Privacy-Aware Tag Recommendation for Image Sharing

Ashwini Tonge¹, Cornelia Caragea¹, Anna Squicciarini²

¹Kansas State University, Manhattan, KS-66506

²Pennsylvania State University, University Park, PA-16801

atonge@ksu.edu, ccaragea@ksu.edu, asquicciarini@ist.psu.edu

ABSTRACT

Image tags are very important for indexing, sharing, searching, and surfacing images with private content that needs protection. As the tags are at the sole discretion of users, they tend to be noisy and incomplete. In this paper, we present a privacy-aware approach to automatic image tagging, which aims at improving the quality of user annotations, while also preserving the images' original privacy sharing patterns. Precisely, we recommend potential tags for each target image by mining privacy-aware tags from the most similar images of the target image obtained from a large collection. Experimental results show that privacy-aware approach is able to predict accurate tags that can improve the performance of a downstream application on image privacy prediction. Crowd-sourcing predicted tags exhibit the quality of the recommended tags.

KEYWORDS

Privacy-aware tag recommendation; Image tagging; Image's privacy

ACM Reference Format:

Ashwini Tonge¹, Cornelia Caragea¹, Anna Squicciarini². 2018. Privacy-Aware Tag Recommendation for Image Sharing. In *HT '18: 29th ACM Conference on Hypertext and Social Media, July 9–12, 2018, Baltimore, MD, USA*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3209542.3209574>

1 INTRODUCTION

Images are constantly shared on social networking sites such as Facebook, Flickr, and Instagram. For instance, it is common to take photos at cocktail parties and upload them to social networking sites without much hesitation for self-promotion and personal sharing. However, when privacy settings are used inappropriately, these photos can potentially reveal a user's personal and social habits, resulting in unwanted disclosure and privacy violations [1, 20, 21, 33]. For example, malicious attackers can take advantage of these accidental leaks to launch context-aware or even impersonation attacks. Thus, several works [20–22, 27, 28, 30, 33] have been developed in an attempt to provide appropriate privacy settings for online images. Prior works on privacy prediction [20, 27, 29, 33] found that the tags associated with images are indicative of their sensitive content. Tags are also important for image-related applications such as indexing, sharing, searching, content detection and social discovery [5, 7]. Yet, the tags are at the sole discretion of users, and they tend

(a) *Private*: Style, Skirt, Corporate
Pretty, Girl, Woman, Elegant

(b) *Public*: Sabrina, Celebrity, News
Famous, Woman, Hollywood

Figure 1: Anecdotal evidence for privacy-aware user tags.

to be noisy and incomplete [25]. Despite that many approaches to automatic image tagging have been developed [6, 12–14], none of these works considers the privacy aspect of an image while making the annotations and hence would not be sufficient for identifying images' private (or sensitive) content.

We posit that visually similar images can possess very different sets of tags if these images have different privacy orientations. For example, Figure 1 shows anecdotal evidence obtained from a Flickr dataset in which visually similar images of private and public classes display different sets of user tags. The picture of a woman that belongs to the private class in Figure 1(a) contains tags such as “Elegant,” “Corporate,” “Style,” and “Pretty,” whereas the picture of a woman that belongs to the public class in Figure 1(b) contains tags such as “Celebrity,” “Famous,” “News,” and “Hollywood.” Images are considered private if they belong to the private sphere (portraits, family, friends, home) or contain information that can not be shared (e.g., private documents) [33]. Figure 1 shows that the tags are correlated to image's privacy patterns [9, 22, 23] and are effective when access to the image content is not allowed since users may be reluctant to share the real images (revealing user's identity through the face, and friends, etc.) for visual content analysis. In such cases, privacy-aware tags can become good indicators of the privacy settings and improve the privacy prediction methods.

To this end, we ask the following questions and address them with our research agenda: *Can we develop an automated approach to recommend accurate image tags that can also take into account the sharing needs of the users for images in question? Can we make precise tag recommendations for newly uploaded images that have an incomplete set of user tags or no tags at all? Can these recommended tags help improve the privacy prediction performance?*

Contributions and Organization. We present a privacy-aware approach to image tagging¹, aimed at improving the quality of user tags, while also preserving the images' original privacy sharing patterns. Precisely, our approach recommends, based on collaborative filtering, potential tags for a target image by mining privacy-aware tags from the most similar images of a target image from a large collection. To evaluate the recommended tags, we employ crowd-sourcing to identify relevancy of the suggested tags to images. The results show that, although the user-input tags are noisy or incomplete, our approach can recommend accurate tags. We also

¹The code is given at <https://github.com/ashwinitonge/privacy-aware-tag-rec.git>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

HT '18, July 9–12, 2018, Baltimore, MD, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5427-1/18/07...\$15.00

<https://doi.org/10.1145/3209542.3209574>

investigate tag recommendation in a binary privacy setting, and show that the predicted tags can exhibit relevant cues for specific privacy settings (*public* or *private*) that can be used to improve the privacy prediction performance.

