

Towards a Social Media-Based Model of Trust and its Application

Erik Boertjes, Bas Gerrits, Robert Kooij, Peter-Paul van Maanen,
Stephan Raaijmakers, and Joost de Wit

Netherlands Organisation for Applied Scientific Research (TNO)
{erik.boertjes, bas.gerrits, robert.kooij, peter-paul.vanmaanen, stephan.raaijmakers,
joost.dewit}@tno.nl

Abstract. In this paper we describe the development of a model for measuring consumer trust in certain topics on the basis of social media. Specifically, we propose a model for trust that takes into account both textually expressed sentiment and source authority, and illustrate it on a specific case: the iCloud cloud computing service of Apple and its reception on Twitter. We demonstrate that it is possible to parameterize a trust function with weights that interpolate between the contribution of sentiment and the authority of the tweet senders. Feedback data containing perceived trust in the iCloud service was gathered from a community of users. On this data, our model was fitted and evaluated. Finally, we show how such a fitted model can be used as a basis for a visualization tool aimed at supporting professionals monitoring trust, or to simulate implications of interventions. Our approach is a first step towards a dynamic trust monitor that is a viable alternative to more rigid, survey-based approaches to measuring trust.

Keywords: consumer trust, modelling, social media

1 Introduction

The Internet is developing increasingly into a vital infrastructure that constitutes the foundation of economic and social processes within our society. This development is unstoppable and will in the coming years bring about changes that go beyond information and data exchange. Development of facilities like “Internet of Things”, the semantic web and cloud computing all contribute to an increased digitization of society. The role of ICT as the foundation for the acceleration of the information society, is clearly expressed in the “Digital Agenda for Europe (2010-2020)”, which was presented by the European Commission early 2010¹. The success of the digital agenda is largely tied to the trust that society puts in the same developments in the digital plane. In the absence of trust, the essential condition for further dissemination and adoption of ICT services and facilities is lacking. For example, in 2008, the global consulting

¹ See http://ec.europa.eu/information_society/digital-agenda/index_en.htm.

firm Booz & Company estimated for the European digital economy in 2012 the economic difference between a scenario with high trust in ICT and one with low trust in ICT to consist of an enormous €124 billion. In this estimate of the economic asset of trust, the digital economy is assumed to consist of advertisements, content, e-commerce and access to ICT. Obviously, with the economical and societal increasing dependence on ICT, also the potential for abuse of this medium increases. Cyber criminals are actively exploiting the vulnerabilities of networks, and terrorist organizations use the internet to exchange information. Personal and sensitive data, profiles and digital identities of organizations and end users are stolen and resold. The consequences are both financial and emotional. It is clear from the discussion above that there is a need for increasing trust in ICT. To achieve this, we need to know what factors determine trust in ICT. It is also necessary to examine the relationship between trust and the adoption of ICT services further. It is important to notice that the actual trust organizations or end users have in ICT may be not in line with the actual risks at hand. This may be causing two types of problems. If there is high trust but the risk is actually high, then the user runs the risk of experiencing damage due to unjustified confidence. If the trust is low but the risk is actually low, then the user will tend not to use the service, hence the adoption of the service falls behind for unfounded reasons.

The aim of this paper is twofold. The first aim is to discuss the dimensions on which trust in ICT depends. This discussion will be mainly based upon the work of Kim et al. ([3]) and Corbitt et al. ([1]), in addition to trend reports that appear on an annual basis² or on a (bi-)quarterly basis³. The second aim this paper is to outline an approach that can be used to assess trust in ICT in near real-time, based on sentiment mining of social media. This approach opens up possibilities for monitoring communities, by either individuals, companies or institutions, in order to assess the current level of trust in certain topics, form or adjust opinions, and to undertake subsequent action, such as active participation in discussions.

2 Dimensions Determining Trust in ICT

According to Kim et al. ([3]) and Corbitt et al. ([1]), six dimensions determine trust in ICT, namely social, institutional, content, product, transactional and technological dimensions. The table below shows the dimensions of trust explained and elaborated into indicators. The indicators in this table affect the trust dimension they belong to, and thus the overall trust in a product or service.

