Building a Signed Network from Interactions in Wikipedia

Silviu Maniu Bogdan Cautis Talel Abdessalem

Télécom ParisTech - CNRS LTCI, Paris, France

{firstname.lastname@telecom-paristech.fr}

ABSTRACT

We present in this paper results on inferring a signed network (a "web of trust") from interactions on user-generated content in Wikipedia. From a collection of articles in the politics domain and their revision history, we investigate mechanisms by which relationships between Wikipedia contributors - in the form of signed directed links - can be inferred based their interactions. Our study sheds light into principles underlying a signed network that is captured by social interaction. We look into whether this network over Wikipedia contributors represents indeed a plausible configuration of link signs. We assess connections to social theories such as *structural balance* and *status*, which have already been considered in online communities. We also evaluate on this network the accuracy of a learned predictor for edge signs. Equipped with learning techniques that have been applied in recent literature on explicit signed networks, we obtain good predictive accuracy. Moreover, by cross training-testing we obtain strong evidence that our network does reveal an implicit signed configuration and that it has similar characteristics to the explicit ones, even though it is inferred from interactions. We also report on an application of the resulting signed network that impacts Wikipedia readers, namely the classification of Wikipedia articles by importance and quality.

Categories and Subject Descriptors

H.2.8 [**Database Management**]: Database Applications—*Data Mining*

General Terms

Algorithms, Experimentation

Keywords

Online communities, social applications, web of trust, signed networks, Wikipedia

1. INTRODUCTION

Large online communities that contribute and share content account nowadays for a significant and highly qualitative portion of the data on the Web. Examples of collaborative applications oriented towards building repositories of quality user-generated content include online encyclopedias (Wikipedia[1], Knol), photo sharing sites (Flickr) or rating sites (Epinions). An important trend in such platforms aims at exploiting user relationships, links between users (e.g., social links), in order to improve core functionalities in the system. For instance, search, recommendation or access control can benefit from socially-driven approaches. This is especially the case when links can be viewed as being *signed*, indicating a *positive* or *negative* attitude; possible meanings for positive links could be trust, friendship or similarity, while for negative links they could stand for distrust, opposition or antagonism. In settings where explicit relationships do not exist, are sparse or are inadequate indicators of one's attitude towards fellow members of the community, it becomes thus important to uncover *implicit* user inter-connections, positive or negative links, from relevant user activities and their interactions.

This paper presents a study of the interaction patterns between Wikipedia contributors and of the relationships that can be inferred from them. For a collection of 563 articles from the politics domain[2], starting from the revision history, we investigate mechanisms by which relationships between contributors - in the form of signed directed links - can be inferred from their interactions. We take into account *edits* over commonly-authored articles, activities such as *votes* for adminship, the *restoring* of an article to a previous version, or the assignment of *barnstars* (a prize, acknowledging valuable contributions).

The signed network we build is based on a local model for user relationships: for a given ordered pair of members of the online community - called in the following the link *generator* and the link *recipient* - it will assign a positive or negative value, whenever such a value can be inferred. This could be interpreted as subjective trust / distrust in a contributor's ability to improve the Wikipedia, and we call the set of such values in the network the "web of trust". In short, our approach aims at converting interactions into indicators of user affinity or compatibility: to give a brief intuition, deleting one's text or reverting modifications (backtracking in the version thread) would support a negative link, while surface editing text or restoring a previous version would support a positive one.

We believe that our work provides valuable insight into principles underlying a signed network that is captured by social interactions. We look into whether the network over Wikipedia contributors, called hereafter WikiSigned, represents indeed a plausible configuration of link signs. First, we assess connections to social theories such as *structural balance* and *status*, which have been tested in similar online communities [8]. Second, we evaluate on

[1] www.wikipedia.org.

[2] http://en.wikipedia.org/wiki/Wikipedia:WikiProject_Politics

DBSocial '11, June 12, Athens, Greece

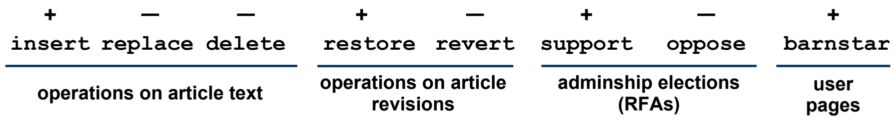

Figure 1: The interaction vector (from a generator to a recipient).

