abstraction for the definition of trust. Gans et al. [8] argue for the importance of giving an explicit consideration for distrust on social networks. They propose a "TCD" model based on the notion of trust, network confidence and distrust. The idea of using "network confidence" as a parameter in social network simulations has interesting applications. Beth, Borcherding and Klein [9] have worked on the problem of determining the trust for an entity in the context of conflicting recommendations about its trustworthiness. They emphasize that the semantics of direct trust values differ from that of the recommended trust values. Their mathematical models for combining conflicting trust scores are based on the non-monotonicy property of trust and use arithmetic mean as the mode of aggregation.

Richardson, Agarwal and Domingos [10] have proposed a framework to represent trust and distrust on the semantic web. Their idea is to compute a subjective trust/distrust score for every user instead of assigning a global trust score to every user. They use "linear pool", "noisy OR" and "logistic regression" for the combination functions and their work is one of the few which is evaluated on a very large dataset from Epinions[5]. Yu and Singh [11] study the problem of adversaries in trust management systems. They provide a model to detect deception in the process of trust/distrust propagation in a networked environment and we

4 http://www.amazon.com/

believe that their models are very applicable to social networks. Kamvar, Schlosser and Garcia-Molina [12] propose a secure method to calculate global trust values for shared files in P2P networks. The goal of their "Eigen Trust" algorithm is to reduce downloads of inauthentic files using global trust scores assigned to each file. Guha et al [13] in their work titled "Propagation of trust and distrust" cover work related to trust propagation in multiple disciplines and claim that their work appears to be first "to incorporate distrust in a computational trust propagation setting". We found that their work was most complete and the trust propagation model suits well to our domain. Hence, our trust propagation approach is very similar to their work.

To the best of our knowledge, no prior work exists in the area of blogosphere to assign sentiments to links (what we term as *link polarity*) and use such polar links to model trust in blog graphs.

3. PROPOSED APPROACH

In this section, we describe our proposed approach and set the basis for experimental validations. We also provide some details on Guha's trust propagation technique wherever appropriate.

3.1 Link polarity

The term *Link Polarity* represents the opinion of the source blog about the destination blog. The sign of polarity (positive, negative or zero) represents whether the bias is for, against or neutral and the magnitude represents how strong or weak the bias is. In order to detect the sentiment based on links, we analyze section of text around the link in the source blog post to determine the sentiment of source blogger about the destination blogger. From our analysis of blog texts and interactions with regular bloggers, we observed that it is not necessary to analyze the complete blog post text to determine the sentiment. In fact, text neighboring the link provides direct meaningful insight into blogger A's opinion about blogger B. Hence, we consider a window of x characters (x is variable parameter for our experimental validations) before and after the link. Note that this set of $2x$ characters does not include html tags.

3.2 Sentiment detection

There has been considerable work on sentiment detection on free-form text. Researchers have experimented with various natural language processing techniques However, we do not need to employ any complex natural language processing techniques since bloggers typically convey their *bias* about the post/blog pointed by the link using fairly standard vocabulary. Hence, we use a corpus of positive and negative oriented words and match the token words from the set of $2x$ characters against this corpus to determine the polarity.

Since our corpus includes words in noun forms, it is essential for us to employ stemming on tokens. We apply stemming mechanism on all such tokens and then convert them into canonical form by eliminating characters such as commas, periods, exclamation marks etc. We observed that bloggers frequently use negation of sentimental words while indicating bias about another blog-post(*"What b <u>says</u> is not bad"),* hence our corpus also includes basic bi-grams of the form "not <positive/negative word>". Our experiments confirmed that the aforementioned simple techniques are very effective in deducing the text sentiment correctly.

3.2.1 Calculation of link polarity

The number of positive and negative words varies to a great extent (typically from 1 to 30 in window size of 750 characters) across multiple posts. Hence, it is necessary to normalize the results over some metric. We adopted the following formula for calculating the link polarity:

Polarity = (Np – Nn) / (Np + Nn)

Np : Number of positively oriented words

Nn : Number of negatively oriented words

Notice that our formula incorporates zero polarity links automatically. The term in the denominator ensures that the polarity is weighed according to the number of words matched against the corpus. This helps to differentiate all such instances where (Np - Nn) is the same but (Np + Nn) varies from a small value (minimum = 2) to a large value (typically, 20). Also, note that we do not incorporate the number of links in our polarity computation. We use summation as the aggregation technique for computing the polarity between two blogs. For our experiments, we choose a domain where "off-the-topic" posts within a single blog are minimal, hence the notion of summing post-post polarity values to generate a blog-blog polarity value works well. We will have to investigate better aggregation techniques for handing more noisy datasets or filter the dataset and then apply the method of summation.