When considering the dimensions in Table 1, one has to bear in mind that several actors influence trust in an ICT service. Although the actors for each service or product may differ, the following general groups can be distinguished:

- Service Providers: The provider of the service used.

² For instance Eurostat data, Ernst & Young Global Information Security Survey (GTISC).

³ For instance reports on phishing by the Anti-Phishing Working Group (APWG; <http://www.antiphishing.org/>) or by McAfee.

- Users: The recipient of a service, this can be both a consumer and another company.
- Technology Providers: All parties who provide the technology needed for the service, such as telecom operators, computer vendors, network infrastructure providers, phone manufacturers, etc.
- Third Parties: All other parties who have a facilitating role in providing a particular service, for instance a payment provider.

Depending on which service is used, various actors play a role. In addition, some players are involved in different roles. This makes the role of the actor in the security and trust in the product not always clear.

Table 1. Six dimensions determining trust in ICT.

<i>Trust dimension</i>	<i>Explanation</i>	<i>Indicators</i>
Social	Social factors affecting trust	Experience, reputation, peer pressure, culture, norms and values
Institutional	“Third parties” and institutional context affecting trust	Reputation, accreditation (trust marks), regulations, legal obligations
Content	(External) characteristics of the product or content	Design, customization (personification), brand visibility
Product	Product features that influence purchase/use decisions	Availability, quality, durability, price
Transaction	Characteristics for trust in transactions involved	Transparency, payment options, cost model, discounts
Technology	Characteristics of infrastructure and software related to security and effectiveness	Availability, integrity, authenticity, confidentiality, reliability

The conceptual framework to assess trust according to the above model is first to monitor the various indicators involved, then to map these, for each dimension, to a certain value. Finally, the six individual values are weighted in order to obtain trust in ICT. Obviously, many obstacles still exist in making such a framework operational. First of all, not all indicators are easily assessable, if at all. Even if data is available, data formats will often differ and originate from different sources. Secondly, the translation from indicators to the actual trust dimension is also far from trivial. Finally, to our knowledge, there is no agreement on how to combine the six dimensions by weighing them, to arrive at the final assessment of trust in ICT. All three issues are items for further research, but we will address a possible integration method in Section 3.

Even though the framework may not be fully applicable yet, there are some data sources available that provide useful input for it. As an example we will consider the dimension 'technology'. Most indicators under this dimension fall under the issue of information security. Security generally consists of the following three categories: availability (the information is available when needed), confidentiality (information is not disclosed to unauthorized individuals or systems) and integrity (data cannot be modified undetectably).

Data on these indicators is available through statistics agencies, such as the European Eurostat, who obtain their data through extensive surveys⁴. In addition vendors (such as Microsoft with their Intelligence Report) and security companies (such as McAfee with their Threats Report) report statistics on an annual or (bi-) quarterly basis. For example, according to Eurostat, the percentage of companies in the European Union (EU) that had their ICT service disrupted due to external attacks, was 4% in 2010. According to the Microsoft Security Intelligence Report 2011, 0.1% of all computers worldwide were part of a botnet.

Mainly through Eurostat, there are also statistics available about the perceived trust in ICT. As examples we mention (for end users in the European Union in 2010): 25% is very worried about viruses in the Internet, 16% does not purchase goods or services through the Internet, for security reasons and 15% does not use on-line banking for the same reason. Unfortunately, for the end users that refrain from purchasing goods and using on-line banking, the surveys do not give insight in what this lack of trust is built upon. In summary, our conceptual framework of trust in ICT is hard to apply in practice for several reasons.

In the next section, we propose a way of assessing trust in ICT in an alternative way, i.e. by using sentiment mining of social media. One of the advantages of such a dynamic, non-survey-based method is that it can be applied in near real-time, while the framework above heavily relies on data that only becomes available on an annual or (bi-)quarterly basis.

