

Privacy Leakage through Innocent Content Sharing in Online Social Networks

Maria Han Veiga
Inst. of Computational Science
University of Zurich, Switzerland
hmaria@physik.uzh.ch

Carsten Eickhoff
Dept. of Computer Science
ETH Zurich, Switzerland
ecarsten@inf.ethz.ch

ABSTRACT

The increased popularity and ubiquitous availability of online social networks and globalised Internet access have affected the way in which people share content. The information that users willingly disclose on these platforms can be used for various purposes, from building consumer models for advertising, to inferring personal, potentially invasive, information.

In this work, we use Twitter, Instagram and Foursquare data to convey the idea that the content shared by users, especially when aggregated across platforms, can potentially disclose more information than was originally intended.

We perform two case studies: First, we perform user de-anonymization by mimicking the scenario of finding the identity of a user making anonymous posts within a group of users. Empirical evaluation on a sample of real-world social network profiles suggests that cross-platform aggregation introduces significant performance gains in user identification.

In the second task, we show that it is possible to infer physical location visits of a user on the basis of shared Twitter and Instagram content. We present an informativeness scoring function which estimates the relevance and novelty of a shared piece of information with respect to an inference task. This measure is validated using an active learning framework which chooses the most informative content at each given point in time. Based on a large-scale data sample, we show that by doing this, we can attain an improved inference performance. In some cases this performance exceeds even the use of the user's full timeline.

Keywords

Privacy, Social networks, Information value

1. INTRODUCTION

User privacy is a topic that has increasingly gained traction with the rise of online social networks (OSN). These platforms allow users to communicate, connect with peers and share content. Originally, OSNs mainly focused on

these core aspects, but nowadays the term also includes platforms which are primarily user-centric, allowing members to broadcast personal thoughts and content. In 2010, [13] find OSNs among the most frequently visited Web sites for a large population of users. Due to their prevalence and abundance in personal content, OSNs lend themselves to the study of human behavior at scale [14].

Recent successful initial public offerings (IPO) and high market valuations underline the monetary value of OSNs. However, the relation between the number of registered users, their online activity, and these valuations is not entirely clear. It has been shown in several studies that user characteristics, such as personality traits [12] or future route intentions [15], can be reliably inferred from corresponding OSN profiles. Although the general value of personal data is widely accepted, there have not been many studies which assign a tangible value to OSN profiles. As a consequence, both for users as well as platform providers, the value of information remains a vague notion, at best. This situation is detrimental both to users who cannot be expected to make informed decisions about privacy controls, as long as they do not know the value and potential risk of disclosing a given information item, as well as platform and service providers who blindly buy and sell user data in bulk instead of saving resources by concentrating on select relevant portions of information.

In this paper, we aim to draw attention to accidental privacy leakage through content sharing in online social networks and make a first step toward describing a formal metric of task-specific informativeness of pieces of shared content.

Our empirical study relies on three popular OSN platforms: (1) Twitter, a microblogging platform whose main content comes in *tweets*, posts limited to 140 characters which can contain text, media (video or images), links to external Web sites, references to other users and *hashtags* (terms starting with the # symbol, which are used to mark keywords or topics in a tweet). (2) Instagram, a photo sharing platform. Its main content are visual in nature along with optional textual descriptors. (3) Foursquare, a location service platform concentrating on the notion of *check-ins*. Check-ins correspond to real-world venues that the user has visited. In addition to the venue name, more information such as location and venue categories are available.

Our investigation is driven by the following research questions:

1. How well can we uniquely identify a user based on matching a set of unseen posts to a user's online foot-

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.

SIGIR PIR '16 Pisa, Italy

© 2016 ACM. ISBN 123-4567-24-567/08/06...\$15.00

DOI: 10.475/123_4

print and is there a benefit in modelling user identities across more than one OSN?

2. For the same user, is the information posted in one OSN indicative of the information contained in another OSN?
3. Can we quantify the amount of new information that a piece of shared content carries, with respect to a concrete inference task?

