Beyond Access Control:
Managing Online Privacy via Exposure

Mainack Mondal, Peter Druschel, Krishna P. Gummadi
MPI-SWS
{mainack, druschel, gummadi}@mpi-sws.org

Alan Mislove
Northeastern University
amislove@ccs.neu.edu

Abstract—**We posit that access control, the dominant model for modeling and managing privacy in today's online world, is fundamentally inadequate. First, with access control, users must *a priori* specify precisely who can or cannot access information by enumerating users, groups, or roles—a task that is difficult to get right. Second, access control fails to separate who *can* access information from who actually *does*, because it ignores the difficulty of *finding* information. Third, access control does not capture if and how a person who has access to some information redistributes that information. Fourth, access control fails to account for information that can be inferred from other, public information. We present *exposure* as an alternate model for information privacy; exposure captures the set of people expected to learn an item of information eventually. We believe the model takes an important step towards enabling users to model and control their privacy effectively.**

I. MOTIVATION

Privacy is traditionally defined as "the ability for people to determine for themselves when, how, and to what extent information about them is communicated to others" [23]. In computing systems, privacy has typically been accomplished via *access control*, which requires enumerating the users, groups, or roles who are or are not able to access information.

The popularity of online social media sites have led to a renewed discussion about whether access control is a satisfactory model for user privacy. These sites now mediate the sharing of personal information, photos, status updates, and contacts of billions of users around the world; some sites even serve as the de-facto Internet portal for a significant fraction of the world's population. In this paper, we focus on the privacy controls these sites provide users to manage access to their content by other users; other works [2], [11] focus on the orthogonal concern of protecting users' content from the site operator.

A. Access control is insufficient

Online social media sites provide privacy controls based on *access control* and require users to allow or deny access to their content by specific users or groups. Recently, there have

been a number of incidents that call into question whether access control is the right mechanism with which to implement privacy. We list a few here; this list is by no means exhaustive:

1. When Facebook introduced the *News Feed*—a feature that automatically presents updates from friends when a user logs in, as opposed to requiring the user to visit the friends' pages—users objected strongly and accused Facebook of privacy violations. Strictly speaking, News Feed did not change the access control policy; all users who viewed content through the News Feed had access to the content before. However the change from a pull mechanism to a push mechanism resulted in users feeling that their privacy had been violated.

2. There was a similar outcry of privacy violations when Facebook introduced *Timeline*, a feature that indexes a user's content by date of upload and allows users to quickly browse content by upload date. As with the News Feed, Timeline did not change the access control policy of any content. Instead, Timeline made accessing old (and potentially embarrassing) content significantly easier.

3. Google's *Street View* project—providing photos of houses and other property taken from public street—has also been accused of violating the privacy of users. In the U.S., there is no legal expectation of privacy on a public street (i.e., Street View photos can legally be posted publicly), but many users feel uncomfortable that Street View has made information easily and widely accessible that previously was visible only to those physically present.

4. Data aggregator *Spokeo* links together public information from different services (e.g., government databases, sites like LinkedIn, etc). While each individual piece of content that Spokeo aggregates is publicly available, users have complained that their privacy is violated when this information is linked together. For example, Spokeo cross-references users' addresses with property records, allowing others to quickly estimate someone's wealth using public information.

While perceptions of privacy and what constitutes a privacy violation are subjective, most people would likely agree that each of the incidents above affect someone's privacy. However, the take-away from all of these incidents and others is that none of them involved a violation of *access control*. As a result, we argue that privacy is not adequately captured by access control alone, and the research community should re-consider how to model and reason about user privacy.

B. Goal: A more inclusive privacy model

In this paper, we carefully reconsider the issue of privacy in the age of the web and social media. We propose a model of privacy based on *exposure*, where the exposure of a piece of information is defined as the set of principals (people) who are expected to eventually know it. Users implicitly reason about the exposure of various pieces of information; a violation of exposure occurs when the set of users who become aware of a piece of information is much different from what the user expected. In fact, recent work [3] by Facebook researchers has shown that such exposure violations are commonplace; e.g., many users significantly underestimate the number of users who actually view their content.

For example, consider the case of a user's public Facebook page being linked to from a high-profile web site such as the New York Times. Strictly speaking, there is no access control violation; the user's profile was previously publicly visible. However, a significant change of exposure occurrs as the set of people expected to see the page increases from a small set of users likely to visit the user's page to the much larger set of New York Times readers. We argue that exposure naturally captures the privacy change of such an incident, and makes clear why access control alone is insufficient.