3 A Social Media-based Model of Trust

In this section a formal model of trust, measured in social media on the basis of sentiment, is described. With tools based on such models, one can inspect the current level of trust in a certain topic and for a certain (online) community, or assess the impact of certain simulated interventions on trust scores. While a social media-based model of trust may be only partially representing the trust of a larger, *offline* community, information from social media can be relevant here: social media are commonly accessible to large proportions of the community, have large degrees of participation, and are quite dynamic, allowing for the monitoring of rapid changes in online expressed mood. Sentiment analysis is one of the themes of text analytics, and it is highly ranked on the research agenda of academia and industry, as exemplified by thriv-

⁴ For instance the annual reports of the Organisation for Economic Co-operation and Development (OECD; <http://www.oecd.org/>).

ing, industry-sponsored scientific conferences such as ICWSM⁵. Roughly speaking, sentiment analysis attempts to detect opinions and subjective utterances, labelling subjective, opinion-bearing utterances with polarity scores such as 'negative', 'positive', or points on a metric scale. For an extensive overview, see Pang and Lee ([5]).

In our trust model, we attempt to model the trust of an audience in a certain ICT service as a function of the exposure to social media (i.c. Twitter messages about the service), in combination with information about the authority of the source of the messages. Our hypothesis is that there is a correlation between trust of a population and the exposure of that population to highly polarized (overtly positive or negative) information expressed by people with high authority. In this work, we simplistically equate 'authority' with a large number of followers but our approach is open to more advanced measures of authority. Figure 1 shows a formal model for combining sentiment information with authority estimates into a trust score, where the top "trust value" is calculated by $T(t)$, the top "input" of the model combined with its "sentiment" is calculated by $\sigma(t)$, and the bottom "input" combined with its "sentiment" is calculated by $\beta(t)$.

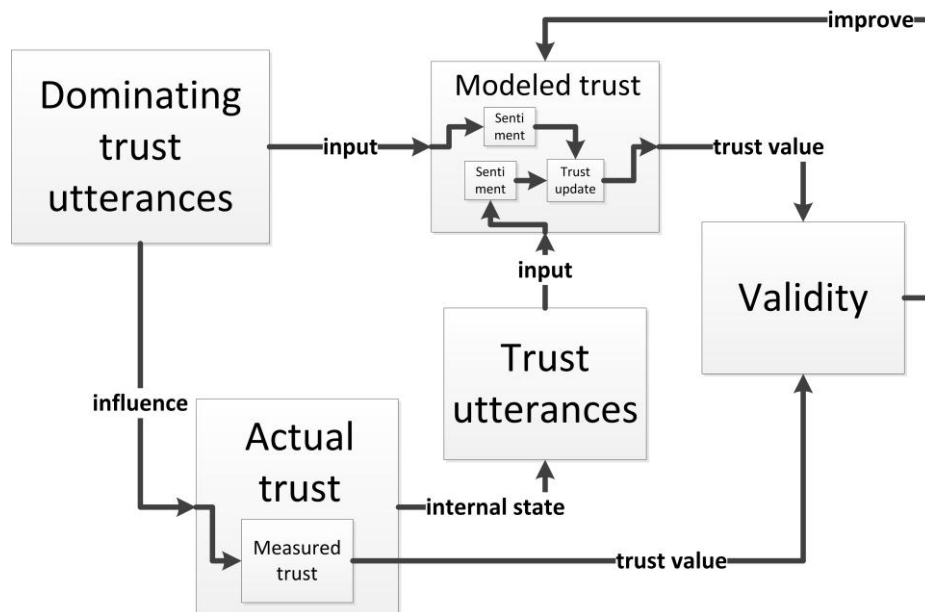


Figure 1. Formal model of trust.

The trust value (or the degree of trust) is a number between 0 and 1 that is an estimate of the tendency of the general public to accept and use a certain ICT service. This estimate is important when one wants to predict future use of the service given the current situation (monitoring and extrapolation), or when one is interested in a possible future introduction of a service (simulation).

