# Concealing Interests of Passive Users in Social Media

Yaroslav Nechaev[1,2], Francesco Corcoglioniti[1], and Claudio Giuliano[1]

[1] Fondazione Bruno Kessler, Via Sommarive 18, 38123 Trento, Italy
[2] University of Trento, Via Sommarive 14, 38123 Trento, Italy
{nechaev,corcoglio,giuliano}@fbk.eu

**Abstract.** User profiling has existed in the social media since their inception and has supported most of their business model. Even if users do not actively share the information about themselves on the social media (so-called passive users), they can still be profiled based on their location and who they follow. In this paper, we present a system that leverages the linking of followed (popular) Twitter users to DBpedia, and the information therein contained, to help users concealing their digital footprint. Specifically, our approach helps a passive Twitter user to stay private by proposing a list of additional profiles to follow that would confuse the social media's inference pipeline and prevent it from inferring useful information about that passive user and his interests.

**Keywords:** Social Media, User Profiling, Linked Open Data, Machine Learning

## 1 Introduction

Currently, an enormous amount of people use social media every day: just recently, in July 2017, Facebook has hit two billion monthly users. Every action of those people is being recorded, analyzed and possibly sold to third parties in one form or another. Additionally, this data is used to acquire a digital footprint of users: what they like or do not like, their level of education, gender, race and much more.

Knowing that, people have learned to be careful about what they post, like or share on social media. Some go even further — they just follow a number of profiles they like and never actually generate any content that could be gathered or analyzed. In the literature, such users are called "passive users". A number of recent studies [3,7,20,26] have demonstrated that, despite their best effort, passive users can still be profiled based on location and the profiles they read. By adopting the paradigm "we are what we read", social media can infer digital footprint of passive users almost as good as for active ones, with the result that protecting the privacy of passive users remains an open issue.

Users can, of course, choose to stop using the social media altogether in an attempt to preserve their privacy. However, in the modern world, it is becoming increasingly hard for an individual to abandon the benefits such services provide just for the sake of privacy. This situation is not unique to the social media: the same happens with recent machine learning-based consumer products such as voice assistants, translation services and even self-driving cars. By using those services, people agree to provide data that is used to train the system or perform inference of user profile attributes.

On the other hand, the privacy of the user is not something that has to be given up in favor of new exciting technologies and services. Companies can choose to protect user's
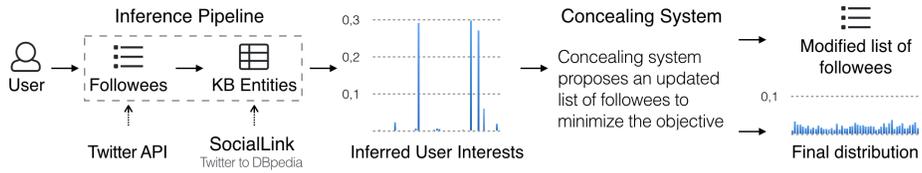
**Fig. 1.** The proposed concealing approach

privacy without degrading the user experience by implementing various techniques, such as differential privacy [6] or by shifting the computation on the user-controlled device [14]. Despite the availability of such techniques, modern social media are reluctant to implement them: their business depends on their ability to learn as much as possible about the target user and exploit this information the best they can to show advertisement and sell auxiliary services.

In this paper, we present a system that helps Twitter passive users to conceal their digital footprint by leveraging the alignments from Twitter profiles to DBpedia entities provided by SocialLink [17] and the categorical information available about those entities. Specifically, our main contribution are two concealing approaches (Greedy and Joint) that help passive users to stay private by proposing additional Twitter profiles to follow (followees) that would turn the user interests distribution inferred from followees close to the uniform one. This has the effect of confounding the social media's inference pipeline, preventing it from inferring useful information about the real interests of the target user. The original list of followees could then be stored on a user-controlled device (using a custom application or a browser plugin) allowing to recreate the original timeline. Our system proposes as few followees as possible to circumvent possible social media-induced limitations, reduce network load and cluttering of the user timeline. Our proposal is highlighted in Figure 1. This task of concealing user interests is inspired by the obfuscation examples provided in Brunton and Nissenbaum's book [5], specifically, the "Bayesian Flooding" idea by Kevin Ludlow [13].