WikiSigned the accuracy of a learning approach for *edge sign prediction*. This amounts to exploiting existing links - in particular *link triads* - to infer new links and could be viewed as *propagation* of signed relationships. Equipped with learning techniques that have been applied in previous literature [7] on explicit signed networks, namely

- Slashdot (friend-foe tags),
- Epinions (trust-distrust tags),
- Wikipedia adminship votes (support-oppose votes)

we obtain good accuracy over the WikiSigned network (better than the one achieved in [7] over a Wikipedia votes network). By cross training-testing we obtain strong evidence that our network does reveal an implicit signed configuration and that these networks have similar characteristics at the local level, even though WikiSigned is inferred from interactions while the other networks are explicitly declared.

There are many opportunities that present to us for exploiting such a network at the application level, e.g., in the management tasks of contributors. We discuss in this paper one application that also impacts the readers, namely the classification of Wikipedia articles by importance and quality. The intuition here is that such article features depend on how contributors relate to one another.

A core contribution of this paper is a thesis: user interactions in online social applications can provide good indicators of implicit relationships and should be exploited as such.

The rest of this article is organized as follows. In Sections 2 and 3 we present the interaction extraction methodology and the reasoning behind building the signed network, respectively. The experimental results on our derived network are presented in Section 4 and the results of the article prediction model are presented in Section 5. In Section 6 we give a brief survey of the related work in the area of deriving a web of trust and Wikipedia interactions. Finally, in Section 7 we outline some possible extensions of this work.

The networks used in this paper are available at `http://www.infres.enst.fr/wp/maniu/datasets`*. The raw Wikipedia data are available upon request.*

2. EXTRACTION METHODOLOGY

The main context of interaction in Wikipedia is the collaborative editing of text on articles. However, the community itself is not restricted to such interactions. In order both to keep a minimum editorial standard and to limit the actions of low-quality contributors in the community, the contributors of Wikipedia participate in high-level interactions that are not directly related to the editing of articles. Project pages, user pages, administrator elections, etc, can serve the purpose of raising the quality level of Wikipedia articles.

We can thus separate the interactions into two main categories: *interactions on article content (text)* and the *community interactions*. We present the extraction methodology for each of them next.

Interactions on article content. For measuring the interactions on article content we have extracted the following measures: the amount of words inserted, deleted and replaced in the text of the article and the number of article revisions (i.e., versions) that have been restored or reverted/discarded.

For establishing these interaction metrics we have extracted the full revision history of a corpus of 563 articles from the Politics domain of the English Wikipedia, composed of 910209 total revisions and a total of 197798 unique contributors (we have not filtered the anonymous contributors or the Wikipedia bots in this extraction).

The revision of a given article A at time t can be seen as a triple:

$$R_A^t = (auth_A^t, txt_A^t, comm_A^t)$$

composed of the author (or the contributor) who issued the changes on the article, the text resulted from the modification and the comment used by the author to describe the modification.

A contributor $auth_A^t$ has two actions at her disposal: she can either *edit* the text of an article or *revert* the text to a previous version of it. We consider these two actions as independent and mutually exclusive (i.e., the author cannot, at time t, both edit and revert the article).

In order to quantify the interactions between authors, we establish, for each revision, the ownership at word level based on the text difference between two consecutive revisions of an article. This is represented as a list O^t of triples of the form $(owner, \Delta_{start}, \Delta_{end})$ for each revision R^t, consisting of the owner id and the span of her ownership (encoded as deltas in words from the start of the document). This list is created using a text diff algorithm that outputs a list of the operations needed to reach txt_t from txt_{t-1}. These operations represent the amount of text (in words) that $auth^t$ has either deleted, inserted, replaced or kept unchanged. Following this, we establish the new ownership list, as follows:

1. for text inserted and replaced, the new author is $auth^t$ and the deltas are the new positions resulted from the text difference algorithm,
2. for deleted text, the previous author and its positions are removed from O^t and the remaining offsets are updated to account for the missing text.

The interaction thus formed is between the author of the current revision and the owners of the text in the previous revision, as follows:

$$text_{auth^t,a}^t = (ins_{auth^t,a}^t, del_{auth^t,a}^t, rep_{auth^t,b}^t)$$

where the components represent a count of the words in each interaction, and $a \in O^t$.