⁵ The International Conference of Weblogs and Social Media, <http://www.icwsm.org>

3.1 Degree of Trust

The following basic formula is used in order to compute the degree of trust at a certain time point t :

$$T(t) = w_T \cdot \sigma(t) + (1 - w_T) \cdot \beta(t)$$

where $T(t)$ is the degree of trust, $\sigma(t)$ the degree of situational trust, and $\beta(t)$ the degree of behavioural trust, at time point t . The weight w_T regulates the balance between the situational and behavioural trust. We will explain these notions now in further detail.

3.2 Situational Trust

The degree of situational trust is the degree of trust that is deducible from the situation. The situation is determined by the opinions of, and the resulting trust utterances by, people with higher authority. The opinion of people with higher authority is more often taken over by the general public than of people with lower authority. Therefore utterances are weighed according to the degree of authority of the utterances' owners.. The above can be formalized as follows:

$$\sigma(t) = \begin{cases} \frac{1}{n(t)} \cdot \left(\sum_{i=1}^{n(t)} \left((1 - w_a) \cdot \lambda_\sigma + w_a \cdot (1 - a_i(t)) \right) \cdot \sigma(t-1) + \right) & \text{if } n(t) > 0 \\ \left((1 - w_a) \cdot \lambda_\sigma + w_a \right) \cdot \sigma(t-1) + \left(1 - ((1 - w_a) \cdot \lambda_\sigma + w_a) \right) \cdot e_d & \text{if } n(t) = 0 \end{cases}$$

where $n(t)$ is the total number of trust utterances from time point $t-1$ until t , and the decay λ_σ regulates the amount of past situational trust that is included when calculating the current situational trust. The weight w_a regulates the effect of momentary authority. Furthermore, $e_i(t)$ and $a_i(t)$ are the momentary sentiment and authority of utterance i from time point $t-1$ until t . And finally, e_d is the default momentary sentiment when no utterances are within the time interval $t-1$ and t (i.e., $n(t) = 0$), which should be put to 0.5 when regression zero sentiment is assumed. The second case of the above equation is equal to the first case when $n(t) = 1$, $e_i(t) = e_d$, and $a_i(t) = 0$ is taken in the first case.

3.3 Behavioral Trust

The degree of behavioural trust is the degree of trust that is deducible from behavior. Behavior is observable via trust utterances of people in general. This can be formalized as follows:

$$\beta(t) = \begin{cases} \lambda_\beta \cdot \beta(t-1) + (1 - \lambda_\beta) \cdot \frac{1}{n(t)} \cdot \sum_{i=1}^{n(t)} e_i(t) & \text{if } n(t) > 0 \\ \lambda_\beta \cdot \beta(t-1) + (1 - \lambda_\beta) \cdot e_d & \text{if } n(t) = 0 \end{cases}$$

where decay λ_β regulates the amount of past behavioral trust that is included when calculating the current behavioral trust. For the initial values of $\sigma(t)$ and $\beta(t)$ the following holds:

$$\begin{aligned} \sigma(0) &= \sigma_0 \\ \beta(0) &= \beta_0 \end{aligned}$$

where initial values σ_0 and β_0 are equal to 0.5 when indifferent trust (neither high nor low trust) is assumed in the case when there are no trust utterances yet.

3.4 Momentary Sentiment and Authority

The momentary sentiment and authority are defined as follows:

$$\begin{aligned} e_i(t) &= w_e + (1 - w_e) \cdot \text{sentiment}_i(t) \\ a_i(t) &= \frac{\text{followers}_i(t)}{\frac{1}{t_{end}} \sum_{t_j=1}^{t_{end}} \left(\frac{1}{n(t_j)} \sum_{p=1}^{n(t_j)} \text{followers}_p(t_j) \right)} \end{aligned}$$

where weight w_e regulates the effect of the momentary sentiment. It furthermore holds that $\text{sentiment}_i(t) \in \{0,1\}$ and $\text{followers}_i(t) \in \mathbb{N}$, for all utterances i and time points t .

