

Policy Aware Social Miner

Sharon Paradesi, Oshani Seneviratne, Lalana Kagal
Decentralized Information Group
CSAIL, MIT
Email: {paradesi, oshani, lkagal}@csail.mit.edu

Abstract—There is a wealth of sensitive information available on the Web about any individual that is generated either by her or by others on social networking sites. This information could be used to make important decisions about that individual. The problem is that although people know that searches for their personal information are possible, they have no way to either control the data that is put on the Web by others or indicate how they would like to restrict usage of their own data.

We describe a framework called Policy Aware Social Miner (PASM) that would provide a solution to these problems by giving users a way to semantically annotate data on the Web using policies to guide how searches about them should be executed. PASM accepts search queries and applies the user's policies on the results. It filters results over data the user owns and provides the user's refutation link on search results that the user does not own. These usage control mechanisms for privacy allow users to break away from siloed data privacy management and have their privacy settings applied to all their data available on the Web.

Keywords—Usage restrictions; Refutations; Social network mining; Web mining

I. INTRODUCTION

There is a large amount of personal information available on the Web about any individual that is generated either by her or by others. Further, more and more users on the Web are generating information that is deemed private. A corporation or a person can use that information to make decisions about that individual. For example, it is becoming fairly common to hear about employers using Web search tools (such as SocialIntelligence [1]) to gather information about a potential hire. Our motivation for building this system stems from the Mosaic Theory, which states in [2] that apparently harmless pieces of information could potentially reveal a damaging picture when pieced together. Individually, the pieces may be of no value because they do not reveal anything particularly significant or dangerous about the person. However, when the pieces are viewed together, the mosaic may present a remarkably different picture of the person. Forming a mosaic of someone is very easy to do using the Web because of the availability of various services like search engines, social networks, etc. However, it is well-known that people usually share information within a specific context or purpose. However, putting the pieces of a person's life together to portray that person's characteristics may violate that person's contextual integrity norms [3]. The problem is that although people

know that such searches are possible and that they can be easily profiled on the Web, they have no way to either control the data that is put on the Web by others or indicate how they would like to restrict usage of their own data.

In order to provide a solution for the users, we are currently developing a tool called Policy Aware Social Miner (PASM) that aims to provide a framework to perform searches for people in a policy-aware manner. With respect to searches, we can broadly categorize the users involved into the following two groups:

- *Data subjects or producers*: These are users who rely on PASM to control how data about them (either self generated on social networks or provided by others) is used or interpreted. For the rest of the paper, we will refer to them as 'data subjects'.
- *Data consumers*: These are users who use PASM to investigate a data subject for a specific purpose (employment, insurance claim processing, etc.) while complying with the latter's policies.

In order for PASM to work, we make the following assumptions in this research:

- 1) Both the data consumer and the data subject are willing to be identified to PASM using an authentication mechanism (e.g., Facebook).
- 2) The data consumer will use PASM for an official search about a specific person and not for a casual Web search.
- 3) The data consumer believes in the principle of accountability and wants to do the right thing by performing necessary searches on the data subject in a manner that respects the data subject's privacy.
- 4) The data consumer is willing to provide a truthful intent (or purpose of search) that accurately reflects the purpose for performing that search.
- 5) The data consumer is interested in and willing to find out both the sides of a story or event documented on the Web by viewing the refutations created by the data subject.

The rest of this paper is organized as follows. In section II, we describe motivating scenarios to make a case for PASM. We then describe the system architecture and implementation of PASM. We explain the design tradeoffs that guided the design of PASM and explain the threat model for PASM in section V. Finally, we describe related research in this

area and conclude the paper.

Before concluding this section, we would like to highlight the main contributions of PASM.

- 1) Provides a framework to enable a policy-aware search on the Web.
- 2) Enables a data consumer to participate in a policy-aware search and view the complete story of an event or incident by being able to access the counterarguments set forth by the data subject.
- 3) Provides a platform for the data subject to annotate data on the Web in the following two ways:
 - a) If they own the data, they can attach policies to it to restrict how that data can be used;
 - b) If they do not own the data, they can write a comment refuting the implications of the data (known as refutation) and link it to third-party data.

II. SCENARIOS AND SOLUTIONS

In this section, we discuss two problematic scenarios related to searches on the Web and then motivate the potential solution used by PASM.