To evaluate our approaches we build a user's interests inference pipeline mirroring state-of-the-art approaches [3,20] that use the list of followees to infer passive users' interests. This pipeline is based on the idea, initially introduced by Piao and Breslin [20], that the followee list could be mapped to the distribution of interest categories for the user by linking followees to DBpedia/Wikipedia and then exploit the categorical information therein contained, averaging it along followees to derive an interests distribution for the target passive user. Our pipeline reiterates on this idea but uses the high-quality alignments of SocialLink [17] instead of the simple heuristics used in state-of-the-art approaches to map Twitter profiles to DBpedia/Wikipedia resources.

We evaluate our Greedy and Joint approaches and a Random baseline against this inference pipeline. We showcase a number of results, demonstrating that the Joint approach is able to achieve almost a perfect uniform distribution, decreasing the average KL-divergence by 94% compared to the Random baseline. The Joint approach solves a joint optimization problem learning the most efficient followee configuration. Additionally, we show the impact of our approaches on the performances of the inference

pipeline using the precision at rank N (P@N) metric and, since we aim at suggesting as few new followees as possible, the average amount of suggested profiles.

Finally, we test our concealing approaches in a more real world setting by evaluating them against the Twitter's *Who To Follow* [9] recommendation system. This system recommends a target user other Twitter profiles to follow based (also) on his/her inferred interests, which can be deduced from produced recommendations leveraging the same DBpedia/Wikipedia alignments and categorical information used in the aforementioned inference pipeline. We show how our approaches can partially equalize those inferred interests, proving that our techniques have general applicability and can be used with little or no knowledge about the target inference pipeline algorithm.

The rest of the paper is structured as follows. In Section 2, we briefly present the related work. Then, we formally define our problem in Section 3, followed by the description of the user's interest inference pipeline used as reference in Section 4. Our concealing approaches are described in Section 5. Finally, we report the evaluation results in Section 6, closing with the conclusion in Section 7.

## 2 Related Work

**User profiling** Profiling of users in social media has been performed since their inception both by the social media themselves and researchers. The inference of many user profile attributes, such as gender [26], age [22], location [23], political affiliation [26], level of education [11], and occupational class [21], has been studied both for active [1,15,16,19,24,25] and passive users [3,7,20,26].

Followee information (i.e., the social graph) of a user plays a key role in user profiling. One of the early studies successfully utilizing followee information to infer a user profile attribute was Sadilek et al. [23]. The usage of GPS data of followees allowed them to pinpoint the exact location (down to coordinates) of a target user with 80% accuracy. No additional information from the user's profile was used. However, their approach required a significant amount of high-quality location data to perform inference, which is in many cases hard to acquire.

Zheleva and Getoor [26] exploited the social graph to infer gender and political affiliation of users. They were able to profile both active and passive users. Additionally, the more profiles of friends were public, the better the accuracy of such inference. Their study is one of the first that raised important questions about the user privacy in the social media. They concluded that the measures typically employed by social media to protect personal data are not sufficient.

More recently, there has been an increasing interest in profiling interests of passive users. Besel et al. [3] proposed an inference pipeline that utilize followee information to infer interests. They link a target user's followees to the corresponding Wikipedia pages using their names and then category information is used to determine interests. Our inference pipeline is constructed following their approach. Piao and Breslin [20] iterated on this idea and produced a better performing system by improving the entity linking step and using a different interest propagation approach.

In summary, social graph-based features have proven to be a useful feature in all cases, confirming the idea that "you are what you follow".

In most approaches for user profiling, a classifier is built for the inference of each individual attribute. However, more holistic approaches were studied as well. For example, Li et al. [12] proposed a two-layered structure. On the first layer they reimplemented typical classifiers for attributes like location or education from previous works. Then they built a Probabilistic Logical Reasoning framework and used the results from the first layer as evidence. Since the accuracy of the first layer is not 100%, the second layer should be able to account for possible errors and handle contradictory knowledge, effectively preventing error propagation. Two reasoning approaches were explored: Markov Logic Networks and Probabilistic Soft Logic. As a result they were able to work with a wide variety of attributes: from gender to general like relationships towards entities.

**Privacy in the social media**  Many user profiling studies cover the topic of users' privacy in social media by warning of the potential risks of sharing private information in user profiles. Privacy problems in social media, however, go beyond user profiling. Bettini and Riboni [4], for instance, produced a comprehensive study on privacy preservation and the technological challenges arising in pervasive systems such as social media.