2 RELATED WORK

We briefly review the related work as follows.

Automatic Image Annotation: Many works on automatic image annotation have been proposed [4, 6, 10–14]. For example, Chen et al. [14] proposed an approach to image tagging, which learned two classifiers to predict tags: one that reconstructs the complete tag set from the tags available during training and the other that maps image features to the reconstructed tag set. Several works on tag recommendation for social discovery [2, 18] and image classification [2, 17, 32] in photo sharing sites (e.g. Flickr) typically trained classifiers for each tag using image’s textual and/or visual features.

Collaborative Filtering: Our approach draws ideas from collaborative filtering (CF), and hence, we briefly review the most relevant works on CF here. Xu et al. [31] designed a CF approach to suggest high-quality tags for Web objects, according to several criteria (coverage, popularity, effort, uniformity). Authors consider that if two tags frequently co-occur when describing a specific object, they should also co-occur in the recommended set of tags. Recently, Peng et al. [15] generated joint item–tag recommendations for users, where the tags represent topics from an item (i.e., a web resource) in which the user may be interested.

Online Image Privacy: Several works analyzed users’ posted data with respect to privacy. For example, Ahern et al. [1] studied the effectiveness of location information and tags in predicting privacy settings of images. They also conducted a study to verify whether the visual features are relevant to an image’s privacy and found that content is one of the discriminatory factors affecting image privacy, especially for images depicting people. This supports the core idea underlying our work: that tags depicting private categories obtained from image content are pivotal for identifying the sensitive content from the search results. For example, tags such as “wedding,” “bride,” “people” describing a wedding event (private category) represent the private class. Jones and O’Neill [8] determined that people are more reluctant to share photos capturing social relationships than photos taken for functional purposes; certain settings such as work, bars, concerts cause users to share less. Zerr et al. [33] developed the PicAlert dataset to help detect private images. Recently, Tonge et al. [27, 29] showed the performance of automatically obtained image tags from the visual content using convolutional neural networks (CNN) for privacy prediction. Yet, these tags depicted objects or scenes given in the image and failed to capture the privacy characteristics of the image while generating the tags. To this end, we recommend privacy-aware tags for online images that have the potential to improve the set of user tags.

3 PRIVACY-AWARE TAG RECOMMENDATION

Our approach to recommending privacy-aware tags for newly posted images on content sharing websites is inspired from collaborative filtering (CF) [19]. Many images posted on the Web in recent years, facilitate the study of potential relationships between images. We leverages these relationships to exchange privacy-aware

Algorithm 1 Tag Recommendation

```

1: Input: A dataset  $\mathcal{D} = \{I_1, \dots, I_n\}$  of images and their tags
    $\{T_1, \dots, T_m\}$ ; a target image  $I$  and its tags  $T$ ;  $k$  the nearest neighbors
   of  $I$  from  $\mathcal{D}$ ;  $r$  the number of tags to recommend.
2: Output: A set  $R$  of recommended tags for  $I$ .
3:  $R \leftarrow \phi$ ; // the set of recommended tags, initially empty.
4:  $S \leftarrow \phi$ ;
5: if  $T = \phi$  then // if the set of tags is empty.
6:    $\mathbf{x} \leftarrow \text{ImageContentEncoding}(I)$ ; // deep features for  $I$ 
7:   for all  $I_j \in \mathcal{D}$  do
8:      $\mathbf{x}_j \leftarrow \text{ImageContentEncoding}(I_j)$ ; // deep features  $I_j$ 
9:      $s_j \leftarrow \text{similarity}(\mathbf{x}, \mathbf{x}_j)$ ; // visual content similarity
10:     $S \leftarrow S \cup (I_j, s_j)$ ; // store  $I_j$  and its similarity with  $I$ 
11:   end for
12: else
13:    $\mathbf{x} \leftarrow \text{ImageTagEncoding}(I)$ ; // get tags’ features of  $I$ 
14:   for all  $I_j \in \mathcal{D}$  do
15:      $\mathbf{x}_j \leftarrow \text{ImageTagEncoding}(I_j)$ ; // get tags’ features of  $I_j$ 
16:      $s_j \leftarrow \text{similarity}(\mathbf{x}, \mathbf{x}_j)$ ; // compute the tags similarity
17:      $S \leftarrow S \cup (I_j, s_j)$ ; // store  $I_j$  and its similarity with  $I$ 
18:   end for
19: end if
20:  $S.\text{similarities.sort}()$ ; // sort images in decreasing order of similarity
21:  $S \leftarrow \text{top } k(I_j, s_j) \text{ entries}$ ; // get  $k$  images with the highest similarities
22:  $W \leftarrow \text{TagRanking}(S)$ ; // rank the tags from  $S$  images
23:  $R \leftarrow r$  tags with the highest scores from  $W$ ;
24: return  $R$ 