By parsing the text of $comm^t$ we can have an indication of the revision R^t being in fact a reversion to a past revision R^{t-x}. When this is the case, we can form new interactions between contributors. First, a *restore* interaction is defined as the interaction between $auth^t$ and $auth^{t-x}$. Then, the ownership list is reverted to O^{t-x}. Finally, we establish *revert* interactions between $auth^t$ and the authors $\{auth^{t-x-1}, ..., auth^{t-1}\}$. For each possible pair of contributors a and b we represent their interaction at time t in this dimension, as follows:

$$rev\_res^t_{a,b} = (revert^t_{a,b}, restore^t_{a,b})$$

where the components encode how many times each type of interaction has occurred.

In our case, the total number of content interactions (textual and reverts-restores) that were established using this model was 30670861.

Community interactions. Using the list of unique contributors, resulted from extracting the article history, we can further crawl the pages of Wikipedia that are not articles (in some sense, the metadata of Wikipedia) to retrieve the contributor user pages. We can thus establish if they have participated in the Wikipedia "Requests for Adminship" elections (RFAs), either as voters or as candidates. This list of community interactions is by no means exhaustive, as contributors can participate in a variety of other interactions. An important one represents the debates between the contributors, present on the talk pages attached to the articles. One could also exploit such interactions (e.g., by means of natural language processing and sentiment detection). This goes beyond the scope of this paper, but we may follow this direction in future research.

By crawling the pages for RFAs (filtering out the pages for which the candidate is not in our contributor list), we can track the votes cast by the contributors from our list, votes that can be either positive or negative. This election interaction will be represented as follows:

$$election_{a,b} = (vote^+_{a,b}, vote^-_{a,b})$$

where a is the voter and b is the candidate for adminship.

Finally, by crawling the user pages of all the contributors in our list, we can retrieve the Wikipedia *barnstars*. Barnstars are prizes that users can give to each other for perceived valuable contributions and are usually present on the receiver's page. We have thus retrieved the user profile pages of all our contributors and extracted this information, resulting in the $barnstars_{a,b}$ measure which denotes the number of barnstars given by contributor a to contributor b.

2.1 Aggregating the interactions

For representing the global interaction between a pair of contributors, we have used an aggregation over the interactions. The aggregation was performed by summing the interactions on article content, as follows:

$$text_{a,b} = (\sum_t ins_{a,b}, \sum_t del_{a,b}, \sum_t rep_{a,b})$$

$$rev\_res_{a,b} = (\sum_t revert^t_{a,b}, \sum_t restore^t_{a,b})$$

(Note that we do not need to aggregate the content interactions, as by definition they give the number of times the respective interaction occurred.)

This aggregation yielded a total number of 17262082 interactions.

3. BUILDING THE SIGNED NETWORK

The four types of interactions presented previously (edits, reverts-restores, election votes and barnstars) can be viewed as an interaction vector from a generator to a recipient. This vector will form the basis for inferring signed edges between users. We describe next how these are further organized and then interpreted as positive or negative units.

Our approach is the following: we give each of the atomic interactions previously identified (text insert, delete and replace; reverts and restores; votes cast and barnstars) a positive or negative interpretation. For instance, for edits on text, we interpret inserts as positive while replacements and deletions of text are seen as negative. Then, the restores of a revision are interpreted as positive interactions, while conversely the reverts of a revision are negative ones. The votes cast in an election are recorded accordingly as positive or negative interactions, while the presence of barnstars is seen as a positive interaction.

Figure 1 summarizes the components of this interaction vector and the sign interpretation of each (positive or negative).

Note that these vectors denote *directed interactions*, from a generator to a recipient, and the presence of interactions in one direction does not necessary imply that interactions in the other direction exist.

Then, for deciding a final link sign, for a given pair (a, b) of contributors, we used the following straightforward heuristic. Each atomic interaction votes with its weight (or its magnitude) by the positive or negative interpretation of the higher-level interaction.

For determining the vote of the textual interactions, we have used Kendall's τ coefficient as follows:

$$\tau_{text} = \frac{ins_{a,b} - (del_{a,b} + rep_{a,b})}{ins_{a,b} + del_{a,b} + rep_{a,b}}$$

giving us a measure within the $[-1, 1]$ interval. In order to better control the link formation for textual interactions, we have used a threshold (both positive and negative) on the τ_{text} coefficient for deciding the vote of the textual interaction. We also recorded the size of the textual interaction, $size_{a,b}$ representing the number of different revisions over which the two contributors interacted. We used a parameter k which acts as a threshold on $size_{a,b}$ and regulates when the vote of textual interactions is taken into account. (Note that if one would only be interested in the sign of the interaction on text, computing the difference between the number of words inserted and the number of words replaced and deleted would suffice.)