Felt and Evans [8] studied how social media themselves can protect users by redesigning their APIs. They devised a privacy-by-proxy technique, where data is revealed to applications only when needed, limiting the scope to prevent data harvesting.

Luo et al. [14] proposed to protect social media users by encrypting the user-related information before it reaches the social media. Their approach aims at achieving a goal similar to ours: not only to prevent the third parties from accessing the sensitive data of the target user but also the social media themselves.

Kevin Ludlow [13] introduced the concept of Bayesian Flooding demonstrating that the social media's advertisement and recommendation systems can be confused by flooding the user's timeline with artificial actions.

**Binarized Neural Networks**  Even though we did not employ neural networks in this paper, our Joint approach was influenced by recent studies on Binarized Neural Networks (BNN) [2,10]. We used the binarization and back-propagation procedures from Hubara et al. [10] to find an optimal solution to our optimization problem.


## 3  Problem Definition

The goal of this paper is to protect the privacy of passive users by modifying their lists of followees in such a way that makes it much harder for the target inference pipeline to profile their interests. Followees are now being universally used by social media as part of the digital footprint of a person and play a crucial role in inferring user profile information such as interests. Even if the user does not post or share content on the social media (passive user), followee data is still available to third parties and the social network itself. The idea is to conceal this information without degrading a user experience, which, in case of modifying his/her followee list, can be achieved by storing the original unmodified list on a user-controlled device and use it to filter the timeline.

In this work we focus on concealing approaches tackling the inference of passive users' interests. Given a user $u$ and his/her list of followees $l_u$, a user's interests inference pipeline $g(l_u)$ is designed to infer this user's interests $\mathbf{c}^u = g(l_u), \mathbf{c}^u \in \mathbb{R}^n, c_i^u \geq 0, \sum_i c_i^u =$

1, where $n$ is the amount of interest categories and $c_i^u$ is the score of interest category $i$ for user $u$. The categories are then ranked based on their score $c_i^u$ and the final list of top $k$ categories is produced to represent the user's interests. Real-world implementations of such inference pipeline will have a set of thresholds to abstain from classifying a user's interests when their ranking is too ambiguous, i.e., $c_i^u$ scores of top categories are very close to each other, making it impossible to reliably determine a user's interests.

An ideal approach for concealing user interests would try to make all the interest categories indistinguishable from each other. Therefore, the goal of our system is to modify the information about the user, i.e., transform user profile $u$ to user profile $u'$ (in our case produce a modified list of followees $l_{u'}$), so that the target inference system produces ambiguous results. Formally, the objective is to minimize the *Kullback-Leibler divergence*[3] between the $\mathbf{c}^{u'} = g(l_{u'})$ and the uniform distribution over possible interest categories:

$$D(u') = D_{KL}(\mathcal{U}\{1,n\}||\mathbf{c}^{u'}) = -\sum_i \frac{1}{n} \log c_i^{u'} + \sum_i \frac{1}{n} \log \frac{1}{n} = -\sum_i \frac{1}{n} \log c_i^{u'} - \log n \quad (1)$$

A possible limitation of such problem formulation is that the social network can impose limitations on the amount of follow requests from the user and a large list of followees can significantly increase the amount of API requests needed to acquire the timeline thus creating a worse user experience. In this case we may require our system to propose as few modifications to the initial followees list as possible and the final objective will be as follows:

$$J(u',u) = (1-\alpha)D(u') + \alpha(|l_{u'}| - |l_u|), \quad \alpha \in [0,1) \quad (2)$$

where $\alpha$ is a parameter that balances the tradeoff between minimizing the KL-divergence (which requires adding followees to equalize inferred interest categories) and minimizing the amount of proposed followees.