```

tags between similar images. The analogy with conventional CF methods is that images correspond to users and tags correspond to items. We base our models on the assumption that *privacy-aware similar images possess similar tags*.

Algorithm 1 describes the process in detail. Recommendations are made for the target image based on the neighboring images’ tags (as a privacy-aware weighted sum of occurrences of tags). A common problem in CF is the *cold start* problem [24]. In our case, this refers to images that have very few tags or no tags at all, and hence, there is not enough information available to find accurate nearest neighbors for a target image, based on tags. However, in our domain, images can be represented using two views: (1) visual content; and (2) tags. We take advantage of both the views. The input of the algorithm is a dataset $\mathcal{D} = \{I_1, \dots, I_n\}$ of images and their tags, $\{T_1, \dots, T_m\}$; a target image I and its set of tags T , which could be empty; k the number of nearest neighbors of I from $\mathcal{D}$; and r the number of tags to recommend. The output of the algorithm is a ranked list of r recommended tags for the target image. The algorithm starts by checking if the set of tags T of the target image I is empty (Alg. 1, line 5). If $T \neq \phi$, the similarities between I and all images in $\mathcal{D} \setminus \{I\}$ are computed based on images’ tags (Alg. 1, lines 13–18). The top k most similar images to I are returned (Alg. 1, lines 20–21) and the candidate set that represents the union of the sets of tags extracted from these k similar images is ranked inside the subroutine for tag ranking (line 22) described in Algorithm 2. The highly ranked r tags from the candidate set are returned as recommended tags for the target image I (Alg. 1, line 23–24). If $T = \phi$ for image I , Alg. 1 recommends r tags based on the similarity computed using image content features (Alg. 1, lines 5–12). For each tag in the candidate set, we compute its score (or weight) as the privacy-aware sum of similarities between the target image and its neighbors (Alg. 2, lines 6–12). This weighting is based on the assumption that a “good” tag is likely to be exchanged

Algorithm 2 Tag Ranking

```

1: function TagRanking( $S$ )
2:    $W \leftarrow \phi$ ; // the set of tags and their scores, initially empty.
3:   for all  $I_j \in S$  do
4:      $T_j \leftarrow I_j.\text{tags}$  // get the set of tags of image  $I_j$ .
5:      $s_j \leftarrow I_j.\text{similarity}$  // similarity of target image and  $I_j$ .
6:     for all  $t \in T_j$  do
7:        $w_t \leftarrow W.\text{scoreOf}(t)$  //  $w_t$  stores the score of  $t$ 
8:       if  $w_t = \text{null}$  then // if tag  $t$  is not in  $W$  already
9:          $W \leftarrow W \cup (t, 0)$  // add  $t$  to  $W$ 
10:      end if
11:       $w_t \leftarrow w_t + s_j \cdot P(t|pr)$  //score of  $t$  weighted by privacy
12:    end for
13:  end for
14:   $W.\text{scores.sort}()$  // sort the scores in  $W$  in the decreasing order.
15:  return  $W$ .
16: end function

