Available Online at www.ijcsmc.com

International Journal of Computer Science and Mobile Computing

A Monthly Journal of Computer Science and Information Technology

ISSN 2320–088X
IMPACT FACTOR: 5.258

IJCSMC, Vol. 5, Issue. 3, March 2016, pg.249 – 255

PRESERVING PRIVACY FOR UPLOADED IMAGES ON THE CONTENT SHARING SITES USING A3P (ADAPTIVE PRIVACY POLICY)

[1]Dr. M.P.Sivaramkumar, [2]Mr. C.B.Sushanth, [3]Mr. P.Karthikeyan, [4]Mr. S.Mohammed Afsar

[1]Department of Information Technology, Prathyusha Engineering College, Thiruvallur, India
[2]Department of Information Technology, Prathyusha Engineering College, Thiruvallur, India
[3]Department of Information Technology, Prathyusha Engineering College, Thiruvallur, India
[4]Department of Information Technology, Prathyusha Engineering College, Thiruvallur, India
[1] sivaramkumar.it@gmail.com, [2] sushanth_fernando@yahoo.com, [3] karthi.palani1995@gmail, [4] apsarcool2@yahoo.com

ABSTRACT: ***Maintaining privacy on social sites has become a major problem as the volume of image sharing is increasing these days. The recent publicized incidents where users accidently share their personal information demonstrated this project. From these incidents, the need of accessing control tools to help users is obvious. Therefore we propose an Adaptive Privacy Policy Prediction (A3P) system to help users create privacy settings for their images. We propose a two-level framework which according to the user's available history on the site which determines the best available privacy policy for the user's images which are uploaded. Our solution relies on an image classification framework for image categories which may be associated with similar policies, and on a policy prediction algorithm to automatically generate a policy for each newly uploaded image, also according to users' social features.***

Index Terms— web-based services, Online information services

1 INTRODUCTION

Images are now one of the key enablers of users' connectivity. Sharing takes place both among previously established groups of known people or social circles (e.g., Google or Picasa), and also increasingly with people outside the users social circles, for purposes of

Social to help them identify new peers and learn about peers interests and social surroundings. However, semantically rich images may reveal content sensitive information [2]. Consider a photo of a student 2012 graduation ceremony, for example. It could be shared within a Google+ circle group, but may unnecessarily expose the students BA position family members and other friends. Sharing images within online content sharing sites, therefore may quickly lead to unwanted disclosure and privacy violations [3], [15]. Further, the persistent .Nature of online media makes it possible for other users to collect rich aggregated information about the owner of the published content and the subjects in the published content [3], [20], [17]. The aggregated information can result in unexpected exposure of one's social environment and lead to abuse of one's personal Information.

Most content sharing websites allow users to enter their privacy preferences. Unfortunately, recent studies have shown that users struggle to set up and maintain such privacy Settings. One of the main reasons provided is that given the amount of shared information this process can be tedious and error-prone. Therefore, many have acknowledged the need of policy recommendation systems which can assist users to easily and properly configure privacy settings [7], [2], [15], .However, existing proposals for automating privacy settings appear to be inadequate to address the unique privacy needs of images [3], [5] due to the amount of information implicitly carried within images, and their relationship with the online environment wherein they are exposed. In this paper, we propose an Adaptive Privacy Policy Prediction (A3P) system which aims to provide users a hassle free privacy settings experience by automatically generating personalized policies. The A3P system handles user uploaded images, and factors in the following criteria that influence one's privacy settings of images:

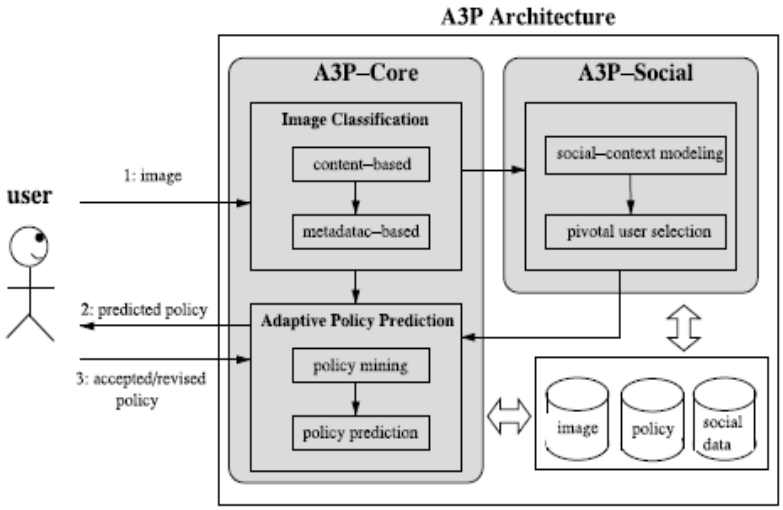

Fig. 1. System overview.