## 4 Interests Inference Pipeline

To evaluate our concealing approaches (presented in Section 5) we have implemented a user's interests inference pipeline ($g(l_u)$ in Section 3) that infers user's interests based on the list of followees. We follow Besel et al. [3] state-of-the-art approach, improving it by removing dependencies from the Wikipedia API and the WiBi Taxonomy. We do that by replacing the Entity Linking heuristics used there with our state-of-the-art resource called SocialLink [17,18] and pre-computing the category distributions over a taxonomy of 49 top categories for all possible entities in English DBpedia/Wikipedia. This enabled us to acquire a simple and robust system that allows testing different approaches for concealing a user's digital footprint. More in detail, the pipeline employs the following three-step procedure:

**Fetch followees** Followees $l_u$ of the target user $u$ are fetched using the Twitter API.

---

[3] http://en.wikipedia.org/wiki/Kullback-Leibler_divergence

**Link followees to DBpedia/Wikipedia**  Each followee profile $f \in l_u$ is linked using our state-of-the-art Linked Open Data resource called SocialLink. SocialLink contains direct alignments between DBpedia entities for persons and organisations and corresponding Twitter profiles. Each alignment in this resource is associated with a confidence score $s_f$ that we use here to appropriately weight the contribution of each followee $f$ to the final user's interest distribution. We want to make sure that our linking procedure is robust since an error at this step will propagate along the pipeline. Therefore, we selected a subset of $m = 101,769$ high-quality alignments by setting custom conservative thresholds on confidence scores (with respect to the default ones used in [17]) that provide 91% precision and 31% recall performance against SocialLink gold standard.

**Produce interest scores**  At this step each aligned followee has to be mapped to a category distribution $\mathbf{f}$, whose elements $f_i$ are the scores for interest categories $i$. Similarly to [3], we use the DBpedia/Wikipedia category graph to propagate the specific categories associated to the followee entity up to the 49 top-level categories here considered, for each of which a relevance score is computed. This process resulted in a list of 3,507,016 scored DBpedia/Wikipedia entries. The interest scores in the category distribution $\mathbf{f}$ are then normalized using a modified softmax function $\sigma(\mathbf{f})$ to produce a valid probability distribution across possible interests, where the normalized score $\sigma(\mathbf{f})_i$ for interest category $i$ is computed as follows:

$$\sigma(\mathbf{f})_i = \frac{z(f_i)}{\sum_{k=1}^{n} z(f_k)}, \quad z(x) = \begin{cases} e^x & \text{if } x \neq 0\,, \\ 0 & \text{if } x = 0 \end{cases} \tag{3}$$

This normalization procedure preserves zero scores for categories that were not observed for the given followee $f$, thus reducing the noise across final user's interests. The interests distribution $\mathbf{c}^u$ for the user $u$ is finally computed as a weighted average of normalized scores for user's followees $f \in l_u$:

$$\mathbf{c}^u = \frac{\sum_{f \in l_u} s_f \sigma(\mathbf{f})}{\sum_{f \in l_u} s_f} \tag{4}$$

The code of our interests inference pipeline along with the precomputed list of scores for each aligned followee can be found in our GitHub repository.[4]

## 5  Concealing Approaches

In order to conceal a user's interests we propose three approaches that calculate an updated list of followees to minimize the objective defined by (2): Greedy approach, Joint approach, and the baseline Random approach. In all cases, the system expects the initial list of followees in input and chooses new followees from the same list of pre-aligned profiles (from SocialLink) that our interests inference pipeline uses.
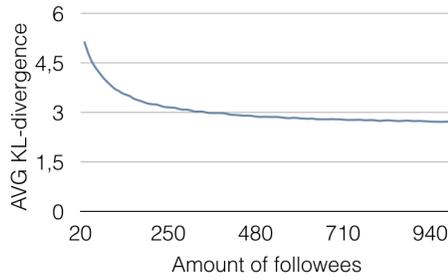
---

[4] http://github.com/Remper/re-coding-ws

**Fig. 2.** Average KL-divergence for different amounts of followees using Random approach
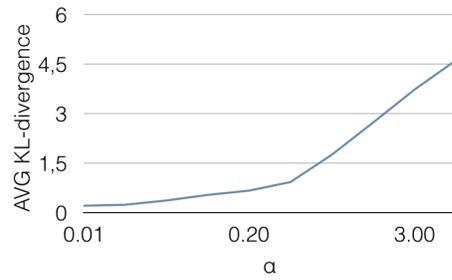
**Fig. 3.** Average KL-divergence for different $\alpha$ values using Greedy approach

**Random approach**    The most trivial direction we can take is to randomly follow new people in hope that the category distribution will become closer to uniform. In this approach, the profiles to follow are randomly drawn from the above-mentioned list and the system will stop when 250 new unique followees are selected. Since the list has its own bias towards certain categories, it can be expected that the more followees we add this way the closer we get to the list's own category distribution. This is why the threshold on the amount of new followees has to be selected carefully in order to provide a positive improvement in terms of our objective. Figure 2 shows how the average KL-divergence changes based on the amount of followees proposed.