We discuss mechanisms that could increase user's control over privacy by moving from *access control* towards *exposure control*, and describe how these mechanisms could be built into today's content sharing systems. Overall, our goal is not to promote concrete proposals, but rather to initiate a discussion of new mechanisms for privacy control.

The remainder of this paper is organized as follows. In Section II, we provide a more formal definition of exposure and discuss and compare exposure control with more traditional access control. In Section III, we describe approaches that could provide users with improved privacy via exposure controls. Section IV explores the feasibility of using exposure control to manage privacy. In Section V, we detail related work and we conclude in Section VI.

II. Defining Exposure

In this section, we propose a simple model of exposure.

Let I be an item of information (e.g., Alice's date of birth is Jan 1, 1980). Informally, I's exposure is the set of principals we expect to *eventually* learn I. The exposure set includes principals who learn I directly from Alice or indirectly from a third person with knowledge of I, and those who infer I from other knowledge available to them.

More precisely, we define the *prominence* $P_I(t)$ as the set of principals who are aware of I at time t.[1] I's exposure $E_I = \lim_{t \to \infty} P_I(t)$. Note that E_I is always finite, because the set of principals (i.e., the world's population) is finite. However, the exposure of most information items, even if they are publicly accessible, is much smaller than the world's population, because they are of interest to only a small community.

[1]Prominence is assumed to be a monotonically non-decreasing function of time. That is, we ignore that people forget or misplace information.

Normally, $P_I(t)$ is unknown for $t > currentTime$. Future values of $P_I(t)$ must be estimated using a probabilistic model, which captures how information spreads among principals; the exposure is given by the expected steady-state prominence predicted by the model. An example of such a model is discussed in Section III.

A. Aspects of exposure

The exposure of an item I is influenced by two factors:

1. The set of principals N_I that meet the preconditions required to learn I. (Preconditions include the expertise, access to the tools, and knowledge of initial leads required to discover or infer I.)

2. The subset of principals in N_I that is sufficiently motivated to actually learn I.

For instance, if learning I requires correlating several pieces of related information, traveling to a particular location or performing a measurement, then it is likely to be learned only by principals with the necessary resources and a strong interest. If, on the other hand, I is online and indexed by a search engine, then it can be learned by anyone with access to the Internet, the expertise to use a search engine, knowledge of appropriate keywords, and sufficient interest to actively issue a search query. Lastly, if I is posted on the front page of the New York Times, then all principals who visit the site on a daily basis will likely learn I serendipitously, even if they are only mildly interested.

The exposure of an item of information may change over time. For instance, when a little-known website is listed on Slashdot, the set of users likely to discover the information contained in it increases dramatically and unexpectedly. Such events cause a discontinuity in the prominence function $P_I(t)$, and thus a potential change of exposure.

B. Comparison with access control

Figure 1 contrasts access control and exposure using a Venn diagram. In the access control model, the set of principals is partitioned into those who are able to access I and those who are not. Access control does not capture how many principals with access permissions actually access the information; nor does it account for principals without permission who nevertheless learn the information, either by inferring it from other information they can access, or from another principal with access.

Exposure captures which principals are likely to actually learn the information, which is more directly relevant to privacy. To illustrate this point, let us reconsider the cases of perceived privacy violations we discussed in Section I-A.

Exposure captures the changes caused by the introduction of the Facebook News Feed: Prior to its introduction, the exposure of an item I on Alice's profile was the number of unique users who visit Alice's Facebook page during I's lifetime, which could be much smaller than the set of users N_I with permission to access I, particularly if Alice chose to make I public. With News Feed, on the other hand, I's exposure includes all of Alice's friends plus any user in N_I

(a) **Access Control:** A_I both over- and underapproximates $P_I(t)$.

(b) **Exposure:** E_I is the set expected to learn I eventually.

Fig. 1. **Access control and exposure of an information item I shown as a Venn diagram.**

who Facebook deems potentially interested in I. I is pushed to these users, who will learn I serendipitously the next time they log into Facebook. Similarly, the introduction of Facebook's Timeline pushes selected information about a person's history to a set principles in N_I that Facebook deems interested. Previously, finding such information would have required a user in N_I to visit Alice's profile and scroll potentially deep down into her historic News Feed.

Google Street View has made available online, in an aggregated and searchable fashion, public information that was previously available only to principals who were physically present at a particular geographic location. Exposure reflects this fact. Spokeo aggregates people's personal information like name, address, data of birth, income, property value and family tree, which is available from different online sources, and makes it available and searchable under the person's name and place of residence. By making it far easier to learn this information, its exposure is increased.