In the next section we will apply this model to a use case: the trust in Apple's newly launched cloud service *iCloud*, and describe our experiments.

4 Model Tuning and Validation

This section describes the efforts to tune and validate the model of trust, specifically for trust in Apple's *iCloud*, our use case.

4.1 Gathering Model Input Data

In order to capture a sufficient amount of data as input for the model described in the previous section, we implemented a procedure for harvesting the streaming timeline of Twitter, on the basis of search keys consisting of tags and keywords. Our data collection was created during a time span of three months (from August 25, 2011 to October 13, 2011), based on the keyword *icloud* and consists of 193,469 tweets of 195,556 different users. For every tweet containing the keyword *icloud*, the following information was stored in a PostgreSQL database:

- the screen name of the sender
- the personal name of the sender
- the creation date of the tweet
- the tweet text
- the number of followers of the sender at the time the tweet was published
- the tags added by the sender to the tweet
- the number of friends the sender had at the time of the tweet
- the geo-location from which the tweet was sent out (if present)
- the number of retweets of the tweet (measured by the Twitter API in a certain interval starting with the broadcast time of the tweet).

In addition, we trained and applied three *text classifiers* to the tweet text:

- a subjectivity classifier that detects whether the tweet contains an opinion or, on the contrary, consists of factual information only;
- a binary sentiment polarity classifier that estimates the polarity of the sentiment (if present) in the tweet text: either positive or negative.
- a binary topic classifier that estimates whether the tweet is about cloud computing or not.

These classifiers consist of *support vector machines* applied to bag-of-word representations (frequency counts, unnormalized) with *radial basis function (RBF) kernels*. Their output was stored in the database as well.

4.2 Gathering Validation Data

As shown in figure 1 the data for validating the trust model are gathered by sampling from the people whose trust is being estimated by the trust model. This was done by questionnaires. For two timestamped events, the official introduction of iCloud, and the death of Steve Jobs, the first twenty tweets were selected, both directly preceding and following these events. These twitter stimuli were sent out through an automated mail procedure to a total number of 61 participants, with two tasks:

- to annotate the sentiment polarity in these tweets on a three-point scale (negative, neutral, positive). This entails a form of affective exposure to the content of these tweets that - to a very limited extent- mimics real-life exposure to this type of information. In addition, the number of followers for every tweet sender at the time of publishing the tweet was listed, with the intent of illustrating somehow the 'authority' of the tweet sender.
- to answer a number of questions in a questionnaire, implemented as a separate (and personal) web page.

So, in total, every participant received 8 questionnaires. In the questionnaires, five questions were posed:

1. Do you know what Apple's iCloud service is? (1=not at all, 5=perfectly)
2. When did you start using iCloud? (1=never used it, 5=from the beginning)
3. Do you think personal files are safe in iCloud? (1=not at all, 5=very safe)
4. Are you enthusiastic about iCloud as a service? (1=not at all, 5=very enthusiastic)
5. How certain are you about the answers above? (1=not at all, 5=very certain)

In order to synthesize a trust value from the answers, we applied the following heuristic. First of all, when the answer to the first or fifth question was 1, the returned answers were excluded. In all other cases, the answers to questions three and four were used to produce a trust value, with the answer to question four serving as a weight factor to the answer of question three.

4.3 Parameter Tuning

The tuning of the five parameters of the trust model was done by means of an exhaustive search with a granularity of 26 (a step size of 0.04 on an interval of 0 to 1 for each of the five parameters). The problem space therefore consisted of $26^5 = 11,881,376$ different parameter settings which needed to be compared to each other. Given that the calculation of the validity of the model for one parameter setting on a regular PC costs roughly one second of time, this would result in a tuning period of 138 days. Instead of going for a heuristic search algorithm to search through the problem space, we chose to implement the problem as a parallel procedure. For larger parameter spaces, this would be a less suitable solution and one could opt for heuristic search algorithms. For instance, an alternative way of parameter estimation could consist of either a stochastic, sampling-based approach (such as a genetic algorithm), or a learning approach with gradient descent and back propagation. For the latter, given a (differentiable) trust function $T(t)$, denoted with τ , we can implement a gradient descent method as follows:

$$\omega^{t+1} = \omega^t - \eta \frac{\partial \tau}{\partial \omega^t}$$

with η a parameter controlling the *learning rate* of the algorithm and

$$\frac{\partial \tau}{\partial \omega^t} = \frac{\partial \tau(\omega)}{\partial \omega}$$

with ω^t the five-dimensional weight vector for the parameters of our model at time t . This implements parameter estimation as a form of supervised learning (regression).

The current parallel search was run on a computer cluster using 2048 different cores in total, which would theoretically reduce the tuning period to two hours. But due to other jobs running on the cluster, the task scheduler allowed the script to run in roughly ten hours, which is still a considerable reduction of computing time.

As a first attempt to see if validation data consisting of four data sets of participants estimating their trust in iCloud is sufficient for tuning the trust model, we first tested whether the found optimal parameter settings were satisfactory using artificially generated validation data. These validation data were generated by normally randomizing the output of the model using predefined parameter settings within four different short time intervals. The predefined parameter settings were $(\omega_a, \omega_e, \omega_T, \lambda_\sigma, \lambda_\beta)_p = (.9, .95, .5, .999, .999)$. The normal randomization was carried out ten times (for ten non-existing participants) using the output of the model as mean and 0.01 as standard deviation. The results of the parameter tuning are shown in figure 2.

The used validation data are indicated by black squares. The found parameters after tuning are were $(\omega_a, \omega_e, \omega_T, \lambda_\sigma, \lambda_\beta)_t = (.44, .96, .12, 1, .84)$. The mean absolute difference between the tuned output of the model (blue line) and the validation data (squares) is 0.00859. The absolute difference between the tuned (dark line) and the artificially generated output of the model (light grey line) is 0.0114. There is no baseline to compare this result against, but it seems satisfactory enough for using the type of validation data that was proposed in the previous section to get a good result in the case when the validation data is not artificially generated (and no artificially generated output of the model is present (i.e., the light grey line in figure 2).

In order to prevent overfitting on the validation data, 2-fold cross validation has been used. This means that we have randomly split the validation data into two sets and the eventual found optimal model parameters are the average between the two found solutions after tuning on the basis of the two different sets. The estimated error of the model for unseen data is calculated by averaging between the mean squared error (MSE) of the model given the found parameters tuning on the first set and testing on the other and the MSE of the model given the found parameters the other way around.




Figure 2. Trust model output after tuning (line) based on normally randomized validation data (squares) using the output of the trust model with predefined parameter settings.

4.4 Results

The results are shown in figure 3. The gathered validation data are depicted as numbers, where each digit is the number of participants indicating his trust in iCloud. Based on this gathered validation data, parameter tuning led to the trust model output as indicated by the line. The shown means per questionnaire (stars) are just for illustration purposes, since tuning was done using each individual data point. The found parameters after 2-fold cross validation are $(\omega_a, \omega_e, \omega_T, \lambda_\sigma, \lambda_\beta)_t = (0.92, 1, 0.3, 0.62, 0.75)$. The average of the MSEs for these parameters is 0.0439.




Figure 3. Trust model output after 2-fold cross validation (line) based on the gathered validation data. The black numbers are the number of participants that indicated the respective trust values in iCloud. The stars are the means of those trust values for each of the four questionnaires.

5 Visualization

Our proposed model of trust was used as a basis for a visualization tool aimed at supporting professionals in their effort to monitor trust in ICT services or to simulate certain implications given different interventions. Below a first version of this tool is described.