Reverts and restores vote for the sign of $rev\_res_{a,b}$ and adminship votes for the sign of $election_{a,b}$. The barnstars can only vote positively or be absent from the vote.

Finally, these votes are aggregated into a link sign from a generator to a receiver, by the sign of the sum of the votes of each interaction type.

In our experiments, we have used a threshold value of 0.5 on τ_{text}, a threshold of 10 for the minimum number of words interacted upon and we variated k. We chose as our most representative network the one given by $k = 2$ (the median value for the number of interactions in our entire corpus is 1, meaning that over half of the contributor pairs interacted on text only once).

The WikiSigned network obtained in this way has 138592 nodes and 740397 edges, of which 87.9% are positive (a link proportion that is very similar to the ones of the existing signed networks). Please note that our mined election network (which can be seen as an explicit signed network) could not have skewed the results, as the total number of election interactions extracted represents less that 10% of the links of WikiSigned.

We present in Table 1 network for $k = 2$ and, for comparison, the ones for other values of k (3, 4 and 5). Also, to better understand the provenance of WikiSigned links, we describe the networks obtained when ignoring the textual interactions or when, instead, using only these interactions (for a value of $k = 2$).

network	nodes	edges	positive	negative
WikiSigned k=2	138592	740397	87.9%	12.1%
WikiSigned k=3	131544	590505	86.2%	13.8%
WikiSigned k=4	126559	497196	84.5%	15.5%
WikiSigned k=5	123070	439644	83.0%	17.0%
textual inter. k=2	73723	568488	96.2%	3.8%
non-text inter.	90188	178354	60.5%	39.5%

Table 1: WikiSigned features for varying k values.

triad	count	P(+)	lrn	bal	stat
t_1	2,513,952	0.98	0.1646		
t_2	125,765	0.88	-0.1589		
t_3	2,581,081	0.96	0.0197		
t_4	81,353	0.85	-0.0300		X
t_5	130,225	0.52	-0.3062		
t_6	44,530	0.32	-0.4268	X	
t_7	77,673	0.44	-0.4093		
t_8	39,642	0.34	-0.1849	X	
t_9	3,705,565	0.96	0.0186		
t_{10}	81,629	0.75	-0.2683		
t_{11}	387,386	0.89	-0.0546	X	
t_{12}	48,940	0.71	0.0575	X	
t_{13}	147,869	0.87	0.0201	X	
t_{14}	112,412	0.63	-0.2011	X	
t_{15}	60,768	0.79	0.0817	X	
t_{16}	33,920	0.38	-0.1388	X	X

Table 2: Statistics on triads. The X symbol marks a contradiction with theory.

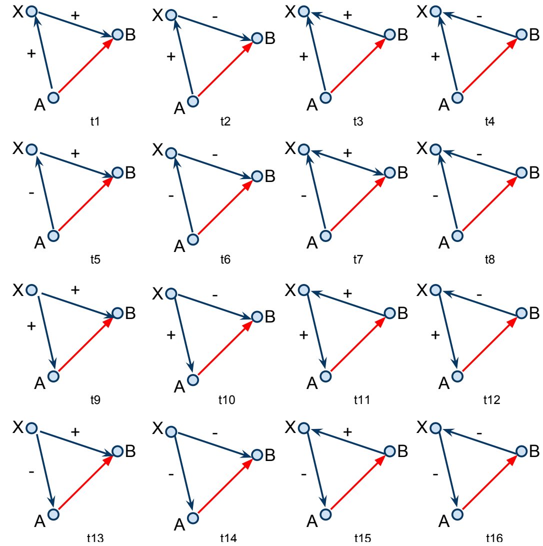

Figure 2: The directed link triads.

4. EMPIRICAL VALIDATION

We present in this section our analysis of the WikiSigned, testing mainly whether its structural properties are consistent with a signed network. For that, we rely on social theories on the formation of links between individuals, which have been tested in similar online communities, and on comparison with explicit networks. First, at the global level, we study the properties of WikiSigned in relation to the theories of structural balance and status. Then, at the local level we study how accurate an edge sign prediction can be performed on WikiSigned. Finally, we consider the indegree and outdegree distributions of contributors and look into how well they into a power-law distribution.