3.3 Trust Propagation

Since blog graphs are not densely connected, we still do not have the trust scores between any given pair of nodes. Hence, we must employ some *sentiment spread* mechanism to calculate trust score between all pairs of nodes from the set of nodes having polar edges between them. There has been considerable amount of work in computer science as well as other disciplines on various aspects of trust definitions, trust metrics, trust propagation models and validation techniques. Guha et al [13] have proposed a framework to spread trust in a network bootstrapped by a known set of trusted nodes. They have evaluated their approach on a large dataset from epinions[4.] Guha's approach uses a "belief matrix" to represent the initial set of beliefs in the graph. This matrix is generated through a combination of known trust and distrust among a subset of nodes. This matrix is then iteratively modified by using "atomic propagations". Finally "rounding" technique is applied on the matrix thus generated so far, to produce absolute values of trust (yes or no) between all pair of nodes. The "atomic propagation" step incorporates direct propagation, co-citation, transpose trust and trust coupling. The overall trust propagation mechanism is represented using matrix operations (additions, multiplications and transpose). We adapt this approach with some modifications for our work. The section on experiments covers our modifications in greater details.

In order to form groups of "like-minded" blogs after the step of trust propagation, we take the approach of averaging trust score for all blog nodes from a predefined set of "trusted" nodes belonging to each community. A positive trust score indicates that the blog node belongs to the community *influenced* by the trusted node of that community. Specifically, we selected top three *influential* democrat and republican bloggers. (We address our

5 http://www.epinions.com/

notion of *influential* blogs shortly). A positive trust score for a blog *foo* from top three democrat blogs indicates that *foo* belongs to the democrat cluster and a negative score indicates that *foo* is a republic blogger. Notice that negative links thus help us to classify a blog into the right cluster even if it is not very well connected within its cluster. In order to determine the *influential* bloggers in each community we experimented with the heuristics of high incoming-degree, high outgoing degree and random subset of all nodes.

4. EXPERIMENTS

We now present the results of our experiments that demonstrate the feasibility and effectiveness of link polarity. Also, we describe the motivation behind choosing the political domain for our experiments and present a representative set of link polarity computations for some of the *influential* blogs.

4.1 Choice of domain

We decided to choose political blogs as our domain; one of the major goals of the experiments was to validate that our proposed approach can correctly classify the blogs into two sets: republican and democrat.

Through some manual analysis of the political blogs, we observed that the link density among political blogs is reasonably high and hence we could deduce the effectiveness of our approach by running our algorithms over fairly small number of blogs. In other words, we do not need to perform a large number of iterations of Guha's atomic propagations; about 20 iterations suffice to *create* polar links with sufficiently accurate polarity values between blogs that did not link to each other.

The dataset from Buzzmetrics [14] provides link structure between blog posts over 1.3 million blog posts. Hence, we needed to aggregate this post-post link structure to a blog-blog link structure. This implied that we should choose such a domain where there would be minimal number of off-the-topic posts from the same blog and political blogs fit this requirement perfectly. (We address this issue of determining link polarity based on specific topics in our discussion section).

From a business model point of view, political blogs are highly effective during election period to determine the trends among voters and a technique that can classify voters into multiple political biases would be extremely beneficial to various sources.

4.2 Parameters for trust propagation

Guha et al propose the notion of "trust and distrust" between two nodes in the graph where the same set of two nodes can trust or distrust each other. We believe that in our domain the initial belief matrix should incorporate both trust and distrust (positive and negative polarities from blog A to blog B). We experimented with all the models of computing final belief matrix and found that "propagated distrust" provides best results on our dataset. The results presented in our previous work [icwsm] used "one step distrust" as the model for final belief matrix computation. We believe that the idea of using "eigenvalue propagation" to determine final trust scores is generic and applies to any domain. Hence we used the same for our experiments.