```

between similar images. The weight of a tag t , w_t , is computed as:

$$w_t = \sum_{j \in S} c_{jt} \cdot s_j \cdot P(t|pr(I)) \quad (1)$$

where S represents the k most similar images of I from $\mathcal{D}$, c_{jt} is 1 if tag t belongs to the tag set T_j of image I_j from S and 0 otherwise, and s_j is the similarity between image I_j and I . The probability $P(t|pr(I))$ is the likelihood of the tag t belonging to the privacy class (i.e., public or private) of the target image I . For instance, if the target image I is of private class then $P(t|pr(I))$ gives the probability of tag t belonging to the set of private images. The likelihood is calculated based on the dataset $\mathcal{D}$. We rely on the privacy likelihood of the tag instead of considering privacy as another parameter (referred as privacy-enforced similarity) in the image similarity because we desire privacy-aware tags without missing out on the high-quality tags. For example, using privacy-enforced similarity, for Figure 1(b) (given its public nature), tags such as “women,” “girl” (inclined to private class) would not be suggested, whereas privacy-aware weights can obtain descriptive tags for both the image’s content and privacy aspect of the image.

4 DATASET

Similar to prior works [20, 27, 33], that identified generic privacy patterns using tags, we verify if the recommended tags are indicative of the privacy classes and also validate their relevancy to the images’ content. Thus, we evaluate the algorithm on Flickr images sampled from the PicAlert dataset [33]. PicAlert contains images on various subjects, which are manually labeled as *private* or *public*. We split the dataset into three subsets. The first two subsets, denoted as DS_1 and DS_2 , for which we recommend tags, contain randomly sampled 3,689 and 500 images, respectively. The third dataset, $PicAlert_{8K}$, is a collection $\mathcal{D}$ of 8,000 images, labeled as private or public, that are used to recommend tags for the target images in DS_1 and DS_2 . The ratio of public to private images in all datasets is 3 : 1. For privacy prediction, we use DS_1 to train Support Vector Machine (SVM) models on the recommended tags and use DS_2 to test these models. For each image I in DS_1 and DS_2 , we randomly split its set of tags into two subsets (i.e., *visible* and *hidden*). The motivation behind using random split is that newly uploaded

image may have an incomplete and noisy set of user-input tags [25]. For both DS_1 and DS_2 , we consider images with a number of user tags greater than 10 to have at least five visible tags to calculate an accurate similarity. After filtering the stop words, numbers, and URL, the size of the vocabulary is $\approx 19,000$.

5 EXPERIMENTS AND RESULTS

We evaluate the tags obtained by the proposed algorithm for images in DS_1 and DS_2 , by transferring tags from their most similar images from $PicAlert_{8K}$ in two settings: 1) *whether these tags hint to specific image privacy settings*; and 2) *whether these tags are good enough to describe the content of an image*. Hence, we adopt two evaluation mechanisms: 1) we examine the performance of models trained on the recommended tags combined with the original tags (when available) for privacy prediction to determine their ability in identifying private content for online image sharing; and 2) we compare the suggested tags against the ground-truth, i.e., the *hidden* set of tags, and also evaluate their quality through crowd-sourcing.

Evaluation Setting. We generate five subsets of visible and hidden tags and report performance averaged over these five splits. For privacy prediction, we use SVM Weka implementation and Boolean features for tags, i.e., 1 if a tag is present and 0 otherwise.

5.1 Evaluation by Privacy Prediction

We study Alg. 1 in the setting where each image in DS_1 has a seed set of tags associated with it, i.e., $T \neq \phi$ (Alg. 1, lines 13-18). The similarity between images is computed between the *visible* tag set of a target image and all available tags from an image in $PicAlert_{8K}$. We experiment with $k = 2, \dots, 10$ and $r = 5, \dots, 20$, where k is the number of similar images, and r is the number of recommended tags (see Alg. 1). We show results for the best value of k i.e. $k = 10$.

Table 1 shows the performance obtained by the SVM models trained on the combination of recommended

tags (rt) and visible tags (vt) (as we increase rt from 5 to 20) for the images in DS_1 and evaluated on the fixed set of visible tags of the images in DS_2 (for consistency). The

Features	Acc.%	F1	Pre.	Re.
vt	74.83	0.743	0.739	0.748
$vt \& rt(5)$	78.20	0.772	0.762	0.783
$vt \& rt(10)$	77.80	0.765	0.754	0.777
$vt \& rt(15)$	77.92	0.767	0.758	0.778
$vt \& rt(20)$	77.43	0.758	0.745	0.771

Table 1: Evaluation by privacy prediction, $T \neq \phi, k = 10$.

results show that the performance of privacy prediction improves when we add recommended tags to the set of visible tags. We get the best performance for $r = 5$ of F1-score of 0.772, whereas the SVM trained on only visible tags achieves 0.743 F1-measure, yielding an improvement of 3% in overall performance. We notice that generally, the performance increases with the decreasing value of r (best performance is given by $r = 5$ and $k = 10$). Due to the diverse nature of the data and a large vocabulary, a large r may introduce noise in the results. In the following experiments, we use $k = 10$.