C. Privacy violations covered by Exposure

Whether the release of information about a person is considered a privacy violation by that person is subjective and deeply rooted in the person's culture, history, situation, the nature of the information, and the specific set of people who learned the information.

In general, however, a person is more likely to feel violated if (s)he is surprised by the fact that certain people have learned the information. There are two relevant cases. A person typically has some expectation about (a) the set of people who know or are likely to learn an item of information, and (b) a specific set of people they expect should not and will not learn the information. A person tends to feel their privacy is violated if the actual exposure of an item includes many more people than the expected exposure [5], or if people in the second set learn the information.

An example of the former case would be if Alice finds out that a picture showing her dancing wildly at a party has been seen by all of her friends, family and colleagues, when she expected that it would become known only to the people who had attended the party. An example of the latter case would be if Alice's work colleagues find out that she is gay (even

though she shares this information freely with her friends and family, and she makes no attempt to hide it from people she encounters in her life outside of work).

Exposure captures these concerns because it reflects the set of people likely to find out the information.

III. Managing Privacy via Exposure

In this section, we discuss how the concept of exposure can be leveraged in practical systems to enable users to manage their online privacy better.
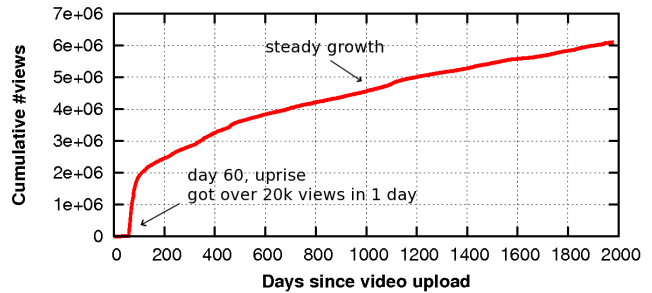
There are two important notes to make before we discuss exposure control. First, our goal is to propose a general methodology that could be broadly applied to control exposure of users' information in a variety of online systems. Thus, our discussion is not specific to any one system. Second, there are several interesting research questions that remain to be investigated through a concrete implementation and deployment of our proposal. However, such an deployment-based study is beyond the scope of this paper and we view our proposal as a call for further research in this direction.

A. Predicting exposure

Modeling and predicting the growth in popularity of information on social Web sites like Facebook photos, Twitter posts or YouTube videos [12], [15], [20] has received significant research attention recently. These studies use empirical data of how information became popular in the past to build models for information propagation that can predict the future popularity of similar information. Popularity growth models are relevant because they can predict he cardinality of exposure. The prediction models vary from very simple models that extrapolate from the historical growth in popularity of a single piece of information; to more complex models that take into account various factors including attributes of the information (e.g., quality, type, and length of a video), historical data about the spread of other similar pieces of information, and the effectiveness of different information dissemination channels (e.g., social links between users or personalized recommendations or search results). More sophisticated models with higher prediction accuracy have been developed over time. While a detailed discussion of these models is beyond the scope of this

3
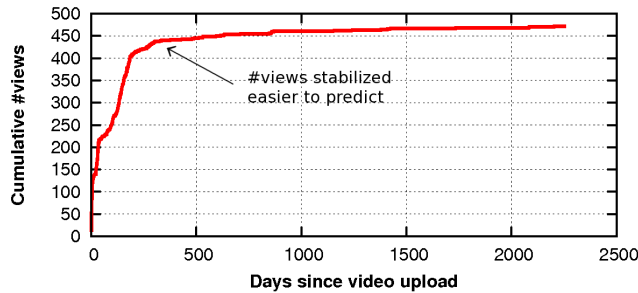
(a) **Niche video with total 471 views**   (b) **Popular video with total 6,101,294 views**

Fig. 2. **Popularity growth patterns of two Youtube videos. The popularity of the niche video stabilize and becomes predictable within a year, but the popularity of the popular video exhibits an uprise and shows steady growth.**

paper, we make two general observations that are relevant to our discussion:

1. Prediction accuracies are higher for less popular (niche) information than they are for more popular information. For example, it is easier to predict the future popularity of YouTube videos with a few hundred views after 1 year than those with few million views after 1 month [20]. As shown in Figure 2(a), the dissemination of niche videos tends to stabilize to a predictable rate sooner than those of popular videos.