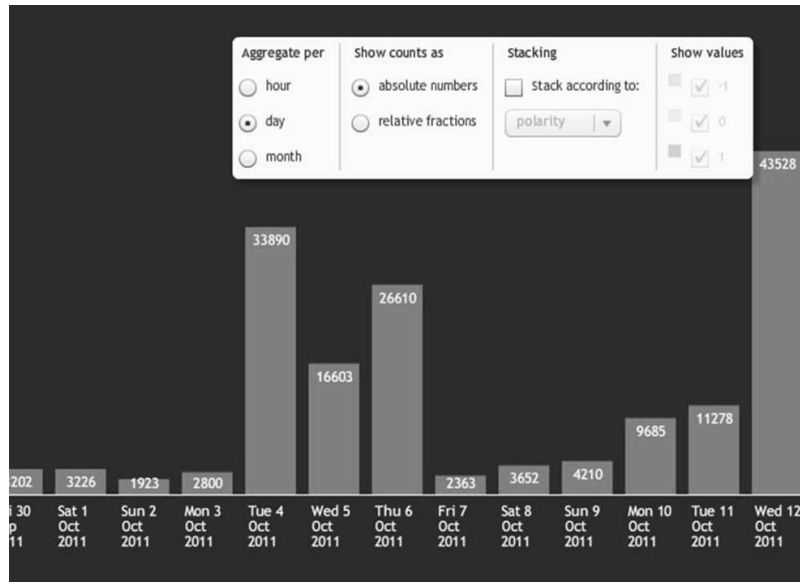


Figure 4. Number of tweets about iCloud per day. On October 4th iCloud is announced; on October 12th iCloud is launched.

The visualization tool serves two purposes. First, it provides insight into the data that we are scraping from Twitter, through simple statistics like the number of tweets over time and sentiment scores. The tool allows to inspect absolute numbers (see figure 4), to break down the numbers according to sentiment (see figure 5), or to see relative fractions (see figure 6). All this can be shown aggregated per hour, day or month.

The other purpose is to show the results of the trust calculations from the model. The tool shows the trust in a certain topic (in our case: iCloud) over time. This can be done in near real time: the tool not only shows the history of trust, but also what the current trust is. In that sense, it has monitor functionality. The bright line in figure 6 illustrates how the trust value over time, as predicted by our model, could be combined with the distribution of sentiment over time. The starting point for the monitor tool is the data scraped from Twitter, which has been automatically annotated with sentiment and authority values. This means that each tweet has been automatically annotated with a value for its sentiment (positive, neutral, or negative) and with a value for its authority (number of followers divided by a certain predetermined maximum). A separate script counts all tweets, aggregating them per hour, per day, and per month. It counts, for instance, how many tweets were about iCloud at a specific hour, and how many of those tweets contained positive, neutral or negative sentiment. The resulting table can directly be visualized by a simple bar chart. The tool is interactive; each time the user selects or deselects an option from the interaction pane (the white pane in the figures), the corresponding data is read from the database and the view is refreshed.



Figure 5. Number of tweets about iCloud per day, broken down in positive (light colored) and negative (dark colored) tweets. Notice the increase in negative tweets right after the announcement of Jobs' death on October 6th.