A. Motivating Scenarios

We use the following two hypothetical scenarios to make the case for PASM. To be consistent in both the scenarios, we assume that the data subject is a fictitious persona called *Alice Metzger* who has a Facebook profile with the username *ametzer1*.

1) *Personal Data*: Alice's adult daughter recently developed psoriasis. Alice's friends on Facebook start writing on her Facebook wall to inquire about the treatment process and to recommend remedies and names of doctors specializing in this area. Though it is her daughter who has the medical condition, Alice fears that anyone who looks at these posts without knowing the complete story may mistakenly conclude that Alice has psoriasis. This is especially problematic if that person happens to be a medical insurance agent, since this information can negatively affect her insurance. Therefore, she creates a policy to protect any information on her Facebook account containing the words "psoriasis", "treatment", or "doctor" to be filtered out from searches for her medical information.

2) *Third-party Data*: Alice comes across a local newspaper article that mentions her recent DUI charge. However, she notices that it incorrectly states that she was convicted of the charges, whereas, in reality, she was not. Since she does not have any way to address this issue online in a way that others who come across this news article in a search would notice, she writes a post on her website, stating the facts and uses PASM to create a link between the offending news article and the post on her website.

B. Potential Solutions

There are many ways one can go about solving the scenarios discussed in this section. For instance, an obvious solution for the first scenario would be to restrict access to the information that Alice considers sensitive. However, not all users are aware of or inclined to set the necessary permissions on social networking sites. It does not help that some social networking sites update their privacy settings frequently. Even if a data subject is scrupulous about setting proper privacy controls for her data on Facebook, if a friend tags her in a photo, then the privacy controls applied are different from what she would normally expect for her own photos. Therefore, we need to protect the data subjects from *usage misuse* rather than *access violations*.

As mentioned in Section I and alluded to in the scenarios described previously, we provide the following solutions to the data subjects – restrictions and refutations as explained below.

1) *Restrictions*: "Respect My Privacy" (RMP) [4] is a policy specification language that helps to create usage policies using which users could govern how their social network data can be used during a search. It offers a pre-defined set of usage policies that are similar in concept to a Creative Commons license in that it suggests means of sharing and viewing of data.

We use RMP because it is a simplified representation that has sufficient coverage of the categories of sensitive usage of data on the Web. The data subject can choose from the RMP restrictions of No-Employment, No-Commercial, No-Financial, No-Depiction and No-Medical.

2) *Refutations*: We define a *refutation* to be a counterargument to a piece of information on the Web. More often than not, a refutation conveys an opposite sentiment compared to the content of the document it is refuting.

Note that we are not trying to restrict or control the general Web search using refutations. Instead, PASM provides both sides of the story to a data consumer during a search by highlighting the data subject's response to a document on the Web. The goal is to prevent the data consumer from arriving at a wrong conclusion by not knowing the facts from both sides of the story.

III. SYSTEM ARCHITECTURE

A high level functional model for PASM is illustrated in Fig. 1. We discuss the functional components of the framework from the data subject and data consumer perspectives.

A. Data Subject's Perspective

Interaction begins when the data subject logs into the system. The identity that the system needs is obtained via Facebook's OAuth authentication mechanism [5]. After being successfully authenticated, the data subject is allowed to create, view or update her policies using PASM.

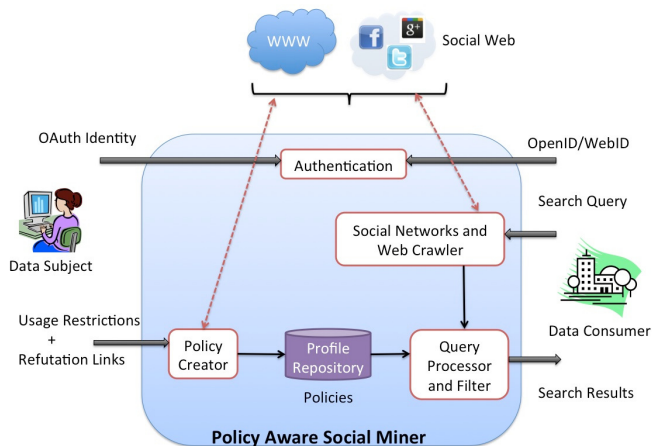


Figure 1. High-level overview of the system. This figure shows how the data subject and consumer interact with PASM and the data sources used in their interactions. The thick arrows show external interaction (of PASM) with the data subject and consumer, while the narrower ones depict the internal flow of control and data. The red, dashed arrows indicate the external (Web and social networks) search performed by PASM.