Global properties of WikiSigned. We first analyze the global properties of WikiSigned, checking whether overall it represents a plausible configuration of link signs. For that, we study the role of "link triads" in our signed network. We used a methodology that has already been employed on explicit networks, in [8, 7], allowing us to compare the properties of our network with the existing ones.

A triad represents the composition between the link from A to B and the possible links to a third party node X. Depending on the direction and sign of the link connecting A, and B, with X, there are sixteen such types of triads.[3] Some of the triads are representations of well-known social scenarios: t_6 is a representation of "the enemy of my enemy", t_1 of "trust transitivity", t_9 is a triad in which X points positively to both A and B. Figure 2 illustrates these triads.

We looked at the distribution of link triads and the proportion of positive $A-B$ links in each type of triad. We found that both measures are very similar with the ones reported in [8] (see columns *count* and *P(+)* in Table 2).

Next, we study the configuration of our network in comparison with two social theories, *status* and *balance*, theories that aim to define and predict the formation of links between individuals. *Structural balance theory* posits that triads which are "balanced" (i.e., have either one or three positive link signs, in an undirected sense) are more prevalent in real-world networks that the other types of triads [5]. *Status* posits that a directed negative link between A and B means that A regards B as having lower "status", while a positive link mean that A regards B as having higher "status" [4]. As such, for the network to have the same properties as the ones predicted by balance theory, in triads t_1, t_3, t_6, t_8, t_9, t_{11}, t_{14} and t_{16} the $A-B$ link should be positive, while it should be negative in the rest of the triad types. For a network to be in line with status, triads t_1, t_4, t_{13}, t_{16} should have a positive $A-B$ link and triads t_6, t_7, t_{10}, t_{11} should have a negative one.[4]

Link prediction in signed networks has been studied in [7], by training a link prediction model using logistic regression learning on a feature vector consisting of the total number of triads of each type that the link participates in. We have used the same methodology for training the model on WikiSigned, with 10-fold cross-validation and a balanced set of negative and positive links with a minimum link embeddedness of 25 (i.e., the total number of triads in which each link participates). Since the positive links represent a large majority of the link signs, naively predicting all link signs as positive would have an accuracy of 0.879 (i.e., the proportion of the positive links in the network). To avoid this bias, we have

[3] As in [8], we encode them by a summation starting at 1 and adding 8 for the $A-X$ link if it is pointing backwards then 4 if the link is negative; for the $X-B$ link we add 2 if the link is backward and 1 if the link is negative. This gives us the id of the triad.

[4] Status predicts link signs only for these triads.

	Epinions	Slashdot	Elections	WikiSigned
Epinions	0.926	0.905	0.787	0.765
Slashdot	0.929	0.806	0.792	0.716
Elections	0.922	0.895	0.814	0.775
WikiSigned	0.882	0.839	0.755	0.852

Table 3: Predictive accuracy in training on the row data and testing on the column data. The first three networks are the datasets used in [7].

randomly selected as our training set 5000 edges for each link sign, via reservoir sampling.

The signs of the coefficients of the trained model are an indication of the influence that each triad type has on the final link sign. Hence, we can compare these signs with the predictions of the two social theories. We perform this comparison on WikiSigned, counting the contradictions with these theories, i.e., the differences between the sign of the learned coefficients for each triads and the prediction of the two social theories. We find that at a global level our interaction-based network is more consistent with the theory of status (two contradictions with the theory, in t_4 and t_{16}) similar to what has been observed in [8] on the Wikipedia election network (see columns *lrn, bal* and *stat* in Table 2; X marks a contradiction with the social theory in the column).

Local properties of WikiSigned. For the local properties analysis of WikiSigned, using the same link prediction model, we tested the accuracy of predicting link signs. The predictive accuracy thus obtained was of 0.852 with an AUC of 0.924.

Furthermore, to better understand how WikiSigned relates structurally to the explicit networks, we have also applied this learning methodology over the three explicit networks considered in [7], asking the following question: how well a predictor learned on one network performs when applied on another network (see Table 3)[5]. First, one can notice that our results that use and apply to explicit networks are almost identical to the ones reported in [8]. WikiSigned performs better than the election network, in that prediction on itself is worse than self-prediction over Epinions and Slashdot, while learning the predictor on WikiSigned and applying it on both Epinions and Slashdot yields good prediction rates, the inverse performing slightly worse. All this indicates that these networks have quite similar characteristics at the local level, even though our network is inferred from interactions while the other three are explicitly declared by users.