2. Most models cannot anticipate the occasional sharp, disruptive jumps in popularity that arise due to unanticipated events, such as when a piece of information goes viral or when it is featured on a prominent site [19]. Figure 2(b) shows an example YouTube video whose number of views experienced a sharp jump on day 60 due to coverage on popular media and blogs.

B. Making the predictions transparent

We argue that system operators (e.g., Facebook or YouTube administrators) should make both past popularity data and predictions for the (cardinality of) exposure of users' information transparent to the user. Currently, some systems provide users with a limited view of the popularity of the information they upload. For example, Facebook and YouTube allows users to check the number of views or "Likes" for their posts. However, no site today explicitly provides estimates of future popularity of a piece of information. For example, neither Facebook nor YouTube offer guidance on how many and which people might see a photo over the next week. We see this as a missed opportunity because (i) the site operators are often in the best position to make such predictions as they have the best access to all the empirical data on how information disseminates through their sites and (ii) such estimates would enable users to (re)calibrate their expectations for the future exposure of their information and check for potential privacy violations.

When providing estimates for the exposure of a piece of information, it would also be useful to estimate the likely exposure through different dissemination channels separately. For example, Facebook might choose to provide estimates of views a user's photo might achieve through updates on personalized news feeds versus graph search versus profile browsing [9]. Doing so would enable users to understand the predicted exposure of their content, and to modify the sharing settings if it does not match their expectations.

C. Controlling exposure

Providing users with more accurate exposure estimates for their information does not by itself eliminate the risk of privacy violations. System designs need to enable users to tune (i.e., increase or decrease) the exposure to the values they desire. Further, systems should be designed to also prevent the actual prominence from deviating significantly from the predictions (after they are tuned to desired values). Below we propose mechanisms to achieve the above two goals.

1) Tuning exposure: When a user finds that the predicted exposure of her information is different from what she desires, there need to be mechanisms that would allow her to tune the exposure. A user could do this in several ways: first, she could enable or disable one or more dissemination channels. For example, on Facebook, one could opt-in/-out of being part of "directory or graph search." Such opt-in/-outs from one or more dissemination channels could help users manipulate their exposure to desired levels.

Second, users can resort to more expansive or restrictive access controls (i.e., who is allowed to see or not see a piece of information) to change the exposure of a piece of information. For example, to increase exposure of a piece of information originally shared with her 1-hop friends, a Facebook user might choose to make it available to 2-hop friends (i.e., friends of friends). To decrease exposure, the user might choose to make it available to only a subset of 1-hop friends (e.g., friends with whom the user shares a common university affiliation). By changing access controls, the user can expand or contract the list of potential viewers and thereby, change the list of predicted viewers. Thus, we envision access control being used in conjunction with exposure control to more closely match the user's expectations.

We propose to design systems that would let users observe the effects of such changes on their information's exposure.

2) Limiting divergence from predictions: Even after a user tunes the exposure to match her expectations, unanticipated events (e.g., the information goes viral and is featured on the front page of a popular site) might cause the actual exposure to deviate significantly from the predictions and consequently,

the desired exposure. As mentioned earlier in Section III-A, no model can accurately predict such occasional disruptive changes to the prominence of a piece of information.

To contain privacy violations in such scenarios, we propose that systems adopt *tripwires* that automatically make a piece of information inaccessible whenever the actual exposure of a piece of information deviates significantly from the predicted exposure and notify the user of the unanticipated divergence. Upon notification, users can explicitly choose to keep the information inaccessible or re-enable access to the information (and readjust the tripwires). Alternately, systems can allow users to specify tripwires that upper-bound the views (e.g., no more than 10 views per day or 50 views in total) to a piece of information.

We believe that tripwire mechanisms can be easily enabled in current systems like YouTube or Facebook. In fact, YouTube already allows users to limit the total number of views to their videos to a preset value of 50 (effectively providing a limited form of exposure control).

IV. FEASIBILITY OF MANAGING PRIVACY VIA EXPOSURE

There are a number of interesting research questions that need to be studied through a concrete implementation and deployment of our proposal of managing privacy via exposure. Such a detailed study is the subject of our future work and beyond the scope of this paper. However, we discuss two important concerns that one might have about the feasibility of a practical deployment of our proposal: (i) our ability to accurately predict the future exposure of a piece of information and (ii) the overheads associated with fine-tuning exposure.