In particular, to address RQ 1, we mimic the following scenario: let us have a collection of users' online footprints and a set of q anonymous posts. We test whether we are able to correctly find the author of the anonymous posts based on seeing part of their online footprint.

As for RQ 2 and 3, we consider that a user may unintentionally expose personal information through seemingly innocuous shared content. For example, when a user shares a venue check-in, it is easy to infer which venue category was visited. However, a post which does not mention a place explicitly might still contain information about a potential behaviour or visit intention. Consider two tweets from the same user: "Lol should start heading to the gym #fitness" and "What a great sunny day!". It is clear that the first tweet contains more information about the user's intention to visit a venue type than the second. To this end, we devise an informativeness metric for shared content. The metric explicitly models the item's relevancy towards a given inference task as well as its novelty in comparison to the previously seen timeline. Such a score can serve as an indication of the amount of novel information disclosure associated with an information item and can, in the future help both service providers as well as privacy advocates in making informed decisions.

This paper makes three novel contributions beyond the state of the art:

1. We formulate a scoring function to quantify the information value of shared content from the perspective of improving the performance of a concrete inference task.
2. We present a practical way of exploiting the theoretical model by integration into an active learning scheme that results in enhanced user model learning rates.
3. In both practical settings, we put particular emphasis on cross-platform models of user identity

The remainder of this paper is structured as follows: Section 2 gives an overview of related work dedicated to user modelling as well as privacy protection in Web scenarios. On the basis of a parallel corpus of OSN profiles belonging to the same natural person, Section 4 discusses unique user identification methods employing intra-platform as well as cross-platform information. Section 5 formally describes an informativeness score for shared OSN content and applies it to the task of predicting user traits manifested on one platform based on the user's activity on other OSNs. Finally, Section 6 concludes with a brief discussion of our main findings as well as an outlook on ongoing and future extensions to this work.

2. RELATED WORK

There is a wealth of work dedicated to privacy protection in Web information systems such as search engines or social networks. A number of early studies investigate the common privacy concerns of information system users [17, 18, 16, 9], finding that general concerns are abundant among Internet users but remain vague and imprecise. Many users are aware of the information collection and behavioral profiling activities undertaken by service providers as well as the wide range of data-driven inference efforts that have been presented by the academic and industrial communities [15, 12]. In spite of this knowledge, however, even technology-affine users cannot reliably quantify the exact risks entailed by careless information disclosure.

Privacy concerns become especially prevalent in mobile computing environments [11]. De Montjoye *et al.* [6] show that as few as four hourly GPS samples are enough to uniquely identify 95% of all individuals in a 500k-user phone log. We encounter an even greater potential for privacy hazards in settings that go beyond raw positional traces, joining them with topical information, *e.g.*, in Web search queries [19] or contextual advertising [1].

To counter such de-anonymization and tracking efforts, various strategies have been proposed. Dwork's concept of differential privacy [7] considers adding ϵ -noise to aggregate queries that prevents singling out individual contributions to the overall aggregate. Similarly, Carpineto and Romano [4] rely on the notion of k -anonymity, ensuring that no query should return less than k individual records. In the domain of personal information, these approaches may not go far enough since certain, frequent, characteristics that would neither be detected under k -anonymity nor differential privacy could cause severe privacy hazards.

This paper, in spirit, follows the reasoning of Howard *et al.* [10] by measuring not just the amount of information contained in a given message, but also with respect to an inference task which can be economically relevant. In this way it attributes an economic dimension to messages, which can be an interesting measure both for industry players as well as for the message's original author. On the basis of a number of concrete classification tasks, this paper aims to close the gap between the rich body of work on empirical analysis of privacy hazards on the one hand and the large range of available privacy protection measures on the other. We argue that only by understanding the concrete implications of information disclosure (*e.g.*, in the form of the value of a piece of information) can users be expected to make educated decisions about the appropriate protection measures they are willing to take.