**Greedy approach**    The greedy approach iteratively selects a new followee from the pre-aligned list by picking the one that will decrease the KL-divergence between the resulting category distribution and the uniform distribution the most. Therefore, it can be seen as a breadth-first search over the space of possible configurations. The algorithm stops when it is no longer possible to select a new profile to follow that will improve the objective score (2). In our experiments the $\alpha$ parameter is set to 0.01. Figure 3 shows how the average KL-divergence changes based on the choice of $\alpha$.

**Joint approach**    Finally, we have devised an approach that directly solves the formulated optimization problem by jointly finding an optimal followee configuration. This approach is inspired by recent studies about Binarized Neural Networks [10] and it effectively "learns" the binary mask $\mathbf{w}^b$ where each element $w_j^b$ corresponds to a possible followee $j$ to add, being 1 if that followee $j$ should be followed and 0 otherwise.

Given the matrix of followees $F \in \mathbb{R}^{m \times n}$ that is obtained by simply stacking row-wise the category distributions $\sigma(\mathbf{f})$ of the pre-aligned followee list (i.e., $m = 101,769$) following the normalization procedure from (3), the $\mathbf{c}^u$ can be rewritten in terms of our binary mask $\mathbf{w}^b$:

$$\mathbf{c}^u = \frac{1}{\sum_i \mathbf{w}_i^b} \mathbf{w}^b F \tag{5}$$

The objective score can be computed as in (2). In order to solve this optimization problem, we define an additional weight vector $\mathbf{w} \in \mathbb{R}^m$, and $\mathbf{w}^b$ will now be computed
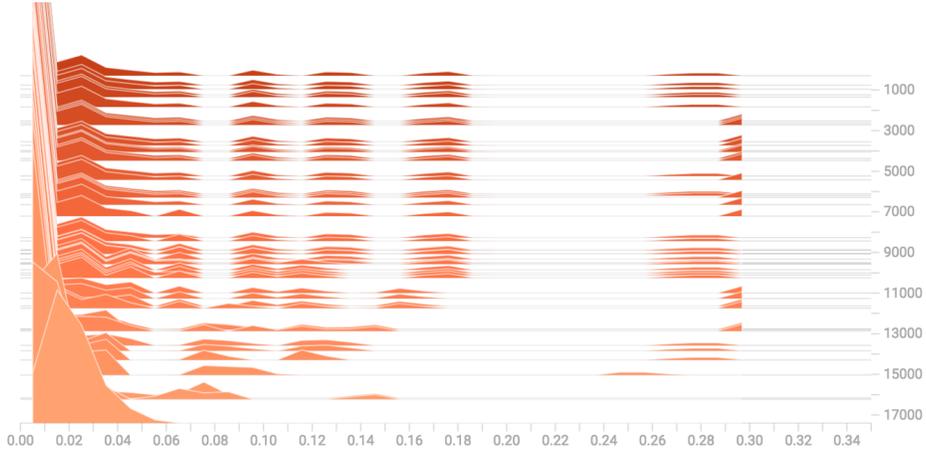
**Fig. 4.** An example histogram of user's categories converging towards uniform distribution over 17k iterations of Joint approach. Each slice represent score distribution among categories for the corresponding iteration. In a perfect scenario all scores should be equal to $1/n = 0.02$.

using a simple deterministic binarization:

$$w_j^b = Bin(w_j) = \begin{cases} 1 & \text{if } w_j \geq 0 \,, \\ 0 & \text{if } w_j < 0 \end{cases} \tag{6}$$

Then $\mathbf{w}$ is learned by gradient descent towards the objective. Note that since the derivative of the *Bin* function is zero almost everywhere, the gradient have to be back-propagated using the *straight-through estimator technique* suggested by Bengio et al. [2]. The $\mathbf{w}$ parameters are initialized by drawing from $\mathcal{N}(2l, l)$, where $l$ is the learning rate, ensuring that the majority of possible followees are followed at the first iteration. Figure 4 shows how a user's category distribution changes during the learning process.