Fitting into power-law distributions. Previous work has found that the distribution of outdegrees and indegrees of nodes in online social networks generally follows a power-law curve [10]. We measured a similar aspect, defining an *absolute degree* of a node as the difference between the positive and negative links pointing to or from that node. For inlinks or outlinks only, we have the related concepts of *absolute indegree* and *absolute outdegree*.

We have plotted the complementary cumulative distribution functions (CCDF) on a log-log scale for the two explicit networks of Epinions and Slashdot and for WikiSigned. The results are presented in Figure 3.

One can see that all the tested networks exhibit power-law characteristics, with various exponents. In the case of WikiSigned, one can see that the distribution follows closely the ideal power-law fit (the dotted line), while in real networks, for big absolute indegrees and outdegrees the curve is below the ideal fit. In our view, this could be due to the fact that the link inference for WikiSigned is based on voting, without including potentially complex link formation conditions that may occur in explicit networks (for instance, the formation of links using the knowledge about other links).

5. EXPLOITING WIKISIGNED AT THE APPLICATION LEVEL

We also investigated the usefulness of having the signed network in applications, by considering how link structure can be exploited in the classification of articles. There are two article features that are explicit on the homepage of the Wikipedia Politics project[6]: the article quality and the article importance (or priority). In our dataset, we have articles that span the top 5 article qualities (Featured Articles, Great Articles, A-class Articles, B-class articles and C-class articles) and all the importances (Top, High, Mid, Low).

For our experiments, we have separated the article qualities and importances into two classes (top-tier and bottom-tier). For the article importance, we have considered the Top, High, Mid as the top tier and the Low importance as the bottom tier, randomly sampling 150 articles for each. As the A-class of articles contains only 8 articles, we have excluded this class for the training, and we have randomly sampled 50 articles from each remaining class. Furthermore, we have categorized as top-tier the FA and GA articles and the B and C-class articles as bottom-tier. This resulted in two equally balanced datasets: 100 for each article quality tier, and 150 for each article importance tier.

We have used the following set of features for each article: the number of *authors*; three features (total, positive and negative) for each of the following: *outgoing links* (links from the authors towards other contributors), *incoming links* (the links from other contributors towards the authors) and *inside links* (links from authors to authors); and the following information about the contributors in the article: the number of *incoming* total positive and negative links (in the entire networks) for the contributors of the article, how many of them have more positive links than negative and vice-versa. The same information is also extracted for *outgoing* links, giving us a total of 18 features for our article prediction model.

We report the predictive accuracy we obtained via logistic regression in Table 4. Following the intuition that more important articles have a larger participation and thus more links, we tested the predictive power of these two values (*contributors* and *contribs.+links*) alone. We found that, while using knowledge about positive or negative links in separation does not provide better accuracy, their combination yields better results (*contribs.+soc. links*). This suggests that the characteristics of an article are not defined solely by the number of contributors, but also by their relationships with other Wikipedia contributors. When we also introduce the information about the contributors (*contribs.+ soc links + rep.*), we see further improvement, especially in the case of quality, which seems to support the intuition that the quality of an article is determined by the "quality" of its contributors.

type	Importance	Quality
contributors	0.691	0.518
contribs.+links	0.743	0.835
contribs.+ soc links	0.749	0.895
contribs.+ soc links + rep.	0.756	0.935

Table 4: Predictive rates for article importance and article quality.

6. RELATED WORK

To the best of our knowledge, this is the first study on inferring a signed network (a "web of trust") directly from user interactions. The work that is closest in spirit to ours uses a semi-supervised approach and existing links to build a predictor of trust-distrust from interactions in Epinions [9]. Several papers deal with edge sign prediction using an existing network, among which [4, 7] (see also the references therein). These approaches use the explicit signed network, either for verifying the accuracy of the predictor or as a basis for the inference of new links. In [3], the authors deal with interactions between contributors of Wikipedia articles, using the

[5]The datasets for Epinions, Slashdot and Wikipedia Elections are available at http://snap.stanford.edu/data/index.html