3. DATASET

As our research questions are concerned with the relationship between parallel user profiles of the same natural person across different OSNs, we rely on the methodology described in [8] to assemble our dataset. We obtain a collection of 618 distinct users who cross-post content from corresponding profiles in multiple social networks, totalling 1.1 million tweets, 18000 Instagram posts and 99000 Foursquare check-ins.

4. USER DE-ANONYMIZATION

Our first use case is concerned with, given a number of

anonymous social media posts q and a collection of users \mathcal{U} , finding the particular user $u \in \mathcal{U}$ that authored the posts. Inspired by general text matching strategies [2], each user’s known previous posts are described in the form of a unigram language model M_u and the likelihood of said user having authored the anonymous text q corresponds to $p(M_u|q)$. Using Bayes’ law, one can write:

$$p(u|q) = \frac{p(q|u)p(u)}{p(q)} \quad (1)$$

And to select the most likely user:

$$\arg \max_{u \in \mathcal{U}} p(q|u) \quad (2)$$

To simplify the expression further, we assume that $p(q)$ is constant for all users and treat $p(u)$ as uniform across all $u \in \mathcal{U}$. Thus, we find the most likely user by estimating $p(q|u)$, the probability of posts q being generated by the language model derived from u ’s available timeline.

These timelines are projected into a n-dimensional TF-IDF weighted vector space. To preserve the natural way in which users write, no further vocabulary pre-processing (such as lemmatization or exclusion of less common words) was applied. Based on this representation, we estimate $p(q|u)$ as the product across all terms t in the vocabulary:

$$p(q|u) \propto \prod_{t \in V} p(t|M_u) \quad (3)$$

As described in Section 3, the dataset contains the on-line footprint of the same user on Twitter, Instagram and Foursquare. To mimic the described task, textual data from one OSN is used as the source of anonymous posts and the textual data from the remaining two OSNs is used to generate the user language model $p(M_u|q)$.

To generate the training data, randomly sampled sections of varying length from the training source are used to generate pairs of the form $(M_u, user)$. For the test set, we remove any form of user mentions to mimic an anonymous post.

Our experiments investigate a number of combinations of (training, test) data sources:

1. (Twitter, Twitter)
2. (Twitter + Foursquare + Instagram, Twitter)
3. (Twitter, Foursquare)
4. (Twitter + Instagram, Foursquare)
5. (Twitter, Instagram)
6. (Twitter + Foursquare, Instagram)

For the first two cases, we split the Twitter timeline into separate training and test portions. However, the more interesting conditions are 3-6, as the source of anonymous posts does not come from the data source used to generate the language model.

Additionally, we vary the amount of available profiling information by successively revealing larger parts of the training data. Furthermore, we also study the impact of changing the size of the set of anonymous posts q .

Multiple training sources are combined (Conditions 2, 4, and 6) in the following way: for a fixed amount of available




Figure 1: F-1 score curve of classifiers for conditions 1-2 varying sample size and training data size




Figure 2: F-1 score curve of classifiers for conditions 3-4 varying sample size and training data size

profiling information, 20% of it is made up from Instagram or Foursquare data (or 40% if these are put in together) and the rest from Twitter, due to the relative abundance of Twitter data.

The performance of our classifier is given by accuracy of predicting the correct user who generated the anonymous posts. The results are averaged across 10 randomization runs.

In Figures 1-3, the micro averaged F-1 score curves for the classifiers built under different conditions are shown.

In Table 1, the results for Conditions 1-2 can be found. We note that the usage of additional OSNs does not improve the de-anonymization performance. This is not too surprising as the source of the anonymous posts come from Twitter. The results for Conditions 3-4 can be found in Table 2. We remark that the classifier’s performance does not improve with the addition of Instagram data as users can cross-post check-ins on Twitter and users can check-in into a venue multiple times.

The more interesting results can be found in Table 3, which presents the results for Conditions 5-6. We note that in this case, as training and test data come from different sources, there is some improvement in the de-anonymization performance when we include Instagram data as an extra training source.