The data subject can create policies over data she owns (on Facebook) and refutations over data she does not own. For the latter case, she can write a post on any platform (such as Twitter, Wordpress and so on) and link it to the offending document using PASM. All this can be done with the help of the *Policy Creator*. The policies are then stored in a centralized *Profile Repository*.

Though the current implementation is tied to Facebook, it is straightforward to extend PASM to operate with other social networks (such as, Flickr, Google+, or LinkedIn) for the following reasons. First, the data subject only needs to specify the keyword(s) that are sensitive to her in the usage restrictions. Therefore, these restrictions are applicable for data across multiple social networks. Second, adding another scheme of authentication for the new social network is easy. Since PASM is modular, as long as you have a mapping for the representation of data items and an authentication mechanism for a specific social network (such as Flickr, Google+, or LinkedIn), you can switch to (or add) a new social network easily. Thus, PASM can be extended to be used with any social networking site.

B. Data Consumer's Perspective

Interaction begins when the data consumer logs into the system using OpenID or WebID. After being successfully authenticated, the data consumer enters a search query containing the data subject's identity, keyword(s) to guide the search, and the intent of the search. The identity of the data consumer along with the intention and query can be logged for future accountability [6].

PASM then searches social networks (Facebook) and the Web using the *Social Networks and Web Crawler* and looks for documents that mention Alice along with the keyword(s) entered. It then decides, using her policies, whether a data item returned from her Facebook account can be shown to the data consumer. If a data item contains the keyword(s) but is not restricted by her policies, it is returned to the data consumer. Otherwise, it is filtered out. Next, given the list of document URLs returned by the Google Custom Search API, PASM checks whether any of them were refuted by Alice. If any such refutations are found, links to the refutations are shown beside the URLs of the corresponding documents in the results. The *Query Processor and Filter* takes care of processing the query and applying the restrictions and refutations to the search results before showing them to the data consumer.

IV. IMPLEMENTATION

In this section, we discuss the actual implementation details. We first discuss the important components built into or used in the system and then demonstrate the working of the system using screenshots.

A. Important Components

1) *Usage Restrictions and Synonym Lookup*: PASM creates usage restrictions for Alice using RMP as described in the previous section. Instead of expecting Alice to enter *all* the terms she finds sensitive in the interface presented to her, PASM performs a simple synonym lookup using Big huge Thesaurus [7] and expands her restrictions to include those synonyms as well. Before displaying data to the data consumer, PASM filters out data returned from from her Facebook account that contain either the keywords she mentioned or their synonyms obtained using the thesaurus. For example, suppose that Alice creates the policy as explained previously using the keywords "psoriasis", "treatment", and "doctor". PASM determines that the synonyms of the word "psoriasis" are "disease of the skin", "skin condition", "skin disease", "skin disorder" and "skin problem". It then expands Alice's policy to include these synonyms as keywords as well.

2) *Data Sources*: Before Alice can interact with PASM as a data subject, she has to first install the application on her Facebook account and thus authenticate herself using her OAuth credentials. When the data consumer searches for Alice using her Facebook username (ametzter1), we map that identifier to the stored OAuth token (stored in the profile repository) and thus obtain her policies and refutations. For the Web search, PASM uses the Google Custom Search API [8]. If any of the URLs of the returned documents match those that Alice refuted, PASM displays her refutations below those documents in the results.

B. System Demonstration

We currently have a working demo of PASM for both a data consumer and a data subject ¹.

1) *Data Subject Interface*: As mentioned earlier, Alice has to first verify her identity using Facebook’s OAuth mechanism by logging into Facebook and installing the application hosted on Facebook and integrated with PASM. Once she does this and visits the interface, she will find her name and Facebook photo on the page to indicate her credentials. She can then choose to view or edit her profile. A *profile* is an abstract representation of the policy (consisting of usage restrictions and refutation links) of a data subject. PASM currently allows users to only add new usage restrictions and refutation links to their profiles. In Fig. 2, we see Alice’s profile containing one usage restriction and one refutation link. In Fig. 3, we see interface using which she can create a new usage restriction by entering keywords and choosing appropriate categories of searches from which to restrict data items containing those keywords.



Figure 2. Viewing Alice’s policy. Note the usage restrictions and refutations present in her policy.

