Chapter 6

Concealing Interests of Passive Users in Social Media

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 chapter, we present a system that leverages the links provided by the SocialLink dataset (Chapter 4) to help social media users to conceal 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. We demonstrate that SocialLink allows novel techniques to be developed that can protect user's digital identity from profiling.

6.1 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 (Besel et al., 2016; Faralli et al., 2015b; Piao and Breslin, 2017; Zheleva and Getoor, 2009) have demonstrated that, despite their best effort, passive users can still be profiled based on location and the profiles they follow, information that is typically publicly available. The pipelines, that exploit the list of followees to infer user interests, are based on the idea initially introduced by Besel et al. (2016). There it was demonstrated that the followee list could be mapped to a distribution of interest categories for the user by linking followees to their corresponding DBpedia/Wikipedia entries — a task that can greatly benefit from a resource like SocialLink— and then exploit the categorical information therein contained to derive an interests distribution for the target passive user. 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 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 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 required for the correct operation of the system (e.g., to train it), but can also be used to perform inference of user profile attributes, in many cases, without the user even being aware of such usage.[1]

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 privacy without degrading the user experience by implementing various techniques, such as differential privacy (Dwork, 2008) or by shifting the computation on the user-controlled device (Luo et al., 2009). 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 chapter, 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 (Nechaev et al., 2017b,d, 2018b, Chapter 4) and the categorical information available about those entities in their knowledge base entries. First, we show how to exploit SocialLink to create a state-of-the-art user's interest inference pipeline mirroring the approaches in the literature that use the list of followees (Besel et al., 2016; Piao and Breslin, 2017). Specifically, we use the high-quality alignments of SocialLink to map Twitter profiles to DBpedia resources, in place of the simple heuristics typically used in state-of-the-art approaches. We then use a custom mapping procedure to obtain the

[1]See, for example, the recent controversy regarding the usage of user's location (http://fortune.com/2019/01/04/la-ibm-weather-channel-app).

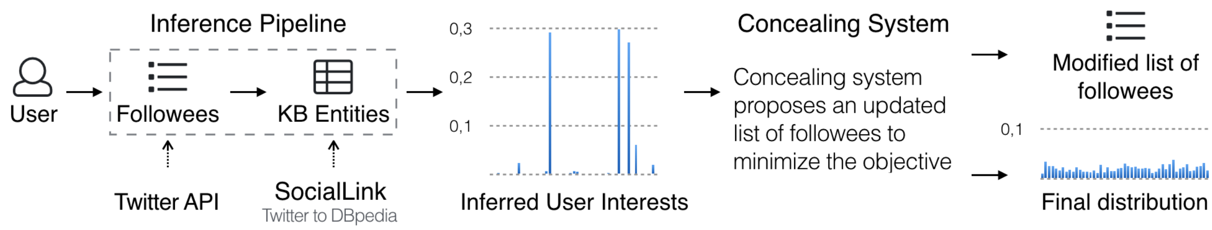

Figure 6.1. The proposed concealing approach

distribution across a taxonomy of 49 interest for each individual followee, which are finally combined to acquire the top interests for the target user.

Based on the interests inference pipeline build on SocialLink, our main contribution here 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 6.1. This task of concealing user interests is inspired by the obfuscation examples provided in Brunton and Nissenbaum's book (Brunton and Nissenbaum, 2015), specifically, the "Bayesian Flooding" idea by Kevin Ludlow (Ludlow, 2012).

We evaluate our Greedy and Joint approaches and a Random baseline against our interest 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 performance 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 real world setting by evaluating them against the Twitter's *Who To Follow* (Gupta et al., 2013) 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 chapter is structured as follows. In Section 6.2, we briefly present the related work. Then, we formally define our problem in Section 6.3, followed by the description of the user's interest inference pipeline used as reference in Section 6.4. Our concealing approaches are described in Section 6.5. Finally, we report the evaluation results in Section 6.6, and we conclude in Section 6.7.

6.2 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 (Zheleva and Getoor, 2009), age (Rao et al., 2010), location (Sadilek et al., 2012), political affiliation (Zheleva and Getoor, 2009), level of education (Li et al., 2014), and occupational class (Preotiuc-Pietro et al., 2015), has been studied both for active (Abel et al., 2011; Michelson and Macskassy, 2010; Mislove et al., 2010; Piao and Breslin, 2016; Siehndel and Kawase, 2012; Zarrinkalam et al., 2016) and passive users (Besel et al., 2016; Faralli et al., 2015b; Piao and Breslin, 2017, 2018; Zheleva and Getoor, 2009).

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. (2012). 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 (2009) 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. (2016) 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 Wikipedia category information is used to determine interests. Our inference pipeline is constructed following their approach. Piao and Breslin (2017) 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 useful 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. (2014) 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 (2015), for instance, produced a comprehensive study on privacy preservation and the technological challenges arising in pervasive systems such as social media.

Felt and Evans (2008) 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. (2009) 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 third parties from accessing the sensitive data of the target user but also the social media themselves.