With respect to RQ 1, we note that it is possible to match




Figure 3: F-1 score curve of classifiers for conditions 5-6 varying sample size and training data size

Table 1: User de-anonymization for different sizes of training and test data, using Twitter as a source for anonymous posts

Train Source	Posts Seen	# anonymous posts			
		1	5	15	20
Twitter	50	18.21	47.80	69.90	74.22
	100	19.19	48.85	71.41	76.96
	200	18.45	48.68	71.57	77.03
	500	19.59	49.32	73.23	77.96
	1000	20.24	49.83	73.34	78.23
Twitter Instagram Foursquare	50	17.56	46.17	69.82	73.24
	100	18.38	47.88	71.24	75.82
	200	18.78	48.35	72.23	77.69
	500	18.16	50.84	72.97	77.87
	1000	19.90	49.39	72.81	76.66

Table 2: User de-anonymization for different sizes of training and test data, using Foursquare data as a source for anonymous posts

Train Source	Posts Seen	# anonymous posts			
		1	5	15	20
Twitter	50	32.43	54.97	61.33	62.34
	100	33.17	58.35	63.24	63.10
	200	34.51	60.85	67.40	67.89
	500	35.81	62.41	67.50	68.94
	1000	36.91	61.58	67.88	67.06
Twitter Instagram	50	32.68	55.72	61.87	62.80
	100	33.16	58.43	64.14	66.01
	200	36.56	60.34	66.84	69.02
	500	37.58	62.24	69.17	69.04
	1000	35.00	62.57	66.92	67.95

user profiles across OSNs based only on their textual data, achieving a maximum accuracy of **77.96%** when using only 500 posts (the equivalent of 20% of the average length of the Twitter timelines at our disposition). Furthermore, the usage of multiple OSNs as training data source seems to improve the classifiers’ performance when the source of anonymous posts and training data are distinct, suggesting there is a consistency in user language and vocabulary across the chosen OSNs. We also observe that, in general, the more anonymous posts are available, the better the performance of the designed classifier becomes.

Table 3: User de-anonymization for different sizes of training and test data, using Instagram data as a source for anonymous posts

Train Source	Posts Seen	# anonymous posts			
		1	5	15	20
Twitter	50	17.44	35.12	49.22	52.32
	100	19.07	40.17	55.51	59.82
	200	20.77	41.71	58.06	60.07
	500	20.75	41.27	58.34	61.74
	1000	21.00	40.15	55.62	60.34
Twitter Foursquare	50	18.05	35.91	51.20	53.17
	100	18.28	38.44	55.01	57.12
	200	21.70	44.81	60.90	64.50
	500	22.55	44.02	60.40	63.78
	1000	21.87	41.85	60.42	63.26

5. INFORMATION VALUATION

Let us again start from an OSN user base \mathcal{U} , in which each user u is defined by the set of his associated timelines $\{M_u^k\}_{k=1}^K$. Further, let S be an OSN such that we can define the set of all posts made by u in S as his timeline, M_u^S . We treat the timeline as a long consecutive piece of text in which each post constitutes a sentence. We use information from the timeline to estimate the probability of a user manifesting a certain property A . This probability is denoted by $p(A|M_u)$, where A denotes “ u shows Property A ” and M_u is the user’s timeline. Due to our definition of timelines, the same method can be used for full timelines or subsets of posts. Regardless of the chosen scope, we now project the timeline into an n -dimensional TF-IDF weighted vector space that allows us to train a classifier \mathcal{C}_A , estimating the final $p(A|M_u)$.

5.1 Measuring informativeness

Our objective is to find a function which quantifies the information carried in a post. On the one hand, we are interested in capturing the relevance of a post with respect to a certain inference task, on the other hand, in order to avoid redundancy or attributing a high score to already seen information, we are interested in capturing the novelty of some content with respect to what is already known. In a spirit similar to [5], we model the information content in two ways:

- Relevance ρ of the post with respect to an inference task or a set of tasks;
- Novelty ν of the post with respect to the user’s previously posted content.