2) *Data Consumer Interface*: In Fig. 4, we see the interface presented to the data consumer while performing a search. According to our earlier scenarios, the data consumer may enter Alice’s Facebook username or id, keyword(s) relevant to the search, and the intent of the search. Let us assume that the data consumer is searching for data containing the keyword “medicine” and truthfully declares that the intent of the search is for medical purposes.

Upon submitting the form, the data consumer would see Fig. 5 which shows data from Alice’s Facebook account that pass the filters created by her policies and also data from the Web obtained using Google Custom Search engine.

To illustrate the impact of Alice’s policies, Fig. 6 shows the results returned by PASM when a search for the same keyword (“medicine”) but a different purpose (employment) was declared. Since the new intent (employment) does not

¹<http://musigma.csail.mit.edu:2020/pasm.html>

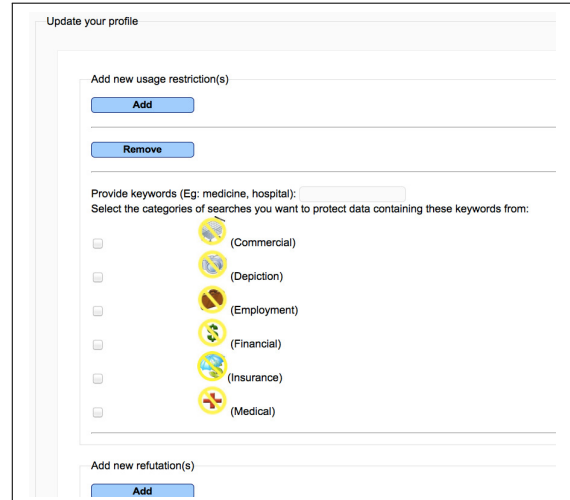


Figure 3. Updating Alice’s policy. This screenshot shows how to create new usage restrictions using PASM.

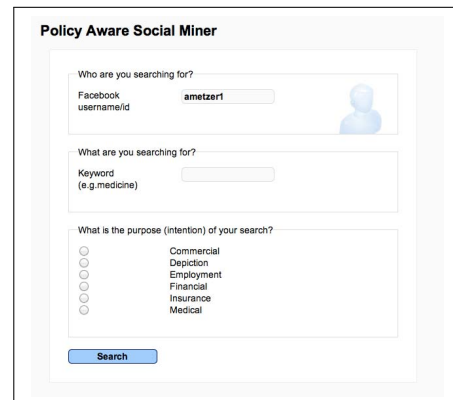


Figure 4. Search interface for the data consumer. A data consumer can search for a specific person by entering that person’s Facebook username or id (Alice’s username is ametz1, as mentioned earlier). One or more keywords along with the actual intention for search are required to properly guide PASM’s search process.

match the filters created by Alice’s policies for that keyword, the data consumer is able to view additional posts from her Facebook account. The additional posts shown in Fig. 6 are ambiguous (it is not clear whether Alice or her family member has psoriasis) and thus demonstrate the dangers of the Mosaic Theory mentioned in Section I.

The Web search performed on Alice using the keyword “medicine” returns the results shown in Figs. 5 and 6. Note that Alice’s refutation link is displayed prominently below the article she wishes to refute. This makes the data consumer aware of her counterargument to that document and thus helps the data consumer understand her side of the story as well.

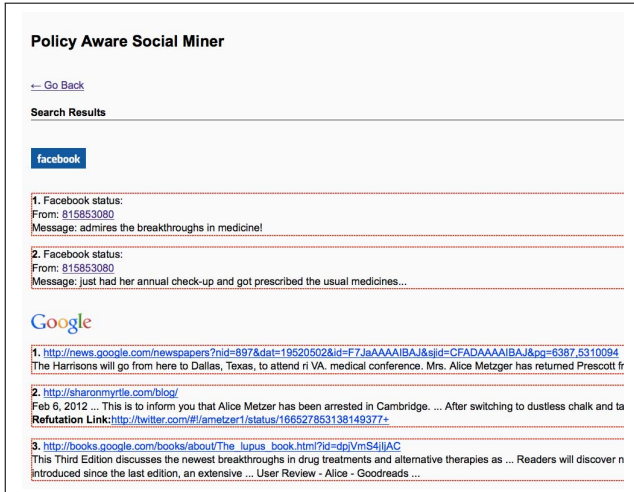


Figure 5. Search Results (intent: medical) using the keyword “medicine”. The search returns data from Facebook (social search) and Google (Web search).