## 6 Evaluation

We evaluate our concealing approaches Random, Greedy, and Joint against the user's interests inference pipeline of Section 4 and the Twitter's Who To Follow recommendation system. We measure four main performance metrics: (i) the average amount of followee suggestions; (ii) the ability to conceal a user's top-N interests categories, measured taking the categories produced without concealing as gold standard and measuring the precision P@N of the inference pipeline in reproducing those categories after concealing is applied; (iii) the average delta between the first and the second category probabilities; and (iv) the average KL-divergence between the category distribution produced by an approach and the uniform distribution across interest categories.

### 6.1 Evaluation Against Interest Inference Pipeline

In order to produce the dataset required to evaluate our concealing approaches, we gathered a list of all the authors from ISWC proceedings for the last three years, extracting

**Table 1.** System evaluation against our interests inference pipeline

| Approach | Suggestions | System performance (P@N) | | | Score diff to 2nd best | Diff from uniform (KL-div) |
|---|---|---|---|---|---|---|
| | | Top 5 | Top 10 | Top 15 | | |
| No mod | 0 | 1.0 | 1.0 | 1.0 | 0.25 | 7.52 |
| Random | 250 | 0.39 | 0.58 | 0.65 | 0.12 | 3.15 |
| Greedy | **27.49** | 0.52 | 0.55 | 0.56 | 0.03 | 0.20 |
| Joint | 76.61 | **0.37** | **0.45** | **0.49** | **0.02** | **0.16** |

the names, titles and abstracts of all the papers associated with each author. After gathering the list, the Social Media Toolkit [18] was used to find the corresponding Twitter profile for each author, leveraging the collected textual context. This resulted in a dataset of 491 Twitter accounts that were used to evaluate our approaches against the user's interests inference pipeline. The choice of the ISWC target audience was made to showcase how easy it is to profile people in a particular community just by having a set of names, keywords, or other related textual context.

The followees list for each author was gathered via the Twitter REST API and provided to our interest inference pipeline to produce an initial interest category distribution. Then, each approach was able to propose a modified list of followees and the category distribution was recalculated. Table 1 shows the resulting performance for each approach and the baseline numbers without applying any concealing technique. The baseline is assumed to be perfect at predicting user interests (have 100% precision at 5, 10 and 15 top categories) since our goal is to hide true user interests from the inference pipeline.

It can clearly be seen that, even though the Random approach already significantly reduces the KL-divergence, it takes a significant amount of suggestions to achieve this result. Random hides the initial top K categories of the user well, but produces a new top category (usually, Sport) that stands out and makes the target inference pipeline consistently producing false positives towards this category, which was not our goal.

The Greedy approach, on the other hand, produces an almost uniform distribution, while providing a relatively small amount of suggestions. It does not hide the top K categories as well as the Random approach, but the target system would have mostly likely abstained to infer user's interests given such category distribution. The results of this approach clearly show that if the concealing system is able to predict the expected score with high accuracy for an arbitrary followees list (in our case the approach had perfect information), it is possible to confuse the target inference pipeline.

Finally, the Joint approach is able to find a more efficient followee configuration than Greedy, producing the best results in almost all metrics at the cost of an increased suggestions count. The amount of suggestions can further be tuned by setting different values for the $\alpha$ parameter. However, we consider the current configuration to be a reasonable tradeoff.

**Table 2.** System evaluation against Twitter's Who To Follow

| Approach | System performance (P@N) | | | Score diff to 2nd best | Diff from uniform (KL-div) |
|---|---|---|---|---|---|
| | Top 5 | Top 10 | Top 15 | | |
| No mod | 1.0 | 1.0 | 1.0 | 0.84 | 12.23 |
| Joint | 0.20 | 0.37 | 0.47 | 0.44 | 8.56 |

### 6.2 Evaluation Against Twitter's Who To Follow

In this scenario we evaluate the Joint approach, resulting the best in the previous evaluation, against the Who To Follow system used in production by Twitter to suggest users to follow based on a target user's profile [9].[5] This system was chosen because it provides a simple way to collect user's interests as measured by Twitter. In order to evaluate against this system, we have created a number of new Twitter users providing as little initial information about the user as possible to the social network,[6] apart the lists of followees, so to simulate the behaviour of passive users.