To assess the practical value of our approach, we built a system prototype and performed an extensive experimental evaluation. We collected and tested over 5,500 real policies generated by more than 160 users. Our experimental results demonstrate both efficiency and high prediction accuracy of our system.

A preliminary discussion of the A3P-core was presented in [7]. In this work, we present an overhauled version of A3P, which includes an extended policy prediction Algorithm in A3P-core (that is now parameterized based on User groups and also factors in possible outliers), and a new A3P-social module that develops the notion of social context refine and extend the prediction power of our system. We also conduct additional experiments with a new data set collecting over 1,400 images and corresponding policies, and we extend our analysis of the empirical results to unveil our system's performance.

The rest of the paper is organized as follows. Section 2 reviews related works. Section 3 introduces preliminary notions. Section 4 introduces the A3P-core and Section 5 introduces the A3P-Social. Section 6 reports the experimental evaluation. Finally, Section 7 concludes the paper.

2 RELATED WORK

Our work is related to works on privacy setting configuration in social sites, recommendation systems, and privacy analysis of online images.

2.1 Privacy Setting Configuration

Several recent works have studied how to automate the task of privacy settings (e.g., [7], [15], [20], [17], [15]).which recommend to users a suite of privacy settings that "expert" users or other trusted friends have already set, so that normal users can either directly choose a setting or only need to do minor modification. Similarly, Danezis[8] proposed a machine-learning based approach to automatically extract privacy settings from the social

context within which the data is produced. Parallel to the work of Danezis, Adu-Oppong et al. [10] develop privacy settings based on a concept of "Social Circles" which consist of clusters of friends formed by partitioning users' friend lists. Ravichandran et al. [15] studied how to predict a user's privacy preferences for location-based data (i.e., share her location or not) based on location and time of day. Fang et al. [15] proposed a privacy wizard to help users grant privileges to their friends. The wizard asks users to first assign privacy labels to selected friends, and then uses this as input to construct a classifier which classifies friends based on their profiles and automatically assign privacy labels to the unlabeled friends. More recently, Klemperer et al. [20] studied whether the keywords and captions with which users tag their photos can be used to help users more intuitively create and maintain access-control policies. Their findings are online with our approach: tags created for organizational purposes can be repurposed to help create reasonably accurate access-control rules. The aforementioned approaches focus on deriving policy settings for only traits, so they mainly consider social context such as one's friend list. While interesting, they may not be sufficient to address challenges brought by image files for which privacy may vary substantially not just because of social context but also due to the actual image content. As far as images, authors in [1] have presented an expressive language for images uploaded in social sites. This work is complementary to ours as we do not deal with Policy expressiveness, but rely on common forms policy specification for our predictive algorithm. IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, VOL. 27, NO. 1, JANUARY 2015 In addition, there is a large body of work on image content analysis, for classification and interpretation (e.g., [14], [15], [19]), retrieval ([12], [13] are some examples), and photo ranking [1], [16], also in the context of online photo sharing sites, such as Flickr [10], [9], [15]. Of these works, Zerr's work [8] is probably the closest to ours. Zerr explores privacy-aware image classification using a mixed set of features, both content and meta-data. This is however a binary classification (private versus public), so the classification task is very different than ours. Also, the authors do not deal with the issue of cold-start problem.

2.2 Recommendation Systems

Our work is related to some existing recommendation systems which employ machine learning techniques.
Chen et al. [9] proposed a system named Sheep Dog to automatically insert photos into appropriate groups and recommend suitable tags for users on Flickr. They adopt concept detection to predict relevant concepts (tags) of a photo. Choudhury et al. [10] proposed a recommendation framework to connect image content with communities in online social media. They characterize images through three types of features: visual features, user generated text tags, and social interaction, from which they recommend the most likely groups for a given image. Similarly, Yu et al. [11] proposed an automated recommendation system for a user's images to suggest suitable photo-sharing groups.
There is also a large body of work on the customization and personalization of tag-based information retrieval (e.g., [21], [23], [1]), which utilizes techniques such as association rule mining. For example, [17] proposes an interesting experimental evaluation of several collaborative filtering algorithms to recommend groups for Flickr users. These approaches have a totally different goal to our approach as they focus on sharing rather than protecting the content.