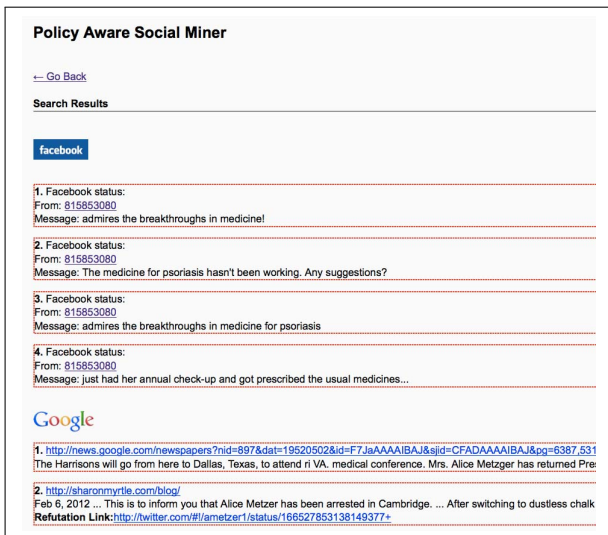


Figure 6. Search Results (intent: employment) using the keyword “medicine”. The search returns data from Facebook (social search) and Google (Web search) (similar to that shown in Fig. 5). However, note the additional results when performing the employment search for this set of keyword(s).

V. DESIGN TRADEOFFS

In this section, we explain the various design decisions we took while building PASM. We can broadly categorize these decisions along the axes of data and threat models.

A. Data Gathering, Storage and Query

There are three ways we can gather and store data from Alice’s social networks and run queries.

- Statically download her data into a local repository as a graph when she installs the application and associate her policies and refutations to this graph.

- Dynamically query Facebook to retrieve her data and form a temporary graph from the returned results for a particular search.
- Bypass the techniques of the Semantic Web and instead deal with the data in the format that Facebook provides (JSON).

PASM utilizes the *second approach* because of the following reasons:

- *Freshness of data:* Since PASM dynamically fetches Alice’s data for each query of each data consumer, the results accurately represent the actual content in her Facebook account at that point in time. For example, suppose that Alice has a post on Facebook containing the phrase “medicine” and “rash”. In our scenarios, the post would pass the filter created by her policy (because “rash” was neither mentioned as a keyword nor was considered a synonym of the existing keywords) and would be shown to the data consumer in all the three techniques. However, if she actually deletes that post on Facebook, then the first technique will still return the post to the data consumer, thus violating her implicit assumption that the post is no longer visible to any data consumer. PASM, however, would not know about that post because Facebook would not have sent it in the results since it does not exist in her account.
- *Efficiency of the search queries:* Since the third technique does not employ the use of a Semantic Web graph structure, the concepts and relationship among those concepts are not apparent. Therefore, when using this approach, we will have to traverse the results multiple times to figure this out. However, PASM can easily issue appropriate queries to the graph and discover the important concepts and relationships quickly.
- *Scalability:* Facebook has around 845 million active monthly users ². Using the first approach, if we try to create a separate graph for each user in a local repository, and only a few of them are used, the system will be unnecessarily overloaded. However, PASM does not store the graph permanently and only uses the graph for a particular search by a data consumer.
- *Security:* Storing the data subject’s information in local repositories (as done in the first approach) would pose a grave risk in case of a security breach on the server. PASM avoids this risk by building temporary graphs stored in memory and deletes them once the search is completed.

B. Threat Model

The major threats affecting PASM concern security risks, false intents and spamming.

²<http://newsroom.fb.com/content/default.aspx?NewsAreaId=22>

1) *Security of Policies*: In PASM, we store the policies of data subjects on a centralized secure server. The policies themselves may potentially be sensitive because they indicate exactly what the data subject wants to protect against. Storing them securely will protect them from privacy implications. Note that PASM is a privileged and trusted system that mediates access to a data subject's information. A data consumer has access to the data subject's information only after successfully authenticating with PASM and only by using the search interface. Even then, the data consumer cannot directly view the policies but only the effects of applying the policies on the underlying data.

2) *False Intent*s: There is currently no mechanism built into PASM to deter a malicious data consumer from declaring a false intent. By stating that the data consumer's actions will be logged, we hope to deter the data consumers from declaring false intents. Therefore, one of our assumptions is that the data consumers will be truthful about their search intents.