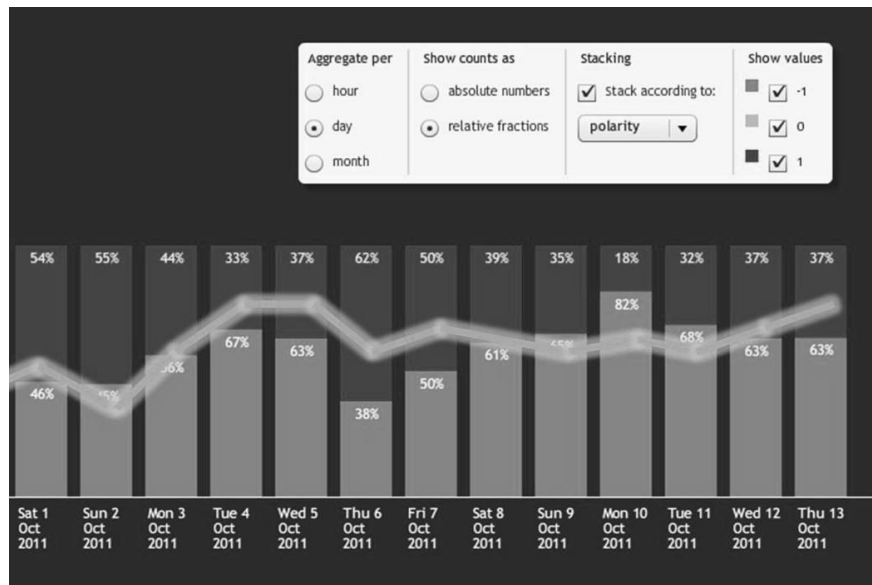


Figure 6. Percentage of positive tweets and negative tweets about iCloud per day. The bright line is a fictive plot of the calculated trust (not based on the actual model).

6 Discussion and Conclusions

We presented a social media-based, parameterized model of trust based on the notions of sentiment and authority, as an alternative for report-based, static models for assessing trust. We demonstrated that this model can be fitted to feedback data gathered from participants in an experiment. In addition, we tested the prognostic capabilities of the model when applied to evaluation data. We have also presented an application of the model in a visualization tool aimed at supporting professionals in their effort to monitor trust or to simulate certain implications of interventions.

From a procedural point of view, we have devised and implemented a fully working automated system that:

- gathers (scrapes) data from social media, and analyzes these data (sentiment analysis, topic classification), storing data and analysis results in a database;
- polls users for feedback on stimuli retrieved from the database;
- fits a model of trust based on measured sentiment, observed authority and self-reported trust;
- offers visual inspection possibilities, combining views on the stored data with model predictions.

Several aspects of our approach are open to improvement or further exploration. The current experimental setup suffers from several drawbacks. Specifically, a ‘tabula rasa’ exposure of participants was not guaranteed. The tweets presented to the test persons were historical, and chances are that they had already been observed by the participants. Furthermore, the current setup does not take into account exogenous (hidden, external) variables to which the participants have been exposed, such as other sources of information (newspapers, blogs, TV), or forms of social influence. Ideally, we would have test persons being exposed to a fixed number of controllable information sources and social influences. While such a laboratory setup is hard to envisage, monitoring devices handed out to test persons, for instance with deep packet inspection and key loggers, would probably yield more reliable information. This would open up the possibility of estimation of a model that weights these several sources of information for the prediction of trust.

Our parameter fitting approach, while powerful, is computationally expensive and required the use of a large computer cluster. Well-known and more practical solutions to finding accurate and robust parameter settings are available, for instance sampling methods (such as genetic algorithms) or learning approaches, such as gradient descent methods. We also suggest the use of many more questionnaires for different time intervals that will also improve the parameter fitting approach. All this emphasizes the need for other parameter fitting methods, which will be addressed in a follow-up of this research.

References

1. Corbitt, B.J., Thanasankit, T., Yi, H.: Trust and e-commerce: a study of consumer perceptions. *Electronic Commerce Research and Applications* pp. 203–215 (2003).
2. Keymolen, E., van den Berg, B., Prins, C., Frissen, V.: *Vertrouwen in hybride ketens. Onderzoeksrapport in het kader van Alliantie Vitaal Bestuur*, Den Haag (2010)
3. Kim, D.J., Song, Y.I., Braynov, S.B., Rao, H.R.: A multidimensional trust formation model in b-to-c e-commerce: a conceptual framework and content analyses of academia/practitioner perspectives. *Decis.Support Syst.* 40(2), 143–165 (Aug 2005)
4. Kim, Y.A., Song, H.S.: Strategies for predicting local trust based on trust propagation in social networks. *Knowl.-Based Syst.* pp. 1360–1371 (2011)
5. Pang, B., Lee, L.: Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval* 2(1-2), 1–135 (2008)
6. Utz, S.: Rebuilding trust after negative feedback: the role of communication. In K. Cook, C. Snyders, V. Buskens and C. Cheshire (eds.); *eTrust: Forming relationships in the online world* (pp. 215-237). New York: Russell Sage Foundation (2009)