5.2 Solution to the Cold Start Problem

Cold start is a challenging problem particularly in many CF approaches, where the absence of items (i.e., tags, in our case) that are used to bootstrap the algorithms may theoretically hinder the

recommendations to be produced. Hence, we evaluate our approach in the setting where we assume that each image in DS_1 has no tags, i.e., $T = \phi$ and recommend tags from visually similar images (Alg. 1, lines 5-12). The similarity between two images is given as the cosine similarity of the corresponding feature vectors. We consider two types of image features extracted from a GoogLeNet CNN [26]: 1) *deep visual feature*, and 2) *deep tags*, due to their prior performance for privacy prediction [27, 30]. We extract visual features $pool_5$ from the layer named as “pool5/drop_7x7_s1”. For deep tags, we use the probability distribution over 1,000 object categories for the input image obtained by applying the softmax function over the last fully-connected layer of the CNN. We consider the top k objects of highest probabilities as *deep tags*. We use the pre-trained GoogLeNet on a subset of the ImageNet dataset [16], which is distributed with the CAFFE framework for CNN [3].

Table 2 shows the privacy prediction performance obtained by the SVM trained on the privacy-aware tags recommended from visually similar images based on $pool_5$ ($pool_5(rt)$) and deep tags (DT(rt)) for the images in DS_1

Features	Acc.%	F1	Pre.	Re.
$pool_5(rt)$	75.74	0.743	0.729	0.757
DT(rt)	74.19	0.731	0.725	0.742
vt	74.83	0.743	0.739	0.748
DT	68.54	0.645	0.619	0.685

Table 2: Visual content similarity ($k = 10$).

and evaluated on the visible tags of the images in DS_2 . The table also shows the performance of the models trained on visible tags alone (vt), if they would be available, and predicted deep tags (DT) of DS_1 , as done in prior work [27]. The results show that the models trained on the recommended tags yield similar results to the models trained on visible tags (user-input tags – if we would know them). We obtain the best F1-score of 0.743 and recall of 0.757 with recommended tags $r = 5$. We observe that the models trained on tags recommended from visually similar images based on $pool_5$ ($pool_5(rt)$) outperform those trained on tags recommended from visually similar images based on deep tags (DT(rt)), which in turn, outperform the models trained on the deep tags (DT). The reason is that deep tags belong to only 1,000 objects due to which many relevant tags (e.g. “walking” and “culture”) can not be captured.

5.3 Quality Assessment of Recommended Tags

In the above experiments, we evaluated the effectiveness of recommended tags for privacy prediction. In this experiment, we determine whether the recommended tags describe an image’s content appropriately. We compare the tags recommended using our privacy-aware weighting scheme against the ground-truth (i.e., *hidden* set of tags). Table 3 shows the performance (Precision@ r) obtained for $r \in \{1, 2, 3, 4, 5, 10\}$ tags recommended for the images in DS_2 when compared against the gold-standard set (GS) of tags (those are hidden from the original user tags). We compute Precision as the total number of *recommended* and *relevant* tags over the number of tags recommended (i.e., r). The results show that the recommended tags achieve precision as high as 0.181 using gold-standard. The gold-standard set is nothing but a subset of user annotated tags, which may not provide all possible tags related to the image content. Hence, the gold-standard set may fail to capture

Visible	Hidden	Recommended Tags	
Beauty	Geisha	People	<i>Culture</i>
Light	Kyoto	Japan	<i>Street</i>
Travel	Japan	Asia	<i>Walking</i>
Couple	Kimono	Geisha	
Woman	Traditional	Kimono	
Vintage	Asia	Kyoto	
	People	Traditional	

Figure 2: Image with recommended tags, $r=10$.

highly relevant tags provided by the recommendation strategy. For example, in Figure 2, tags relevant to the image content (shown in *italic*) are recommended, but do not appear in the user-input tags.

Crowd-sourcing can be used to address the above limitation.