3) *Spamming*: We define *spam* in refutations as any content not related to the document being refuted. The most obvious type of spam is created by advertisers placing their ads as refutations. Another type of spam involves a person creating a refutation to a document that is tangentially related to it. This latter scenario can be used to garner sympathy for the spammer's cause though the readers of the original document may not be interested in it.

PASM avoids these kinds of spamming through the use of authentication and providing directed results. First, data subjects can create refutations only after authenticating themselves to PASM. Therefore, data subjects cannot post under the guise of someone else. Second, even if some data subjects are malicious and post spam refutations, their refutations are only displayed to a data consumer when a specific search is executed for them. Therefore, if a malicious spammer creates a spam refutation for a document talking about Alice, that refutation will be shown to a data consumer only when the data consumer explicitly searches for the spammer (and is not shown when the data consumer searches for Alice).

VI. RELATED WORK

A. Data

With the proliferation of data among various data stores on the Web, there is an additional challenge in creating a unified technique to preserve user's privacy expectations. Clifton et al. in [9] propose a privacy framework for data sharing and integration. However, the framework does not involve user input regarding their expectations and policies. On the other hand, PASM involves users in the process to make them aware of how their data is being used.

PASM uses a pull-based search where a data consumer wanting to learn about Alice would use the PASM interface and obtain relevant information. Passant and Mendes [10]

describe a push-based service called PubSubHubbub where news of data updates is proactively pushed to users. A push-based service would have to be careful about logging the intent of searches when data consumers perform searches with multiple intents (e.g., medical insurance agents). In those situations, the system may push information to the data consumer without knowing that the data consumer needs it and regardless of whether the intent of the search is still the same.

B. Annotation systems on the Web

Reputation.com is a service aimed to replace malicious reviews with truthful, positive feedback. However, our goal is not to hide or to resolve any issue related to the data subject. Instead, PASM aims to present both sides of the story to the data consumer.

Google originated the concept of Sidewiki, through which users were able to annotate webpages. Because it was an open platform, it became widely susceptible to spamming. As described in the previous section, we prevent spamming in PASM because the refutations by a data subject are shown to a data consumer only when the data consumer does an explicit search for that particular data subject. Thus, even though spammers may post to the site, their posts will not be displayed to data consumers unless the data consumers are searching for information about those spammers.

Other researchers [11], [12], [13], [14], [15] explore annotations as a medium for optimizing Web searches and for classification tasks. The primary difference between these approaches and PASM is the audience. PASM is meant for users to annotate documents on the Web related to them because the annotations will be displayed to data consumers that look for their name specifically.

Ennals et al. in [16], [17] discuss how disputes on the Web happen and how to show users that certain content on the Web is disputed by other sources on the Web. Users can also highlight pieces of text to show that they disagree with it. The system, DisputeFinder, maintains a centralized database with the collection of disputed claims and shows them to the users if the users browse one of the pages. PASM differs from DisputeFinder because we primarily focus on counterarguments related to personal issues.

C. Social Screening Systems and Reputation Systems

SocialIntelligence [1] is an FTC-approved commercial system that searches for information about an employee and notifies the employer whether the employee has passed the test or not. We provide a comparable platform to perform searches on the Web. While SocialIntelligence is focused to serve the needs of data consumers by providing them data about the data subjects in a way that does not reveal discriminatory information (race, gender etc), we enable data subjects to become more involved in these searches by helping them protect their data and share their counterarguments.

We provide a similar framework that can be extended to facilitate inferences. However, we also enable data subjects to create usage policies to protect themselves from potential incorrect inferences by data consumers when searching for data about the data subjects on the Web.

VII. CONCLUSION AND FUTURE WORK

In this paper, we demonstrate a need for and describe a system called Policy Aware Social Miner (PASM). With the help of PASM, a data subject is given the ability to create policies over her data and refutations for data that she does not own on the Web. Further, PASM provides a way for the data consumers to execute policy-aware searches. Looking forward, we envision working on the following improvements that will be useful extensions to PASM:

- 1) We currently use string matching techniques to filter data that matches Alice's restrictions and refutations. It would be helpful to investigate the semantics of the keywords by using DBpedia and other structured data sources.
- 2) Use richer representations of policy coupled with a reasoner to support more finer grained usage restrictions.
- 3) Explore extending the policies and refutations applied to a pull-based service to one that is based on pushing (e.g., PubSubHubbub).
- 4) Enable third party refutations and refutations of refutations and handle associated challenges of determining the legitimacy of the claims.
- 5) For better security, users can adopt a decentralized approach and host their own policies. One possible approach is to use data.fm [18], which is a read/write linked data service.
- 6) Make the interface more user-friendly for the data subjects to build policies.