3 A3P FRAMEWORK

3.1 Preliminary Notions

Users can express their privacy preferences about their content disclosure preferences with their socially connected users via privacy policies. We define privacy policies according to Definition 1. Our policies are inspired by popular content sharing sites (i.e., Facebook, Picasa, Flickr), although the actual implementation depends on the specific content-management site structure and implementation.
Definition 1. A privacy policy P of user u consists of the following Components :

- Subject (S): A set of users socially connected to u.
- Data (D): A set of data items shared by u.
- Action (A): A set of actions granted by u to S on D.
- Condition (C): A boolean expression which must be satisfied in order to perform the granted actions.

In the definition, users in S can be represented by their identities, roles (e.g., family, friend, coworkers), or organizations (e.g., non-profit organization, profit organization). D will be the set of images in the user's profile. Each image has a unique ID along with some associated metadata like tags "vacation", "birthday". Images can be further grouped into albums. As for A, we consider four common types of actions: {view, comment, tag, download}. Last, the condition component C specifies when the granted action is effective.
C is a Boolean expression on the grantees' attributes like time, location, and age. For better understanding, an example policy is given below.

Example 1. Alice would like to allow her friends and coworkers to comment and tag images in the album named "vacation album" and the image named "summer.jpg" before year 2012. Her privacy preferences can be expressed by the following policy:
P: ½{friend, coworker}, {vacation_album, summer.jpg},
{comment, tag}, (date< 2012)_.

3.2 System Overview

The A3P system consists of two main components: A3P-core and A3P-social. The overall data flow is the following. When a user uploads an image, the image will be first sent to the A3P-core. The A3P-core classifies the image and determines whether there is a need to invoke the A3P-social. In most cases, the A3P-core predicts policies for the users directly based on their historical behavior. If one of the following two cases is verified true, A3P-core will invoke A3Psocial:
(i) The user does not have enough data for the type of the uploaded image to conduct policy prediction;
(ii) The A3P-core detects the recent major changes among the user's community about their privacy practices along with user's increase of social networking activities (addition of new friends, new posts on one's profile etc). In above cases, it would be beneficial to report to the user the latest privacy practice of social communities that have similar background as the user. The A3P-social groups users into social communities with similar social context and privacy preferences, and continuously monitors the social groups. When the A3P-social is invoked, it automatically identifies the social group for the user and sends back the information about the group to the A3P-core for policy prediction. At the end, the predicted policy will be displayed to the user. If the user is fully satisfied by the predicted policy, he or she can just accept it. Otherwise, the user can choose to revise the policy. The actual policy will be stored in the policy repository of the system for the policy prediction of future uploads.

4 A3P-CORE

There are two major components in A3P-core: (i) Image classification and (ii) Adaptive policy prediction. For each user, his/her images are first classified based on content and metadata. Then, privacy policies of each category of images are analyzed for the policy prediction. Adopting a two-stage approach is more suitable for policy recommendation than applying the common one-stage data mining approaches to mine both image features and policies together. Recall that when a user uploads a new image, the user is waiting for a recommended policy. The two-stage approach allows the system to employ the first stage to classify the new image and find the candidate sets of images for the subsequent policy recommendation. As for the one-stage mining approach, it would not be able to locate the right class of the new image because its classification criteria need both image features and policies whereas the policies of the new image are not available yet.

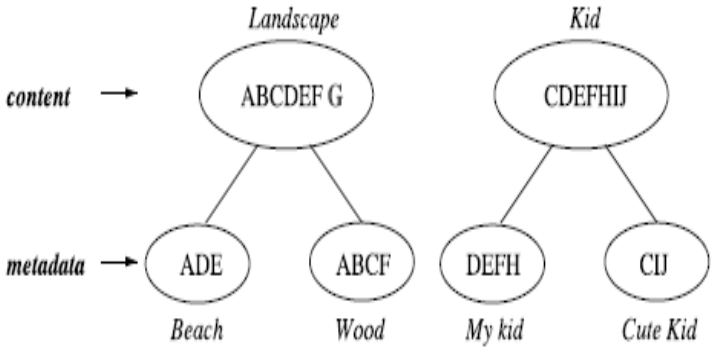

Moreover, combining both image features and policies into a single classifier would lead to a system which is very dependent to the specific syntax of the policy. If a change in the supported policies were to be introduced, the whole learning model would need to change.