We experimented with various values of the "alpha vector" to confirm that Guha's conclusion of using the values they proposed {0.4, 0.4, 0.1, 0.1} yields best results. Our experiments indeed confirmed that this set of values yield the most accurate results. We do not provide the results of our comparisons here, since this is not the contribution or the primary motivation of our work. Further, Guha et al recommend performing "atomic propagations" approximately 20 times to get best results. Since, we can not guarantee that such numbers would work in our domain; we took the approach of iteratively applying atomic propagations till convergence. Our experiments indeed indicate a value close to 20 (which we believe is dominated by the diameter of the graph under consideration), after which the final trust scores do not seem to improve. Finally, we do not incorporate the extra step of "rounding" in Guha's work since the sign of trust is sufficient to determine if the blog under consideration belongs to the democrat or republican set.

4.3 Parameters for link polarity

As explained in section 3, we used various window sizes around the links to fetch the token words to be used for sentiment detection. After some manual analysis of political blogs, we decided to experiment with 1000, 750, 500, 250 and 50 characters before and after the link under consideration. We expected to get some insights into what would be the right window size (and hence, the right number of words around links that yield more signal then noise) by varying this parameter.

4.4 Datasets

We studied the effectiveness of our approach over a graph of 300 blogs created from the link structure of buzzmetrics [14] dataset. We observed that in-degree as a heuristic works better over out-degree and random heuristics for selection of *influential* nodes for the seed set. Hence all the results that follow are based on the in-degree heuristic. Lada A. Adamic provided us with a reference dataset of 1490 blogs with a label of democrat or republican for each blog. Their data on political leaning is based on analysis of blog directories. Some blogs were labeled manually, based on incoming and outgoing links and posts around the time of the 2004 presidential election. Buzzmetrics does not provide a classified set of political blogs. Hence, for our experiments we used a snapshot of Buzzmetrics that had a complete overlap with this reference dataset to validate the classification results obtained by our approach.

4.5 Effect of Link Polarity

The results in Figure 1 indicate a clear improvement on classifying republican and democrat blogs by applying polar weights to links followed by trust propagation. We get a "cold-start" for democrat blogs and we observe that the overall results are better for republican blogs than democrat blogs. The results being better for republican blogs can be attributed to the observations from [15] that republican blogs typically have a higher connectivity then democrat blogs in the political blogosphere.

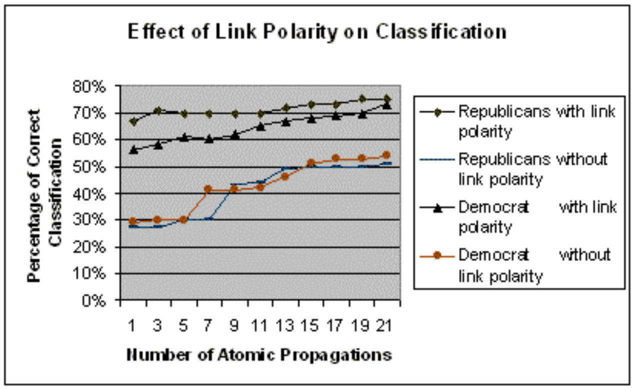

Figure 1: Using polar links for classification yields better results than plain link structure.

We are aware of the fact that the results need to be improved further, however it is interesting to note that there exists an upward swing in the accuracy using polar links. Thus, our idea of using trust propagation to *create* polar links between blogs that do not link to each other directly, helps to classify them. This clearly demonstrates the potential of our approach. We would like to note that the linear curve should not be generalized as a typical characteristic of blogosphere, it might be due to certain attributes of our dataset. We briefly discuss about further analysis of such trends in the discussion section (section 5).