ACKNOWLEDGMENTS

The authors would like to thank Hal Abelson, Joe Pato, Susan Landau and other members of the Decentralized Information Group for their input.

REFERENCES

- [1] "Social intelligence," <http://www.socialintel.com/>.
- [2] D. Pozen, "The mosaic theory, national security, and the freedom of information act," *Yale Law Journal*, 2005.
- [3] H. F. Nissenbaum, "Privacy as contextual integrity," *Washington Law Review*, vol. 79, 2004.
- [4] T. Kang and L. Kagal, "Enabling Privacy-Awareness in Social Networks," in *Intelligent Information Privacy Management*, vol. 2010. AAAI, 2010, pp. 98–103.
- [5] "Oauth specification," <http://oauth.net/>.
- [6] D. Weitzner, H. Abelson, T. Berners-Lee, J. Feigenbaum, J. Hendler, and G. Sussman, "Information accountability," pp. 82–87, June 2008. [Online]. Available: <http://dspace.mit.edu/bitstream/handle/1721.1/37600/MIT-CSAIL-TR-2007-034.pdf?sequence=2>
- [7] J. Watson, "Big huge thesaurus," <http://words.bighugelabs.com/api.php>.
- [8] "Google custom search api," <http://www.google.com/cse/>.
- [9] C. Clifton, M. Kantarcioğlu, A. Doan, G. Schadow, J. Vaidya, A. Elmagarmid, and D. Suciú, "Privacy-preserving data integration and sharing," in *Proceedings of the 9th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery*, ser. DMKD '04. New York, NY, USA: ACM, 2004, pp. 19–26. [Online]. Available: <http://doi.acm.org/10.1145/1008694.1008698>
- [10] A. Passant and P. N. Mendes, "sparqlpush: Proactive notification of data updates in rdf stores using pubsubhubbub," 2010.
- [11] A. Zubiaga, R. Martínez, and V. Fresno, "Getting the most out of social annotations for web page classification," *Proceedings of the 9th ACM symposium on Document engineering*, p. 74, 2009. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1600193.1600211>
- [12] S. Bao, G. Xue, X. Wu, Y. Yu, B. Fei, and Z. Su, "Optimizing web search using social annotations," *Proceedings of the 16th international conference on World Wide Web WWW 07*, vol. pages, no. C, pp. 501–510, 2007. [Online]. Available: <http://portal.acm.org/citation.cfm?doi=1242572.1242640>
- [13] W. Choochaiwattana and M. B. Spring, "Applying Social Annotations to Retrieve and Re-rank Web Resources," in *Information Management and Engineering 2009 ICIME 09 International Conference on*. Ieee, 2009, pp. 215–219.
- [14] R. Sanderson and H. V. D. Sompel, "Making web annotations persistent over time," in *Proceedings of the Joint Conference on Digital Libraries JCDL*. ACM Press, 2010, pp. 1–10. [Online]. Available: <http://portal.acm.org/citation.cfm?id=1816123.1816125>
- [15] M. G. Noll and C. Meinel, "Exploring social annotations for web document classification," *Proceedings of the 2008 ACM symposium on Applied computing SAC 08*, p. 2315, 2008. [Online]. Available: <http://portal.acm.org/citation.cfm?doi=1363686.1364235>
- [16] R. Ennals, D. Byler, J. M. Agosta, and B. Rosario, "What is disputed on the web?" *Proceedings of the 4th workshop on Information credibility WICOW 10*, p. 67, 2010. [Online]. Available: <http://portal.acm.org/citation.cfm?doi=1772938.1772952>
- [17] R. Ennals, B. Trushkowsky, and J. M. Agosta, "Highlighting disputed claims on the web," *Proceedings of the 19th international conference on World wide web WWW 10*, no. Figure 3, p. 341, 2010. [Online]. Available: <http://portal.acm.org/citation.cfm?doi=1772690.1772726>
- [18] "data.fm," <http://data.fm/>.