We employ crowd-sourcing as follows: we

use two annotators from the Amazon Mechanical Turk to determine if the recommended tags are relevant to the image content. For each tag, annotators were asked to choose between: *relevant*, *irrelevant* and *not sure*. To calculate precision values, we consider a tag as *Relevant* if at least one annotator marked it as *relevant*, i.e., the tags can be subjective and one annotator can observe more in an image than the other. Table 3 shows the performance obtained through crowd-sourcing (CS). Note that the results of crowd-sourcing are higher than those obtained by relying only on *gold standard* to compute the performance. Precisely, through crowd-sourcing, the precision increased from 0.181 (GS) to 0.855, reassuring that the generated tags are relevant to image’s content. The difference in the results can be justified as user tags tend to be noisy, incomplete, and may not relate to the image content [25].

r	GS	CS
1	0.177	0.855
2	0.181	0.761
3	0.181	0.755
4	0.172	0.703
5	0.174	0.691
10	0.155	0.633

Table 3: Quality evaluation of suggested tags.

6 CONCLUSIONS

We proposed privacy-aware image tagging, based on collaborative filtering, that can improve the original user-input tags while preserving the images’ privacy. Although user tags are prone to noise, we were able to integrate them in our approach and recommend accurate tags. Importantly, we simulated the recommendation strategy for newly-posted images, which had no tags attached. This is a particularly challenging problem, as in many CF approaches, the absence of items (tags in our case) may theoretically hinder the recommendations to be produced, due to the lack of enough information available to find similar images to a target image. We achieve better performance for privacy prediction with recommended tags than the original set of user tags, which in turn indicate that the suggested tags comply to an image’s privacy. We also conducted a user evaluation of recommended tags to inspect the quality of our privacy-aware recommended tags. The results show that the proposed approach is able to recommend highly relevant tags. In future, it would be interesting to study the algorithm for multiple sharing needs of users such as friends and family, by considering privacy likelihood with respect to multi-class privacy settings.

7 ACKNOWLEDGMENTS

This research is supported by NSF grant 1421970 and 1421776.