4.6 Effect of window size on polarity determination

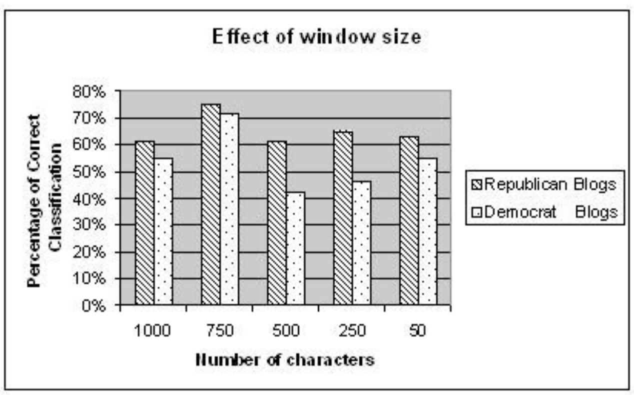

Figure 2: The correctness of classification depends on the optimal window size (around 750 characters) and decays on both sides of the optimal window.

The results in figure 2 indicate that 750 characters was the most appropriate window size for our dataset. If the window size is too small, our system becomes susceptible to short non-opinionated phrases around the link (e.g. *here is what xyz* says) which leads to a zero match of token words to corpus words in text surrounding link. On the other hand, if the window size is too large, our system becomes susceptible to analyzing text unrelated to the opinion expressed around the link. Another source of misinterpretation is the presence of other links in our window. Hence, we stop extending the window from the link whenever we hit the window size *x* or another link having a different domain name.

4.7 Effect of influential node selection

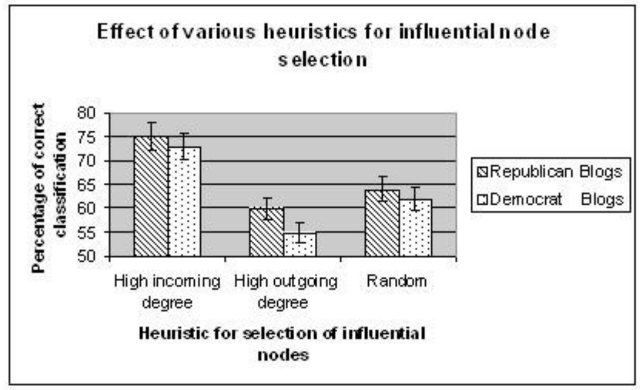

Figure 3: The correctness of classification is dominated by the heuristic used for influential node selection.

Since trust propagation is inherently a "push" model in which the trust/distrust is pushed from a subset of nodes to all nodes, high out-going degree seems to be the best heuristic for influential node selection. However, as the results in figure 3 indicate, high incoming degree is in fact the best heuristic. The random selection demonstrates intuitive results. In order to ensure true randomness in the process of random selection, we selected 3 nodes at random, repeated this process 10 times and averaged the results.

4.8 Sample polarity computations

The table in figure 4 depicts polarity values computed between some pairs of *influential* democrat and republican blogs. We present this data as a quick measure of demonstrating the potential of our work and make the following observations.

1. Trust propagation was effective in predicting the accurate polarity for DK-AT, even though our text processing did not yield the correct polarity initially. Thus, the errors introduced due to shallow NLP techniques are compensated by the step of trust propagation. The following illustration should make this claim clear.

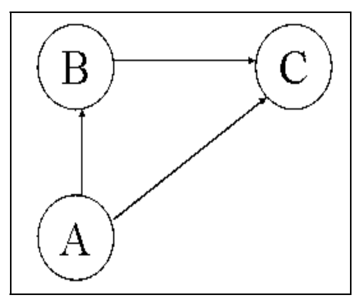

Figure 4: Direct Propagation as a unit step in trust propagation. If A trusts B and B trusts C then A should trust C with a fixed probability.

Suppose that our link polarity computation technique computes correct polarities for links A -> B, B -> C and incorrect polarity for A -> C. Without loss of generality, let us assume that the correct polarities are positive and the incorrect polarity is negative. Since the notion of "Direct Propagation" involves multiplying the initial belief matrix with itself to get the new set of beliefs ($M_{ik} = M_{ij} * M_{jk}$) , the sign of polarity for A -> C link will toggle from negative to positive if the collective evidence of positive polarity links is greater in magnitude than the magnitude of negative polarity link.