A. Accuracy of exposure predictions

There is ongoing research on predicting information propagation and dissemination in online systems. These works leverage the ability of sites to gather and analyze detailed historical information about the exposure of billions of pieces of information posted by hundreds of millions of users to make accurate predictions. For example, in a recent study [3], Facebook researchers were able to predict the audience size of a new post by a user within an 8% error margin, using data such as the number of friends and the likes and comments the post received. To illustrate the ability to make such predictions in different scenarios, we conducted a small-scale study in three different real-world scenarios, each using different access control policies—(i) public posts on sites like YouTube, (ii) posts limited to members of Facebook groups and (iii) personal posts limited to one's Facebook friends.

In each of these scenarios, we used linear regression [17] for predicting future popularity using past information and then measure the *relative error* of our prediction. Relative error is defined as $\left|1 - \frac{\text{Predicted\_value}}{\text{Actual\_value}}\right|$. The lower the relative error, the more accurate the predictions. In each of the scenarios, we show how system operators can use different types of past information to predict the future popularity of a content with low error, i.e, high accuracy.

Scenario 1: Predicting future popularity of public YouTube posts. We analyzed the past number of daily view for publicly posted YouTube videos to predict their future popularity.


Fig. 3. **Prediction of number of people accessing a content in three scenarios: scenarios 1, 2 and 3 corresponds to prediction of public YouTube video views, Facebook group post views and photo likes on personal Facebook photos respectively. The relative error of prediction decreased from scenario 1 to 3.**

Specifically, we collected data about the number of daily views that 50 randomly chosen YouTube videos obtained in their first 6 months. We used this historical data to predict how many views the videos will get the immediate next day.

Scenario 2: Predicting future popularity for new posts in Facebook groups. Next, we tried to predict the number of views for a new post in an open Facebook group [8] using the popularity of older posts in the same group [7]. We collected popularity data for the posts of 50 open Facebook groups and predicted the number of views for the latest post in each of these groups using the popularity of older posts.

Scenario 3: Predicting future popularity of personal posts limited to one's Facebook friends. We collected data about the pictures shared by 50 Facebook users (randomly selected users of a Facebook application [10] created by authors) with their friends along with the number of "Likes" on those pictures. Using this data, we predicted the number of "Likes" a user would get on a future photo shared with the same access control policy. We performed this prediction for the latest photo of each user.

We present the distribution of relative errors in predictions for each of these scenarios in Figure 3. Note that, intuitively, the set of people who can learn about the content decreases from scenario 1 to scenario 3. Figure 3 shows that consequently the relative error decreases from scenario 1 to scenario 3. However, even in the case of scenario 1, where the videos are public and the information can spread through multiple possibly unknown channels, for 75% of the videos the relative error is less than 0.1 (i.e. actual value is within ±10% of the predicted value).

Our study, while conducted at a very small scale, hints at the potential for accurately predicting future popularity in different real-world scenarios. We plan to conduct larger-scale studies in the future. Furthermore, there is significant ongoing research on predicting information propagation and dissemination in online systems and new techniques are being proposed. As the field advances, we expect the accuracy of predictions to improve as well.

B. Overheads associated with fine-tuning exposure

At first glance, it would appear that supporting exposure control would require users to check and fine-tune the exposure for every single piece of their information separately, which raises

usability concerns. In practice, users might want to organize all their pieces of information into a small number of groups, each with a different level of desired exposure. So when a new piece of information is uploaded, they can easily set its exposure to the desired level by choosing the appropriate exposure group. The overheads involved here would be no greater, if not lower, than the overheads involved with configuring privacy settings of uploaded content in social media sites today.

V. RELATED WORK

We now provide a brief overview of related work.

A. Defining privacy

Legal studies have long attempted to define privacy [21]–[23], with each of these approaches capturing different aspects of how privacy is perceived. Present discussions of privacy in social networks as well as in the Internet mostly adapt [13] the privacy definition presented by Westin [23]: *Privacy is the ability for people to determine for themselves when, how, and to what extent, information about them is communicated to others.* Westin's definition captures user's expectations of privacy, and all privacy management models (including ours) try to encode and respect these expectations.

B. Privacy is more than access control

Recent work [1], [4], [13] argues that access control is insufficient to meet the definition of privacy through anecdotal examples. Boyd et al. [4] presented an example where Facebook users felt their privacy was violated when Facebook launched News Feed. As mentioned in Section I-A, the News Feed change did not violate any access control policy, but the users still perceived a privacy violation. To address their privacy concern Facebook quickly reacted [16] to provide more fine grained access control mechanisms. However, researchers have shown [3], [18] that in spite of these mechanisms, Facebook users still severely underestimate the number of users who access their content. Other work [1] has discussed privacy violations in the context of data aggregators, showing that these aggregators were collecting only public data and were not violating any access control.