4.1 Image Classification

To obtain groups of images that may be associated with similar privacy preferences, we propose a hierarchical image classification which classifies images first based on their contents and then refine each category into subcategories based on their metadata. Images that do not have metadata will be grouped only by content. Such a hierarchical classification gives a higher priority to image content and minimizes the influence of missing tags. Note that it is possible that some images are included in multiple categories as long as they contain the typical content features or metadata of those categories.

Moreover, Fig. 2 shows an example of image classification for 10 images named as A, B, C, D, E, F, G, H, I, J, respectively. The content-based classification creates two categories: "landscape" and "kid". Images C, D, E and F are included in both categories as they show kids playing outdoor which satisfy the two themes: "landscape" and "kid". These two categories are further divided into subcategories based on tags associated with the images. As a result, we obtain two subcategories under each theme respectively. Notice that image G is not shown in any subcategory as it does not have any tag; image A shows up in both subcategories because it has tags indicating both "beach" and "wood".

4.1.1 Content-Based Classification

Our approach to content-based classification is based on an efficient and yet accurate image similarity approach. Specifically, our classification algorithm compares image signatures defined based on quantified and sanitized version of Haar wavelet transformation. For each image, the wavelet transform encodes frequency and spatial information related to image color, size, invariant transform, shape, texture, symmetry, etc. Then, a small number of coefficients are selected to form the signature of the image. The content similarity among images is then determined by the distance among their image signatures. Our selected similarity criteria include texture, symmetry, shape (radial symmetry and phase congruency [19]), and SIFT [15]. We also account for color and size. We set the system to start from five generic image classes: (a) explicit (e.g., nudity, violence, drinking etc), (b) adults, (c) kids, (d) scenery (e.g., beach, mountains), (e) animals. As a preprocessing step, we populate the five baseline classes by manually assigning to each class a number of images crawled from Google images, resulting in about 1,000 images per class. Having a large image data set beforehand reduces the chance of misclassification. Then, we generate signatures of all the images and store them in the database. Upon adjusting the settings of our content classifier, we conducted some preliminary test to evaluate its accuracy. Precisely, we tested our classifier it against a ground-truth data set, Image-net.org [17]. In Image-net, over 10 million images are collected and classified according to the wordnet structure. For each image class, we use the first half set of images as the training data set and classify the next 800 images. The classification result was recorded as correct if the synset's main search term or the direct hypernym is returned as a class. The average accuracy of our classifier is above 94 percent. Having verified the accuracy of the classifier, we now discuss how it is used in the context of the A3P core. When a user uploads an image, it is handled as an input query image. The signature of the newly uploaded image is compared with the signatures of images in the current image database. To determine the class of the uploaded image, we find its first m closest matches. The class of the uploaded image is then calculated as the class to which majority of the m images belong. If no predominant class is found, a new class is created for the image. Later on, if the predicted policy for this new image turns out correct, the image will be inserted into the corresponding image category in our image database, to help refine future policy prediction. In our current prototype, m is set to 25 which is obtained using a small training data set.

5 A3P-SOCIAL

The A3P-social employs a multi-criteria inference mechanism that generates representative policies by leveraging key information related to the user's social context and his general attitude toward privacy. As mentioned earlier, A3Psocial will be invoked by the A3P-core in two scenarios. One is when the user is a newbie of a site, and does not have enough images stored for the A3P-core to infer meaningful and customized policies. The other is when the system notices significant changes of privacy trend in the recommendation process.

5.1Analysis of Users' Characteristics

We are also interested in examining whether our algorithm performs better for users with certain characteristics. Therefore, we study possible factors relevant to the performance of our algorithm. We used a least squares multiple regression analysis, regressing performance of the A3P-core to the following possible predictors:

- _ Frequency of social network use was measured on a frequency rating scale (1 ¼ daily; 2 ¼ weekly; 3 ¼ monthly; 4 ¼ rarely; 5 ¼ never) with the item 'How often do you access Social Network Sites?'
- _ Privacy settings take timwas measured on a LikertScale (5-point rating scale, where 1 ¼ strongly agree and 5 ¼ strongly disagree) with the item 'Changing privacy settings for images uploaded on a social site can be very time consuming.'
- Frequency of sharing pictures was measured using three items (a ¼ 0:69) rated on a Likert scale.
- Frequency of changing privacy settings was measured using four items (a ¼ 0:86) rated on a Likert scale.
- Privacy concern was measured using four items (a ¼ 0:76) rated on a Likert scale. An example item is 'I have had concerns about my privacy due to shared images on social network sites.'