2. Trust propagation retained the sign of polarity if the initial computed sign of polarity was correct (e.g., AT-DK). In fact, trust propagation helped in assigning correct polarities to non-existent links (e.g., AT-IP).
3. The numbers in italics indicate the instances where trust propagation failed to assign correct sign to the polarity. However, notice that none of these had any polarity value to start with, so even if trust propagation did not assign the right sign to the link; it helped the clustering process for other blogs by establishing a connection between these blogs. We plan to work on a detailed analysis of such failures in order to get an insight into the effectiveness of our heuristics for link polarity determination. A preliminary analysis indicates that such failures are most likely due to the fact that there are fewer than three links between most blogs in our dataset, hence averaging over such small dataset leads to incorrect sentiment prediction occasionally.

From-To	Number of links	Polarity before trust propagation	Polarity after trust propagation
MM–MM	0	N/A	3.53
MM–DK	0	N/A	-2.9
MM-IP	0	N/A	2.2
MM-AT	0	N/A	*1.09*
DK-MM	0	N/A	-2.9
DK-DK	0	N/A	2.02
DK-IP	0	N/A	*1.71*
DK-AT	20	0	8.51
IP-MM	8	1	2.2
IP-DK	6	0	*1.71*
IP-IP	0	N/A	1.06
IP-AT	0	N/A	-7.19
AT-MM	0	N/A	*1.09*
AT-DK	5	0.342	8.51
AT-IP	0	N/A	7.19
AT-AT	0	N/A	3.57

MM–http://michellemalkin.com, DK–http://dailykos.com
IP–http://instapundit.com, AT–http://atrios.blogspot.com

Figure 5: Polarity values for some influential blogs in our dataset

4. Our validation techniques did not involve computing trust score for a blog *foo* from *influential* blogs in both communities. This implies that polar links help us by providing multiple ways to find like-minded blogs for *foo*. Thus, AT – IP polarity can correctly classify AT even if AT – MM polarity is incorrect. However, we are working on finding more sophisticated techniques to perform such validations in graphs having more than two communities and hence, we did not rely on non-scalable methodologies for our validations.

4.9 Main Stream Media Classification

Our test dataset from Buzzmetrics contained information about links from blog posts to main stream media news sources. As described in the previous sections, our experiments on determining the left or right inclination of blogs provide results with high accuracy. Hence, we decided to classify the main stream media sources using blog - media links. This serves as the evidence for the fact that our approach is not constrained to just the blogs - blog links but can be applied to other domains as well. Also, in the view of 2008 presidential elections, classification of main stream news sources has interesting business value.

4.9.1 Approach

Our approach for classification of main stream sources contains the same steps as described in section 3. Precisely, we compute the polarity for blog - media links and use the same trust propagation model to create a fully connected graph with polar links. Since we do not have a labeled dataset of left and right leaning main stream sources, we do not validate our results formally. Instead, we used human subjects and resources from web to assess the quality of classification. This further required us to consider only the popular media sources for our experiments, so that our human experts could provide meaningful comments on the results. Thus, the size of graph (in terms of nodes as well as links) for this experiment is significantly smaller than the precious experiments.

4.9.2 Results

Figure 6 represents the polarity values from influential republican and democrat blogs to media sources. The inclination of the media source can be interpreted from these results as follows:

If the polarity from republican blogs is positive and polarity from democrat blogs is negative, the media source has a right-leaning inclination

Else if the polarity from republican blogs is negative and polarity from democrat blogs is positive, the media source has a left-leaning inclination

Else if the sign of polarity from both republican and democrat blogs is same, the inclination of media source depends on the respective magnitudes.

We make the following observations from this data.
1. We can classify 24 out of 27 sources correctly.
2. Well-known left and right leaning sources like “guardian”, “foxnews”, “cnn”, “latimes”, “truthout” and “mediamatters” can be classified correctly.
3. The main outliers are “the nation” and “boston globe”. Our preliminary analysis shows that this is due to incorrect polarity computations. These errors in sentiment detection could not be compensated in the step of trust propagation due to the small size of graph.
4. Both left and right leaning blogs talk negatively about “nytimes” and “abcnews” and positively about “rawstory” and “examiner”.