After the creation of each account, the initial Twitter recommendations of users to follow were gathered. These users represent popular Twitter accounts that are always recommended to new Twitter users in a certain location. Then, a number of users were followed from the pre-aligned list with the intention to give a clear bias towards some interest category. After that, we gathered again the users to follow recommended by the network. Finally, the Joint approach was used to propose the modified list and the network's suggestions were gathered one more time. Overall, three lists of Twitter user recommendations were gathered, with the initial list acting as a filter to clean the other two lists from the location-based and general popularity-based suggestions.

The remaining two lists were mapped to distributions of interest categories using the formula presented in (4). Then the same metrics used in Section 6.1 were computed to evaluate results. Unfortunately, since the Who To Follow box had to be gathered manually and fresh Twitter accounts had to be created every time, the evaluation was significantly downscaled compared to what we initially have hoped to measure. However, it can clearly be seen in Table 2 that, even though the concealing approach does not have any information on how the target user's interests inference system works, it is often able to conceal the user's true category distribution.

In order to achieve better results, a training set can be gathered to modify the followee matrix $F$ using the same manual gathering approach employed during the evaluation. However, this is currently beyond the scope of this paper.

## 7 Conclusions and Future Work

In this paper we have presented approaches for concealing a social media user's digital footprint by proposing further people to follow, leveraging the alignment from Twitter

---

[5] http://support.twitter.com/articles/227220

[6] All user accounts were created using fresh email accounts using an IP address that can be tracked down to Microsoft Azure cloud datacenter in Cheyenne, USA.

users and DBpedia/Wikipedia entries provided by SocialLink [17]. We showed that by using simple techniques and without degrading user experience, passive Twitter users can prevent the network or a third party system to infer their interests based on the knowledge of who they follow. As our approach relies only on social graph information, which is present in any social media, we believe it can be generalized and ported to other platforms like Facebook and Instagram.

Even though the discussion about the privacy of users online has been a hot topic lately, social media are reluctant to implement industry standard techniques such as differential privacy and on-device computation, wanting instead to preserve their ability to sell ads and promote their services. In this situation, we believe it is increasingly important to explore various ways users can protect their digital identity.

In future work, further user profiling scenarios can be explored targeting different attributes (e.g. location) and profiling use cases. We also believe that much can be done to explore the privacy of active users, for example, making sure that the metadata shared with the social media is limited. Regarding our use case, more sophisticated approaches can be explored that are able to learn from the output of the interests inference pipeline to produce more accurate results.

## References

1. Abel, F., Gao, Q., Houben, G.J., Tao, K.: Analyzing user modeling on Twitter for personalized news recommendations. In: Proc. of 19th Int. Conf. on User Modeling, Adaption, and Personalization (UMAP). pp. 1–12. Springer-Verlag, Berlin, Heidelberg (2011)
2. Bengio, Y.: Estimating or propagating gradients through stochastic neurons. CoRR abs/1305.2982 (2013)
3. Besel, C., Schlötterer, J., Granitzer, M.: Inferring semantic interest profiles from Twitter followees: Does Twitter know better than your friends? In: Proc. of 31st Annual ACM Symposium on Applied Computing (SAC). pp. 1152–1157. ACM (2016)
4. Bettini, C., Riboni, D.: Privacy protection in pervasive systems: State of the art and technical challenges. Pervasive and Mobile Computing 17(Part B), 159 – 174 (2015), 10 years of Pervasive Computing' In Honor of Chatschik Bisdikian
5. Brunton, F., Nissenbaum, H.: Obfuscation: A user's guide for privacy and protest. The MIT Press (2015)
6. Dwork, C.: Differential privacy: A survey of results. In: Proc. of 5th Int. Conf. on Theory and Applications of Models of Computation (TAMC). pp. 1–19. Springer-Verlag, Berlin, Heidelberg (2008)
7. Faralli, S., Stilo, G., Velardi, P.: Recommendation of microblog users based on hierarchical interest profiles. Social Network Analysis and Mining 5(1), 25 (2015)
8. Felt, A., Evans, D.: Privacy protection for social networking APIs (2008), http://www.cs.virginia.edu/felt/privacybyproxy.pdf
9. Gupta, P., Goel, A., Lin, J., Sharma, A., Wang, D., Zadeh, R.: WTF: The Who to Follow Service at Twitter. In: Proc. of 22nd Int. Conf. on World Wide Web (WWW). pp. 505–514. ACM, New York, NY, USA (2013)
10. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized Neural Networks. In: Proc. of Advances in Neural Information Processing Systems (NIPS), pp. 4107–4115. Curran Associates, Inc. (2016)