The model results are shown in Table 5. We can observe that the content of concern variable was the biggest predictor of performance of our algorithm (standardized b ¼ 0:461, p < 0:001). This suggests the importance of content in determining the privacy level of uploaded images to social network sites. Privacy concern was also a significant predictor of performance (standardized b ¼ 0:329, p < 0:01) with increased performance for those users who felt that images uploaded to social network sites allowed for exposure of personal information. Surprisingly, none of the other predictors were significantly related to performance of the A3P-core. We expected that frequency of sharing pictures and frequency of changing privacy settings would be significantly related to performance, but the results indicate that the frequency of social network use, frequency of uploading images and frequency of changing settings are not related to the performance our algorithm obtains with privacy settings predictions. This is a particularly useful result as it indicates that our algorithm will perform equally well for users who frequently use and share images on social networks as well as for users who may have limited access or limited information to share.

In the second round of experiments, we analyze the performance of the A3P-Social component by using the first set of data collection. For each user, we use the A3PSocial to predict policies and compare it with a base-line which does not consider social contexts but bases recommendation only on social groups that have similar privacy strictness level for same type of images.

Using the base-line approach, we note that regardless of the individual privacy inclination of the users, the best accuracy is achieved in case of explicit images and images dominated by the appearance of children. In both cases, users maintain more consistent policies, and our algorithm is able to learn them effectively. The largest variability, and therefore worse results occur for images denoting scenery, where the error rate is 15.2 percent.

Overall, the accuracy achieved by grouping users by strictness level is 86.4 percent. With A3P-Social, we achieve a much higher accuracy, demonstrating that just simply considering privacy inclination is not enough, and that "social-context" truly matters. Precisely the overall accuracy of A3P-social is above 95 percent. For 88.6 percent of the users, all predicted policies are correct, and the number of missed policies is 33 (for over 2,600 predictions). Also, we note that in this case, there is no significant difference across image types. For completeness, we compared the performance of the A3P-Social with alternative, popular, recommendation methods: Cosine and Pearson similarity [15]. Cosine similarity is a measure of similarity between two vectors of an inner product space that measures the cosine of the angle between them. In our case, the vectors are the users' attributes defining their social profile. The algorithm using Cosine similarity scans all users profiles, computes Cosine similarity of the social contexts between the new user and the existing users. Then, it finds the top two users with the highest similarity score with the candidate user and feeds the associated images to the remaining functions in the A3P-core.

Pearsons similarity instead measures how highly correlated are two variables, and is usually used to correlate users' ratings on recommended products. To adapt, we replaced the users rating from the Pearson similarity with self-given privacy ratings, that is, we tested similarity based on how users rate their own privacy inclinations. The data we use for this assumption is the response to three privacy-related questions users provide on their presession survey during data collection (the questions are adapted from the well-known privacy-index measures from Westin). Accordingly, we use Pearson similarity to find other users who are similar to this new user. With Pearson, we obtain an accuracy of 81.4 percent. We note however that 2-components accuracy is only about 1.77 percent of the missed policies, and even less 1-component. A similar result is obtained with Cosine similarity, where we achieved 82.56 percent accuracy, with again less than 2 percent accuracy for 2-components match and about 0.05 percent for 1-component. In sum, A3P social appears to be always superior to other methods. Note however that we cannot use A3P-social alone without A3P-core since the A3P-social does not factor in the evolution of an individual's privacy preferences.

Also A3P-social is more costly to be executed than A3P-core since the A3P-social analyzes information from a community rather than a single user.

6 CONCLUSION

We have proposed an Adaptive Privacy Policy Prediction (A3P) system that helps users automate the privacy policy settings for their uploaded images. The A3P system provides a comprehensive framework to infer privacy preferences based on the information available for a given user. We also effectively tackled the issue of cold-start, leveraging social context information. Our experimental study proves that our A3P is a practical tool that offers significant improvements over current approaches to privacy.