C. New models to manage privacy

Existing work provides privacy management for specific situations where access control may not be sufficient. Differential privacy [6] tries to limit the individual information leaked while querying a database for aggregate statistics. The recently proposed semantic privacy model [14] argues that users should specify "semantically" how they want the user data to be accessed and then a system should translate those semantic privacy preferences to syntactic privacy, e.g. via access control. While both of these approaches are significant steps towards better privacy control, neither directly addresses the issue of which users *actually learn* content.

VI. CONCLUSIONS

We argue that access control, the traditional model for managing privacy, is inadequate in today's online world. In this paper, we propose an alternative model for information

privacy based on exposure. A key difference compared to access control is that exposure captures the principals who *actually learn* a piece of information rather than who *can directly access* a piece of information. We believe exposure is an intuitive measure that captures the privacy implications of publishing information much better than access control.

We discussed how the concept of exposure can be leveraged in practical systems to enable users to manage their online privacy better. A key challenge we face here is predicting the future exposure of a piece of information and allowing users to control its exposure. We argue that existing literature on information dissemination can be leveraged to quantify and predict exposure. The huge volumes of data that many system operators collect about their users and their activities can be exploited to help users better understand and control the exposure of their information.

REFERENCES

[1] Privacy is not Access Control (But then what is it?). http://33bits.org/2010/02/13/privacy-is-not-access-control/.

[2] R. Baden, A. Bender, N. Spring, B. Bhattacharjee, and D. Starin. Persona: An Online Social Network with User-defined Privacy. In *SIGCOMM*, 2009.

[3] M. S. Bernstein, E. Bakshy, M. Burke, and B. Karrer. Quantifying the invisible audience in social networks. In *CHI*, 2013.

[4] D. Boyd. Facebook's Privacy Trainwreck. *Convergence: The International Journal of Research into New Media Technologies*, 14(1), 2008.

[5] Deleting My Teenage Tweets: A Student Journalists Perspective. http://ajr.org/delete-tweets/.

[6] C. Dwork. Differential privacy. In *ICALP*, 2006.

[7] How do I know who's seen each post or message in a group? https://www.facebook.com/help/409719555736128.

[8] What are the privacy options for groups? https://www.facebook.com/help/www/220336891328465.

[9] F. Figueiredo, F. Benevenuto, and J. M. Almeida. The tube over time: characterizing popularity growth of youtube videos. In *WSDM*, 2011.

[10] Friendlist Manager. http://apps.facebook.com/friendlist_manager, 2012.

[11] S. Guha, K. Tang, and P. Francis. NOYB: Privacy in Online Social Networks. In *WOSN*, 2008.

[12] L. Hong, O. Dan, and B. D. Davison. Predicting popular messages in Twitter. In *WWW*, 2011.

[13] L. Kagal and H. Abelson. Access Control is an Inadequate Framework for Privacy Protection. In *W3C Workshop on Privacy for Advanced Web APIs)*, 2010.

[14] B. Krishnamurthy. Privacy and online social networks: can colorless green ideas sleep furiously? *IEEE Security & Privacy*, 11(3), 2013.

[15] K. Lerman and T. Hogg. Using a model of social dynamics to predict popularity of news. In *WWW*, 2010.

[16] An Open Letter from Mark Zuckerberg. https://blog.facebook.com/blog.php?post=2208562130.

[17] Linear regression. http://en.wikipedia.org/wiki/Linear_regression.

[18] Y. Liu, K. P. Gummadi, B. Krishnamurthy, and A. Mislove. Analyzing Facebook privacy settings: User expectations vs. reality. In *IMC*, 2011.

[19] Rebecca Black phenomenon. http://youtube-trends.blogspot.de/2011/03/rebecca-black-phenomenon.html.

[20] G. Szabo and B. A. Huberman. Predicting the popularity of online content. *Communications of ACM*, 53(8), Aug. 2010.

[21] United Nations General Assembly. Universal Declaration of Human Rights (UDHR). http://www.un.org/en/documents/udhr/index.shtml.

[22] S. D. Warren and L. D. Brandeis. The Right to Privacy. *Harvard Law Review*, 4(5), 1890.

[23] A. Westin. *Privacy and Freedom*. Bodley Head, 1970.