11. Li, J., Ritter, A., Hovy, E.H.: Weakly supervised user profile extraction from Twitter. In: Proc. of 52nd Annual Meeting of the Association for Computational Linguistics (ACL). pp. 165–174. The Association for Computer Linguistics (2014)
12. Li, J., Ritter, A., Jurafsky, D.: Inferring user preferences by probabilistic logical reasoning over social networks. CoRR abs/1411.2679 (2014)
13. Ludlow, K.: Bayesian flooding and Facebook manipulation (2012), http://www.kevinludlow.com/blog/1610/Bayesian_Flooding_and_Facebook_Manipulation_RD/
14. Luo, W., Xie, Q., Hengartner, U.: FaceCloak: An architecture for user privacy on social networking sites. In: Proc. of Int. Conf. on Computational Science and Engineering. vol. 3, pp. 26–33 (Aug 2009)
15. Michelson, M., Macskassy, S.A.: Discovering users' topics of interest on Twitter: A first look. In: Proc. of 4th Workshop on Analytics for Noisy Unstructured Text Data (AND). pp. 73–80. ACM, New York, NY, USA (2010)
16. Mislove, A., Viswanath, B., Gummadi, K.P., Druschel, P.: You are who you know: Inferring user profiles in online social networks. In: Proc. of 3rd ACM Int. Conf. on Web Search and Data Mining (WSDM). pp. 251–260. ACM, New York, NY, USA (2010)
17. Nechaev, Y., Corcoglioniti, F., Giuliano, C.: Linking knowledge bases to social media profiles. In: Proc. of 32nd ACM Symposium on Applied Computing (SAC). pp. 145–150. ACM (2017)
18. Nechaev, Y., Corcoglioniti, F., Giuliano, C.: SocialLink: Linking DBpedia entities to corresponding Twitter accounts. In: Proc. of 16th Int. Semantic Web Conf. (ISWC). Springer (2017)
19. Piao, G., Breslin, J.G.: Exploring dynamics and semantics of user interests for user modeling on Twitter for link recommendations. In: Proc. of 12th Int. Conf. on Semantic Systems (SEMANTICS). pp. 81–88. ACM, New York, NY, USA (2016)
20. Piao, G., Breslin, J.G.: Inferring user interests for passive users on Twitter by leveraging followee biographies. In: Proc. of 39th European Conference on IR Research (ECIR). pp. 122–133. Springer, Cham (2017)
21. Preotiuc-Pietro, D., Lampos, V., Aletras, N.: An analysis of the user occupational class through Twitter content. In: Proc. of 53rd Annual Meeting of the Association for Computational Linguistics (ACL). pp. 1754–1764. The Association for Computer Linguistics (2015)
22. Rao, D., Yarowsky, D., Shreevats, A., Gupta, M.: Classifying latent user attributes in Twitter. In: Proc. of 2nd Int. Workshop on Search and Mining User-generated Contents (SMUC). pp. 37–44. ACM, New York, NY, USA (2010)
23. Sadilek, A., Kautz, H., Bigham, J.P.: Finding your friends and following them to where you are. In: Proc. of 5th ACM Int. Conf. on Web Search and Data Mining (WSDM). pp. 723–732. ACM, New York, NY, USA (2012)
24. Siehndel, P., Kawase, R.: Twikime!: User profiles that make sense. In: Proc. of 11th Int. Semantic Web Conf. (ISWC) – Posters & Demonstrations Track. vol. 914, pp. 61–64. CEUR-WS.org, Aachen, Germany (2012)
25. Zarrinkalam, F., Fani, H., Bagheri, E., Kahani, M.: Inferring implicit topical interests on Twitter. In: Proc. of 38th European Conf. on IR Research (ECIR). pp. 479–491. Springer, Cham (2016)
26. Zheleva, E., Getoor, L.: To join or not to join: The illusion of privacy in social networks with mixed public and private user profiles. In: Proc. of 18th Int. Conf. on World Wide Web (WWW). pp. 531–540. ACM, New York, NY, USA (2009)