REFERENCES

- [1] Shane Ahern, Dean Eckles, Nathaniel S. Good, Simon King, Mor Naaman, and Rahul Nair. 2007. Over-exposed?: privacy patterns and considerations in online and mobile photo sharing. In *CHI '07*.
- [2] Hong-Ming Chen, Ming-Hsiu Chang, Ping-Chieh Chang, Ming-Chun Tien, Winston H. Hsu, and Ja-Ling Wu. 2008. SheepDog: group and tag recommendation for flickr photos by automatic search-based learning. In *MM '08*.
- [3] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. 2013. DeCAF: A Deep Convolutional Activation Feature for Generic Visual Recognition. *CoRR* (2013).
- [4] Yansong Feng and Mirella Lapata. 2008. Automatic Image Annotation Using Auxiliary Text Information. In *ACL-08: HLT*. Columbus, Ohio.
- [5] Yue Gao, Meng Wang, Huanbo Luan, Jialie Shen, Shuicheng Yan, and Dacheng Tao. 2011. Tag-based Social Image Search with Visual-text Joint Hypergraph Learning. In *MM '11*. ACM, New York, NY, USA, 1517–1520.
- [6] Matthieu Guillaumin, Thomas Mensink, Jakob Verbeek, and Cordelia Schmid. 2009. TagProp: Discriminative Metric Learning in Nearest Neighbor Models for Image Auto-Annotation. In *ICCV*.
- [7] Livia Hollenstein and Ross Purves. 2010. Exploring place through user-generated content: Using Flickr tags to describe city cores. *J. Spatial Information Science* 1, 1 (2010), 21–48.
- [8] Simon Jones and Eamonn O'Neill. 2011. Contextual dynamics of group-based sharing decisions (*CHI '11*). 10. <https://doi.org/10.1145/1978942.1979200>
- [9] Peter Klemperer, Yuan Liang, Michelle Mazurek, Manya Sleeper, Blase Ur, Lujo Bauer, Lorrie F. Cranor, Nitin Gupta, and Michael Reiter. [n. d.]. Tag, you can see it! Using tags for access control in photo sharing. In *CHI '12*.
- [10] Victor Lavrenko, R. Manmatha, and Jiwoon Jeon. 2004. A Model for Learning the Semantics of Pictures. In *NIPS '04*. MIT Press.
- [11] Wee Leong, Rada Mihalcea, and Samer Hassan. 2010. Text Mining for Automatic Image Tagging (*COLING '10*).
- [12] Jing Liu, Mingjing Li, Qingshan Liu, Hanqing Lu, and Songde Ma. 2009. Image Annotation via Graph Learning. *PR* (Feb. 2009), 11.
- [13] Ameesh Makadia, Vladimir Pavlovic, and Sanjiv Kumar. 2008. *A New Baseline for Image Annotation*. Springer Berlin Heidelberg, 316–329.
- [14] Kilian Q. Weinberger, Minmin Chen, Alice Zheng. 2013. Fast Image Tagging. In *ICML*.
- [15] Jing Peng, Daniel Dajun Zeng, Huimin Zhao, and Fei-yue Wang. 2010. Collaborative Filtering in Social Tagging Systems Based on Joint Item-tag Recommendations. In *CIKM '10*. ACM, 809–818. <https://doi.org/10.1145/1871437.1871541>
- [16] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2014. ImageNet Large Scale Visual Recognition Challenge.. In *arXiv:1409.0575*.
- [17] Jose San Pedro and Stefan Siersdorfer. 2009. Ranking and classifying attractiveness of photos in folksonomies (*WWW '09*). ACM, NY, USA.
- [18] Neela Sawant. 2011. Modeling tagged photos for automatic image annotation (*MM '11*). ACM, 2.
- [19] Yue Shi, Martha Larson, and Alan Hanjalic. 2014. Collaborative Filtering Beyond the User-Item Matrix: A Survey of the State of the Art and Future Challenges. *ACM Comput. Surv.*, Article 3 (May 2014), 45 pages. <https://doi.org/10.1145/2556270>
- [20] Anna Squicciarini, Cornelia Caragea, and Rahul Balakavi. 2014. Analyzing Images' Privacy for the Modern Web (*HT '14*). ACM, NY, USA, 136–147.
- [21] Anna Squicciarini, Cornelia Caragea, and Rahul Balakavi. 2017. Toward Automated Online Photo Privacy. *ACM Tran. on the Web* 11, 1, Article 2 (2017).
- [22] Anna Squicciarini, Andrea Novelli, Dan Lin, Cornelia Caragea, and Haoti Zhong. 2017. From Tag to Protect: A Tag-Driven Policy Recommender System for Image Sharing. In *PST '17*.
- [23] Anna Squicciarini, Smitha Sundareshwaran, Dan Lin, and Josh Wede. 2011. A3p: adaptive policy prediction for shared images over popular content sharing sites. In *HT '11*. ACM, 261–270.
- [24] Xiaoyuan Su and Taghi M. Khoshgoftaar. 2009. A Survey of Collaborative Filtering Techniques. *Adv. in AI*, Article 4 (2009), 1 pages.
- [25] H. Sundaram, L. Xie, M. De Choudhury, Y.R. Lin, and A. Natsev. 2012. Multimedia Semantics: Interactions Between Content and Community. *IEEE* 100, 9 (2012).
- [26] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2014. Going Deeper with Convolutions. *CoRR* abs/1409.4842 (2014).
- [27] Ashwini Tonge and Cornelia Caragea. 2016. Image Privacy Prediction Using Deep Features. In *AAAI '16*.
- [28] Ashwini Tonge and Cornelia Caragea. 2018. On the Use of "Deep" Features for Online Image Sharing. In *Companion Proceedings of The Web Conf.* 1317–1321.
- [29] Ashwini Tonge, Cornelia Caragea, and Anna Squicciarini. 2018. Uncovering Scene Context for Predicting Privacy of Online Shared Images. In *AAAI '18*.
- [30] Lam Tran, Deguang Kong, Hongxia Jin, and Ji Liu. 2016. Privacy-CNH: A Framework to Detect Photo Privacy with Convolutional Neural Network Using Hierarchical Features. In *AAAI '16*.
- [31] Zhichen Xu, Yun Fu, Jianchang Mao, and Difu Su. 2006. Towards the semantic web: Collaborative tag suggestions. In *Collaborative Web Tagging Workshop*.
- [32] J. Yu, D. Joshi, and J. Luo. 2009. Connecting people in photo-sharing sites by photo content and user annotations. In *ICME 2009*. IEEE.
- [33] Sergej Zerr, Stefan Siersdorfer, Jonathon Hare, and Elena Demidova. 2012. Privacy-aware image classification and search. In *ACM SIGIR*. ACM, NY, USA.