We form our informativeness score as a convex combination between these two quantities, thus introducing a mixture parameter $\lambda \in [0, 1]$. Now, for each newly authored post m , we can define an informativeness function $\mathbb{I} : \mathbb{R}^n \times \mathbb{R}^n \times \mathcal{C} \rightarrow \mathbb{R}^+$ as follows:

$$\mathbb{I}(m, M_u, \mathcal{C}) = \lambda \nu(m, M_u) + (1 - \lambda) \rho(m, \mathcal{C}) \quad (4)$$

5.1.1 Relevance

Measuring the relevance of shared content can be intuitively thought of as determining which piece of shared content contains features that are important for the classifier’s

decision. A popular choice of such a function describing feature importance is the *Gini Importance* (Ig). For a feature θ , the Gini Importance for a classifier \mathcal{C} is defined as:

$$Ig_{\mathcal{C}}(\theta) = \sum_T \sum_{\tau} \Delta i_{\theta}(\tau, T). \quad (5)$$

Where τ is a node, T a decision tree and $\Delta_i(\tau)$ the decrease in Gini Impurity. The Gini Importance indicates how often a particular feature θ was selected for a split, and how large its overall discriminative value was for a particular classification problem. We estimate the overall relevance ρ of a post by summing up the importance scores across features contained in the post $\vec{m} \in \mathbb{R}^n$:

$$\rho(\vec{m}, \mathcal{C}) = \sum_{i=0}^n Ig_{\mathcal{C}}(m_i)$$

5.1.2 Novelty

For a fixed $u \in \mathcal{U}$ and OSN S , let $\vec{m}_1, \vec{m}_2 \in \mathbb{R}^n$ be the vector representation of shared contents $m_1, m_2 \in M_u^S$, the user’s timeline.

Informally, the function *novelty* : $\mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^+$ should have the following properties:

1. \vec{m}_2 should have low novelty if it is contained in \vec{m}_1 .
2. \vec{m}_2 should have low novelty if it is similar to \vec{m}_1 .
3. \vec{m}_2 should have high novelty if it is distinct from \vec{m}_1 .

Let \vec{m}_1 be denoted as $(m_{1,1}, \dots, m_{1,n})$. The proposed function to measure novelty is the following: Let $\nu: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^+$, be a non-symmetric function defined by:

$$\nu(\vec{m}_1, \vec{m}_2) = \frac{\sum_{i=1}^n \exp(-\alpha(|m_{i,1}| + |m_{2,i}| - 1))}{\sum_{i=1}^n 1[m_{2,i} \neq 0]} \quad (6)$$

For each word, the novelty function decays with the number of times that the word appears. By regulating α we can control how many times we have to observe a word to not consider it novel anymore.

5.2 Experimental Setup

As a concrete example of the general data-driven label prediction problem introduced previously, we turn towards the task of predicting whether a person will visit a particular type of location (*e.g.*, an Italian restaurant or a golf course) based on their social network timeline(s). These timelines are projected into a TF-IDF-weighted vector space. The vocabulary is curated by: removing all links and user mentions, stop words, words which occur less than 5 times and, when possible, word lemmatisation using WordNet. For the prediction task, we use the AdaBoost algorithm [3] with decision trees as weak learners as our classifier since they generally work well without refined parameter tuning. The classifier’s performance is evaluated under 10-fold cross-validation.

We begin by training one binary classifier per venue type that decides whether or not a user’s timeline suggests they are likely to visit that type of location. For every test user u , we initiate the procedure by randomly sampling a single post from their timeline S_k and create a truncated timeline. Then, at each iteration, we sample a constant number d of




Figure 4: Percentage of users visiting different venue types

posts from the timeline, add them to the truncated timeline and make a prediction using this iteratively updated input vector. The procedure is iterated until user u has no more posts left (or until the truncated timeline reaches a fixed amount of posts). We obtain an ordered sequence of predictions: $(y_0, y_1, \dots, y_{n_{end}})$, where n_{end} represents the number of iterations. When $d = 1$, *i.e.*, we add one post at a time, we simulate the situation in which an existing user profile is updated over time as new content is being shared.