REFERENCES

[1] A. Acquisti and R. Gross, "Imagined communities: Awareness, information sharing, and privacy on the facebook," in Proc. 6th Int. Conf. Privacy Enhancing Technol. Workshop, 2006, pp. 36–58.
[2] R. Agrawal and R. Srikant,"Fast algorithms for mining association rules in large databases," in Proc. 20th Int. Conf. Very Large Data Bases, 1994, pp. 487–499.
[3] S. Ahern, D. Eckles, N. S. Good, S. King, M. Naaman, and R. Nair, "Over-exposed?: Privacy patterns and considerations in online and mobile photo sharing," in Proc. Conf. Human Factors Comput. Syst., 2007, pp. 357–366.
[4] M. Ames and M. Naaman, "Why we tag: Motivations for annotation in mobile and online media," in Proc. Conf. Human Factors Comput. Syst., 2007, pp. 971–980.
[5] A. Besmer and H. Lipford, "Tagged photos: Concerns, perceptions, and protections," in Proc. 27th Int. Conf. Extended Abstracts Human Factors Comput. Syst., 2009, pp. 4585–4590.
[6] D. G. Altman and J. M. Bland ,"Multiple significance tests: The bonferroni method," Brit. Med. J., vol. 310, no. 6973, 1995.
[7] J. Bonneau, J. Anderson, and L. Church, "Privacy suites: Shared privacy for social networks," in Proc. Symp. Usable Privacy Security, 2009.
[8] J. Bonneau, J. Anderson, and G. Danezis, "Prying data out of a social network," in Proc. Int. Conf. Adv. Soc. Netw. Anal. Mining., 2009, pp.249–254.
[9] H.-M. Chen, M.-H. Chang, P.-C. Chang, M.-C. Tien, W. H. Hsu, and J.-L. Wu, "Sheepdog: Group and tag recommendation for flickr photos by automatic search-based learning," in Proc. 16th ACM Int. Conf. Multimedia, 2008, pp. 737–740.
[10] M. D. Choudhury, H. Sundaram, Y.-R. Lin, A. John, and D. D. Seligmann, "Connecting content to community in social media via image content, user tags and user communication," in Proc. IEEE Int. Conf. Multimedia Expo, 2009, pp.1238–1241.
[11] L. Church, J. Anderson, J. Bonneau, and F. Stajano, "Privacy stories: Confidence on privacy behaviors through end user programming," in Proc. 5th Symp. Usable Privacy Security, 2009.
[12] R. da Silva Torres and A. Falc~ao, "Content-based image retrieval: Theory and applications," Revista de Inform_atica Te_orica e Aplicada, vol. 2, no. 13, pp. 161–185, 2006.
[13] R. Datta, D. Joshi, J. Li, and J. Wang, "Image retrieval: Ideas, influences, and trends of the new age," ACM Comput. Surv., vol. 40, no.2, p. 5, 2008.
[14] J. Deng, A. C. Berg, K. Li, and L. Fei-Fei, "What does classifying more than 10,000 image categories tell us?" in Proc. 11th Eur. Conf. Comput. Vis.: Part V, 2010, pp. 71–84. [Online]. Available: http://portal.acm.org/citation.cfm?id=1888150.1888157
[15] A. Kapadia, F. Adu-Oppong, C. K. Gardiner, and P. P. Tsang, "Social circles: Tackling privacy in social networks," in Proc. Symp. Usable Privacy Security, 2008.
[16] L. Geng and H. J. Hamilton, "Interestingness measures for data mining: A survey," ACM Comput. Surv., vol. 38, no. 3, p. 9, 2006.
[17] Image-net data set. [Online]. Available: www.image-net.org, Dec. 2013.
[18] S. Jones and E. O'Neill, "Contextual dynamics of group-based sharing decisions," in Proc. Conf. Human Factors Comput. Syst., 2011, pp. 1777–1786. [Online]. Available: http://doi.acm.org/10.1145/1978942.1979200
[19] A. Kaw and E. Kalu, Numerical Methods with Applications: Abridged., Raleigh, North Carolina, USA: Lulu.com, 2010.
[20] P. Klemperer, Y. Liang, M. Mazurek, M. Sleeper, B. Ur, L. Bauer, L. F. Cranor, N. Gupta, and M. Reiter, "Tag, you can see it!: Using tags for access control in photo sharing," in Proc. ACM Annu. Conf. Human Factors Comput. Syst., 2012, pp. 377–386