[6]http://en.wikipedia.org/wiki/Wikipedia:WikiProject_Politics

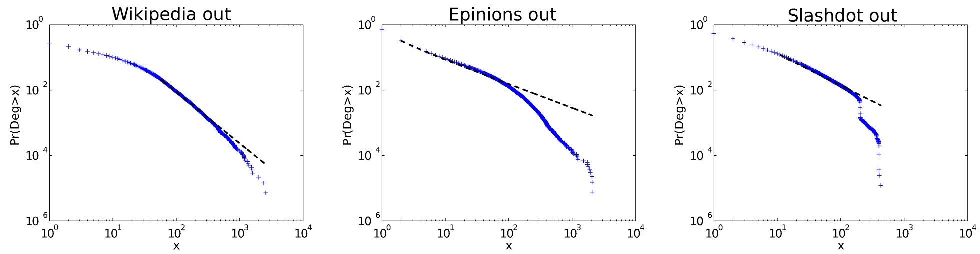

(a) Absolute outdegree: WikiSigned, Epinions and Slashdot.

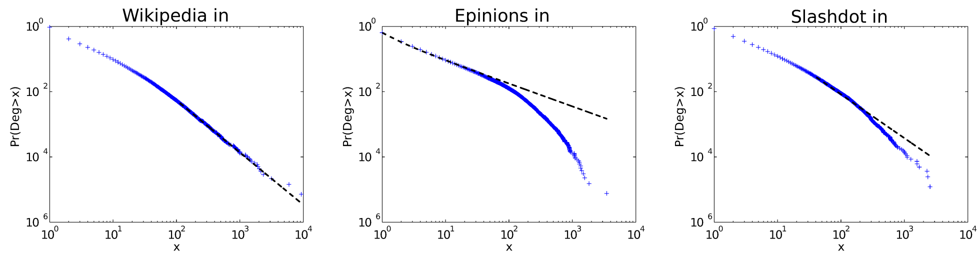

(b) Absolute indegree: WikiSigned, Epinions and Slashdot.

Figure 3: Absolute indegree and outdegree power-law fits and CCDFs for different signed networks.

concept of an "edit network" to measure the degree of polarization in articles. In [2, 1], a contributor reputation system and a measure of trustworthiness of text are derived based on their interactions over Wikipedia content. Another paper that experiments with reputation systems using the editing interactions between contributors is [6].

Previous publication. An extended abstract of this work, using a smaller Wikipedia dataset, was presented in a poster paper at WWW2011.

7. FUTURE WORK

We intend to use the link prediction methods validated by our results to further enrich the WikiSigned network. At the application level, one goal is to establish and exploit a reputation system for contributors, for example based on exponential ranking on the derived links (while also taking into account the negative links). Another goal is to propose a text-trust system that is similar and comparable to the one in [1].

A natural extension to our interaction model is to add other types of interactions. One could for instance exploit commenting and debating interactions by using natural language processing and sentiment detection techniques. Also, by taking into account the timeline of interactions, one could think of approaches that can model also the evolution of interaction measures.

8. ACKNOWLEDGEMENTS

We are grateful to the anonymous reviewers for their comments. This work was supported by the French ANR project ISICIL and the EU project ARCOMEM FP7-ICT-270239.

9. REFERENCES

[1] B. T. Adler, K. Chatterjee, L. de Alfaro, M. Faella, I. Pye, and V. Raman. Assigning trust to Wikipedia content. In *WikiSym*, 2008.

[2] B. T. Adler and L. de Alfaro. A content-driven reputation system for the Wikipedia. In *WWW*, 2007.

[3] U. Brandes, P. Kenis, J. Lerner, and D. van Raaij. Network analysis of collaboration structure in Wikipedia. In *WWW*, 2009.

[4] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *WWW*, 2004.

[5] F. Heider. Attitudes and cognitive organization. *Journal of Psychology*, 1946.

[6] S. Javanmardi, C. Lopes, and P. Baldi. Modeling user reputation in wikis. *Stat. Anal. Data Min.*, 2010.

[7] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *WWW*, 2010.

[8] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Signed networks in social media. In *CHI*, 2010.

[9] H. Liu, E. Lim, H. W. Lauw, M. Le, A. Sun, J. Srivastava, and Y. Kim. Predicting trusts among users of online communities: an Epinions case study. In *EC*, 2008.

[10] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *IMC*, 2007.