We aggregate results taking into account the varying timeline length across users in the following way: the maximum timeline length $n_{l_{max}}$ is calculated. Then, for each user whose timeline is shorter than $n_{l_{max}}$, the last prediction $y_{n_{end}}$ is repeated to generate a sequence of predictions of length $n_{l_{max}}$. As a performance baseline, we randomly sample posts to be added. In the active method, instead, we select the posts which are most informative according to our metric presented in the previous section.

Foursquare offers a hierarchical taxonomy of places that display very different relative popularity. Figure 4 plots the percentage of our user population that has visited each of the more than 500 categories. We can note that the distribution is heavily skewed. While virtually all users have, at some point, visited places that fall into broad categories such as *[Arts & Entertainment]* or *[Food]*, others are so specific that they remain almost empty (*e.g.*, *[South Tyrolean Restaurants]* or *[Hunting Supplies]*). For the purpose of venue prediction, we are forced to make a subselection of categories that are neither so broad that the prediction task would become trivial nor so specific that the classifier would not find sufficiently many training examples. For this reason, we focus on those categories that were visited by 25% to 35% of the population, giving us a set of 37 venues (highlighted in yellow in the figure).

5.3 Results

5.3.1 Cross OSN

Each user’s timeline is initialized with a single post to which we iteratively add additional, randomly sampled, posts to form an updated user model. Figure 5, shows the classifier performance as a function of the number of posts available to the user model. From this overview, we note three recurring slope patterns. Some venue-specific classifiers quickly reach their optimal performance after as few as 750 posts have been observed (top figure), for others, significantly more iterations are required (center figure), and lastly, for some particular venues, the classification accuracy hardly benefits at all from using more posts (bottom figure). We refer to these three situations as quick-to-learn, slow-to-learn and hard-to-learn venues, respectively. Figure 6 gives a complete

overview of the relative frequency of mentions of the chosen venue categories and their affiliation to the three slope types. The general tendency seems to be that frequently mentioned venues tend to be quicker to learn than rare ones, while hard-to-learn venues appear to be randomly spread across the observation frequency range.

5.3.2 Active resource selection

After having confirmed the intuitive assumption that (within the limits of our three slope types) more data results in more accurate predictions, we now proceed to describing an active selection scenario in which we expand the user model by the most informative posts according to our metric rather than random ones. To this end, we fix the novelty parameter α at 0.5, meaning that after a word appears 5 times, its novelty becomes negligible. Table 4 highlights this method’s performance at different settings of λ . 50 posts are actively selected for this experiment and we note that our selection scheme biased towards informativeness delivers significantly better F_1 performance¹ than the random selection baseline at all parameter settings. The overall best performance was obtained at a setting of $\lambda = 0.1$, where the score mixture is dominated by feature importance while still taking into account novelty.

Table 4: Average classifier performance across all 37 venue types, for different λ .

λ	F_1 -score	Precision	Recall
baseline	15.36	52.78	8.99
0.0	42.10	44.21	40.18
0.1	44.19	46.64	41.97
0.2	44.16	46.57	41.98
0.3	43.63	46.00	41.48
0.4	43.66	46.09	41.47
0.5	43.23	45.50	41.18
0.6	43.80	46.31	41.56
0.7	43.55	46.85	40.69
0.8	42.78	45.72	40.19
0.9	42.21	45.99	39.00
1.0	18.95	43.38	12.12

Let us return to our previously introduced categorization of learning curve slope types. Table 5 shows the influence of λ on the performance of the three slope categories. We observe clearly diverging tendencies between quick and slow-to-learn venues. While quick-to-learn venues benefit from low novelty contributions, their more slowly evolving counterparts benefit from novelty-biased informativeness scores. Examples can be seen in Figure 7. Again, hard-to-learn venue types do not show any noticeable response to different choices of λ , as long as the relevance component is not fully turned off.

Furthermore, for some particular venues, the classifier attained a better performance when using only 50 actively selected posts than when using the full timeline of the users. Some examples can be found in Table 6.