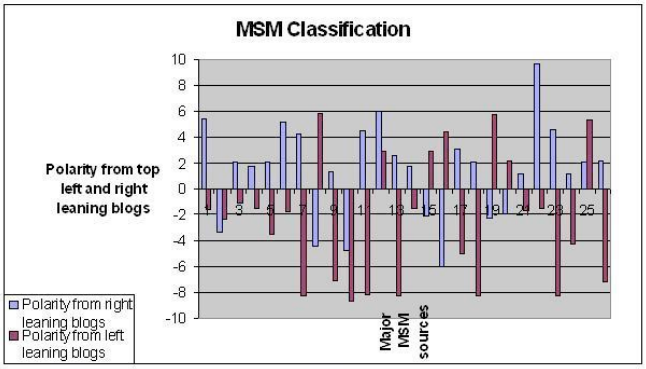

1	http://www.washingtonpost.com	14	http://www.truthout.org
2	http://www.nytimes.com	15	http://today.reuters.com
3	http://news.yahoo.com	16	http://mediamatters.org
4	http://news.bbc.co.uk	17	http://www.townhall.com
5	http://www.msnbc.msn.com	18	http://www.timesonline.co.uk
6	http://www.cnn.com	19	http://www.guardian.co.uk
7	http://www.usatoday.com	20	http://www.salon.com
8	http://www.latimes.com	21	http://www.thenation.com
9	http://www.boston.com	22	http://apnews.myway.com
10	http://www.abcnews.go.com	23	http://www.xaminr.com
11	http://www.foxnews.com	24	http://www.humaneventsonline.com
12	http://www.rawstory.com	25	http://www.dailybulltin.com
13	http://www.cbsnews.com	26	http://www.spectator.org

Figure 6: Polarities of well-known news sources from left and right leaning blogs

5. DISCUSSION

We are aware that we need to analyze results for our approach on a larger dataset. We are also investigating better techniques of validating our results and exploring various heuristics to determine topic of the link. Thus, topic as an extra attribute to the link would give us a fine-grained detail on positive or negative sentiment about a topic over a link and we believe that there are interesting applications of what we would like to term as “topical link polarity”. We are working towards new clustering techniques that incorporate polarity of links in the distance measure matrix and some of our preliminary results further confirm the effectiveness of link polarity. The idea of using link polarity suits well for all such domains where there exists a distinct set of different opinions (e.g. sports, windows vs. linux etc) and we believe that it has potential for deducing sub-communities from communities as well.

While we are optimistic about our approach, we would like to note that the traditional clustering techniques [16, 17] should be preferred over our approach when the graph is strongly connected. As explained before, the key contribution of our approach lies in classifying the *marginal* nodes (which either do not link or link very sparingly to the tightly connected cluster nodes). The idea of link polarity can help in predicting the swings in such *marginal* nodes and the temporal analysis of such swings can be very beneficial for advertising applications.

This paper presents preliminary results of our on-going work that demonstrates the effectiveness and feasibility of using polar links as evidence for clustering blogs into communities. Most trust models in use today rely on having *biased* links between nodes and our *polar* links can fit in such models perfectly. The focus of our future work is to make effective use of such polar links in various trust models to determine trustworthy nodes on web graphs.

The main contribution of our work lies in applying trust propagation models over polar links. We believe that the idea of *polar links* is very generic and can be applied to different domains. We demonstrated one such application in the domain of political blogosphere where we used natural language processing to deduce the link polarity. We would like to emphasize that the specific techniques to generate polar links is orthogonal to our main contribution. The idea of “Link Polarity” is very subjective to the domain under consideration. Hence in the discussion that follows, we give some insights into how our work can be extended to a very different domain of research conferences.

Co-authorship is an influential factor in the domain of research conferences. Suppose that the goal is to build a recommendation system for publications that assigns a *quality score* to the paper under review. Thus, the reader now has more metadata/feedback about the paper than just the contents of the paper. The system would be based on the data of papers, their authors and the affiliated universities from publicly available sources like DBLP. The reader can assign trust/distrust (*bias*) scores to the subset of researchers and universities that he is “influenced” with. This score can serve as a measure of explicit user-driven ``Link Polarity". The system can use metadata such as how many times the author of the paper under review has published to a well-known conference, how respected is the research community in the affiliated university and such, to generate more ``polar" links. Using the trust propagation models discussed in our work, the system can then compute the trust/distrust score for the paper under consideration. This application can easily be extended to detect “conflict of interest” as well.