Regarding RQ 2, we show again that there is consistency in terms of content shared across OSNs, in particular, we show that it is possible to predict venue type visits based on what is shared on Twitter and Instagram. Furthermore,

¹Average F_1 score for all classifiers is computed using the average of precision and recall across all classifiers.




Figure 5: Venue-dependent classifier learning rates across training iterations.




Figure 6: Average percentage of venue type mentions per twitter timeline, when user has visited the venue type with category assignment of quick, slow or hard to learn for venue types.

Table 5: Average F_1 -scores across slope types.

λ	Quick	Slow	Hard
0.0	47.81	38.51	36.70
0.2	50.16	40.40	38.43
0.4	49.83	40.10	37.23
0.6	49.59	40.18	38.48
0.8	46.97	41.19	36.97
1.0	17.87	9.70	30.48

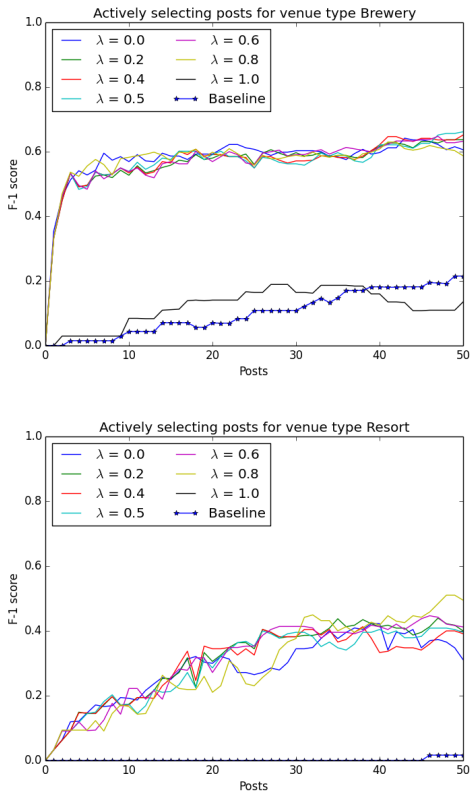


Figure 7: Example of venue types which benefit from novelty component. Top figure: venue type *Brewery*. Bottom figure: *Resort*

with respect to RQ 3, we show that using our designed metric, we can find the posts which are most relevant to predict venue type visits. In particular, using the active learning framework with our information measure of content in posts as selection criteria, we observe overall quicker learning rates and in some cases, we can use a significantly reduced the number of posts to attain a classifier performance which is comparable to using the full timeline of the user.

6. CONCLUSION

In this paper, we studied privacy hazards pertaining to cross-platform social network usage. Individually innocuous posts can lead to leakage of critical information when aggregated along or across a user’s OSN profiles. We quantify this effect in two experiments: (1) uniquely identifying users in an anonymous pool and (2) predicting user properties manifested on one OSN platform based on content from other parallel profiles.

In the user de-anonymization task, we note that it is possible to match user profiles across OSNs based only on their textual data, with as little as 10% to 20% of the user’s full timeline. Furthermore, the inclusion of multiple OSNs as training data sources has been shown to improve the classifiers’ performance when the source of anonymous posts and training data are distinct. This suggests that there is a consistency in user language and vocabulary across OSNs.

In the information valuation task, we propose a general-purpose metric of textual informativeness in order to model the value of shared information items both for service providers (predictive power) as well as the user (potential privacy haz-

Table 6: Classifier performance using a truncated timeline versus a full timeline.

	Sushi Rest.	Cocktail Bar	Gastropub	Brewery	Nightclub
Full timeline	52.86	31.48	38.53	50.88	43.40
Truncated + Random Selection	11.49	4.97	5.36	21.51	20.11
Truncated + Active Selection	61.72	55.84	54.79	66.15	54.23

ards). We show experimentally that the metric reflects the relative importance of posts with respect to the inference task being performed. When actively selecting a subset of posts per user, this method was always able to beat a random selection baseline. While choosing posts according to their relevance seems to lead to better performance in general, we noted that only for some venues there was a noticeable benefit in including a strong novelty component in the information scoring function.

This work focused on showing the privacy hazards that arise from sharing content which seems uninformative or harmless. In the future, we are excited to extend this line of work by a dedicated investigation of information valuation scores on the user side (*e.g.*, of an OSN) as it would greatly help people understand their own digital footprint and enable them to recognize moments of critical information disclosure. Furthermore, part of this work focused on proposing a metric for information valuation with respect to an inference task. We are interested to extend this line of work by an *in-vivo* study of monetary efficiency of advertisers as a consequence of introducing an informativeness-aware resource selection scheme in their real-time bidding (RTB) pipelines.

7. REFERENCES

- [1] Ashish Agarwal, Kartik Hosanagar, and Michael D Smith. Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of marketing research*, 48(6):1057–1073, 2011.
- [2] Adam Berger. *Statistical Machine Learning for Information Retrieval*. PhD thesis, Pittsburgh, PA, USA, 2001. AAI3168516.
- [3] Leo Breiman. Arcing classifier (with discussion and a rejoinder by the author). *Ann. Statist.*, 26:801–849, 1998.
- [4] Claudio Carpineto and Giovanni Romano. Semantic search log k-anonymization with generalized k-cores of query concept graph. In *Advances in Information Retrieval*, pages 110–121. Springer, 2013.
- [5] Charles L.A. Clarke, Maheedhar Kolla, Gordon V. Cormack, Olga Vechtomova, Azin Ashkan, Stefan Büttcher, and Ian MacKinnon. Novelty and diversity in information retrieval evaluation. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’08, pages 659–666, New York, NY, USA, 2008. ACM.
- [6] Yves-Alexandre de Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports*, 3, 2013.
- [7] Cynthia Dwork. Differential privacy. In *Encyclopedia of Cryptography and Security*, pages 338–340. Springer, 2011.
- [8] Maria Han Veiga and Carsten Eickhoff. A cross-platform collection of social network profiles. In *Proceedings of the 39th SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2016.
- [9] Weiyin Hong and James YL Thong. Internet privacy concerns: An integrated conceptualization and four empirical studies. *MIS Quarterly*, 37(1):275–298, 2013.
- [10] Ronald Howard et al. Information value theory. *Systems Science and Cybernetics, IEEE Transactions on*, 2(1):22–26, 1966.
- [11] Mark J Keith, Samuel C Thompson, Joanne Hale, Paul Benjamin Lowry, and Chapman Greer. Information disclosure on mobile devices: Re-examining privacy calculus with actual user behavior. *International journal of human-computer studies*, 71(12):1163–1173, 2013.
- [12] Michal Kosinski, David Stillwell, and Thore Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 110(15):5802–5805, 2013.
- [13] Ravi Kumar and Andrew Tomkins. A characterization of online browsing behavior. In *Proceedings of the 19th International Conference on World Wide Web, WWW ’10*, pages 561–570. ACM, 2010.
- [14] David Lazer, Alex Pentland, and et al. Computational social science. *Science*, 323(5915):721–723, 2009.
- [15] Wen Li, Carsten Eickhoff, and Arjen P. de Vries. Want a coffee?: predicting users’ trails. In William R. Hersch, Jamie Callan, Yoelle Maarek, and Mark Sanderson, editors, *SIGIR*, pages 1171–1172. ACM, 2012.
- [16] Naresh K Malhotra, Sung S Kim, and James Agarwal. Internet users’ information privacy concerns (iupc): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, 2004.
- [17] H Jeff Smith, Sandra J Milberg, and Sandra J Burke. Information privacy: measuring individuals’ concerns about organizational practices. *MIS quarterly*, pages 167–196, 1996.
- [18] Kathy A Stewart and Albert H Segars. An empirical examination of the concern for information privacy instrument. *Information Systems Research*, 13(1):36–49, 2002.
- [19] Robert West, Ryen W White, and Eric Horvitz. Here and there: Goals, activities, and predictions about location from geotagged queries. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*, pages 817–820. ACM, 2013.