6. CONCLUSION

We describe a novel approach for classifying blogs into predefined sets by applying positive or negative weights to links connecting the blogs. We validated our approach against a labeled dataset and the results are impressive. We use shallow natural language processing for the text around the links to determine the sentiments of one blog about another. This simple way of sentiment detection augmented by propagating trust using well-known trust models classifies the blogs with high accuracy. The results demonstrate the potential of using polar links for trust determination problems on web graphs and our future work will be focused on addressing this problem.

7. ACKNOWLEDGMENTS

We thank Nielson BuzzMetrics for making the blog dataset available. We thank Lada A. Adamic and Natalie Glance for allowing us to use their labeled dataset of democrat and republican blogs for our experimental validations.

8. REFERENCES

[1] M. E. J. Newman. Fast algorithm for detecting community structure in networks. Physical Review E, 69:066133, 2004.

[2] P. D. Turney, “Thumbs up or thumbs down: semantic orientation applied to unsupervised classification of reviews,” in ACL ’02: Proceedings of the 40th Annual Meetingon Association for Computational Linguistics, (Morristown, NJ, USA), pp. 17–424, Association for Computational Linguistics, 2001.

[3] B. Pang, L. Lee, and S. Vaithyanathan, “Thumbs up? Sentiment classification using machine learning techniques,” in *EMNLP ’02: Proceedings of the ACL-02 conference on Empirical methods in natural*

language processing, (Morristown, NJ, USA), pp. 79–86, Association for Computational Linguistics, 2002.

[4] K. W. Church and P. Hanks, “Word association norms, mutual information, and lexicography,” *Comput. Linguist.*, vol. 16, no. 1, pp. 22–29, 1990.

[5] M. A. Hearst, “Direction-based text interpretation as an information access refinement,” Text-based intelligent systems: current research and practice in information extraction and retrieval, pp. 257–274, 1992.

[6] Talmy L., “Force dynamics in language and thought. in parasession on causatives and agentivity”, Chicago Linguistic Society 21st Regional Meeting, 1985.

[7] J. Huang and M. S. Fox, “An ontology of trust: formal semantics and transitivity,” in *ICEC ’06: Proceedings of the 8th international conference on Electronic commerce*, (New York, NY, USA), pp. 259–270, ACM Press, 2006.

[8] Gans, M. Jarke, S. Kethers, and G. Lakemeyer, “Modeling the impact of trust and distrust in agent networks,” in *AOIS-01 at CAiSE-01*, 2001.

[9] T. Beth, M. Borcherding, and B. Klein, “Valuation of trust in open networks,” in *ESORICS ’94: Proceedings of the Third European Symposium on Research in Computer Security*, (London, UK), pp. 3–18, Springer-Verlag, 1994.

[10] M. Richardson, R. Agrawal, and P. Domingos, “Trust management for the semantic web.,” in *International Semantic Web Conference*, pp. 351–368, 2003.

[11] B. Yu and M. P. Singh, “Detecting deception in reputation management,” in *AAMAS ’03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, (New York, NY, USA), pp. 73–80, ACM Press, 2003.

[12] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, “The eigentrust algorithm for reputation management in p2p networks,” in *WWW ’03: Proceedings of the 12th international conference on World Wide Web*, (New York, NY, USA), pp. 640–651, ACM Press, 2003.

[13] Guha R, Kumar R, Raghavan P, Tomkins A. Propagation of trust and distrust. In: *Proceedings of the Thirteenth International World Wide Web Conference*, New York, NY, USA, May 2004. ACM Press, 2004.

[14] Nielsen Buzzmetric, www.nielsenbuzzmetrics.com

[15] Lada A. Adamic and Natalie Glance, "The political blogosphere and the 2004 US Election", in Proceedings of the WWW-2005 Workshop on the Weblogging Ecosystem (2005)

[16] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. Physical Review E, 69:026113, 2004.

[17] J. M. Kleinberg. Bursty and hierarchical structure in streams. Data Min. Knowl. Discov., 7(4):373-397,2003.

[18] Anubhav Kale *et al.*, "Modeling Trust and Influence in the Blogosphere Using Link Polarity", *Proceedings of the International Conference on Weblogs and Social Media (ICWSM 2007)*, March 2007