Ludlow (2012) 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 do not employ neural networks in this chapter, our Joint approach was influenced by recent studies on Binarized Neural Networks (BNN) (Bengio, 2013; Hubara et al., 2016). We used the binarization and back-propagation procedures from Hubara et al. (2016) to find an optimal solution to our optimization problem.

6.3 Problem Definition

The goal of this chapter 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 the 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 number 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*[2] between the $\mathbf{c}^{u'} = g(l_{u'})$ and the uniform distribution over possible interest categories:

$$D(u') = D_{KL}(\mathcal{U}\{1,n\}||\mathbf{c}^{u'}) = -\sum_i \frac{1}{n} \log c_i^{u'} + \sum_i \frac{1}{n} \log \frac{1}{n} = -\sum_i \frac{1}{n} \log c_i^{u'} - \log n \quad (6.1)$$

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:

$$J(u', u) = (1-\alpha)D(u') + \alpha(|l_{u'}| - |l_u|), \quad \alpha \in [0, 1) \quad (6.2)$$

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

where α 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.

6.4 Interests Inference Pipeline

To develop our concealing approaches (presented in Section 6.5) we have implemented a user's interests inference pipeline ($g(l_u)$ in Section 6.3) that infers user's interests based on the list of followees. We follow Besel et al. (2016) 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, SocialLink, which we have introduced in details in Chapter 4, and pre-computing the category distributions over a taxonomy of 49 top categories for all possible entities in English DBpedia/Wikipedia. This enables 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.

Link followees to DBpedia/Wikipedia Each followee profile $f \in l_u$ is linked using the SocialLink resource described in Chapter 4. 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 from SocialLink $v2.0$[3] by setting custom conservative thresholds on confidence scores (with respect to the default ones used in Chapter 4) 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 Besel et al. (2016), 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

[3] SocialLink $v3.0$ was not available at the time our experiments were conducted.

valid probability distribution across possible interests, where the normalized score $\sigma(\mathbf{f})_i$ for interest category i is computed as follows:

$$\sigma(\mathbf{f})_i = \frac{z(f_i)}{\sum_{k=1}^{n} z(f_k)}, \quad z(x) = \begin{cases} e^x & \text{if } x \neq 0 \ , \\ 0 & \text{if } x = 0 \end{cases} \tag{6.3}$$

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 $f \in l_u$:

$$\mathbf{c}^u = \frac{\sum_{f \in l_u} s_f \sigma(\mathbf{f})}{\sum_{f \in l_u} s_f} \tag{6.4}$$

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]

6.5 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 (6.2): 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.

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 6.2 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

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

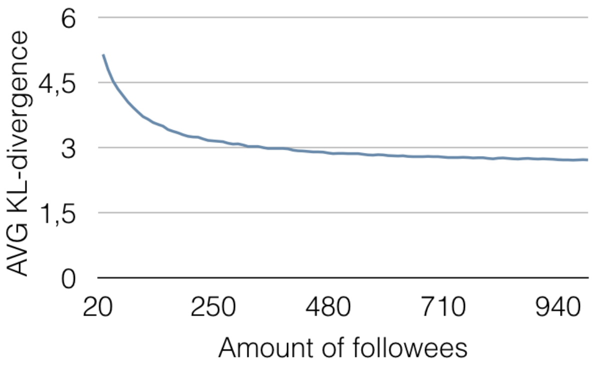

Figure 6.2. Average KL-divergence for different amounts of followees using Random approach

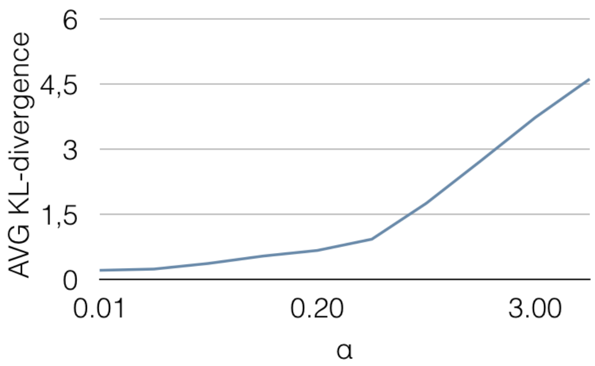

Figure 6.3. Average KL-divergence for different α values using Greedy approach

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 (6.2). In our experiments the α parameter is set to 0.01. Figure 6.3 shows how the average KL-divergence changes based on the choice of α.

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 (Hubara et al., 2016) 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 (6.3), the $\mathbf{c}^u$ can be rewritten in terms of our binary mask $\mathbf{w}^b$:

$$\mathbf{c}^u = \frac{1}{\sum_i \mathbf{w}_i^b} \mathbf{w}^b F \tag{6.5}$$

The objective score can be computed as in (6.2). 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 using a simple deterministic binarization:

$$w_j^b = Bin(w_j) = \begin{cases} 1 & \text{if } w_j \geq 0 \text{ ,} \\ 0 & \text{if } w_j < 0 \end{cases} \tag{6.6}$$

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 (2013). The $\mathbf{w}$ parameters are initialized by drawing from $\mathcal{N}(2l, l)$, where l is the learning rate,

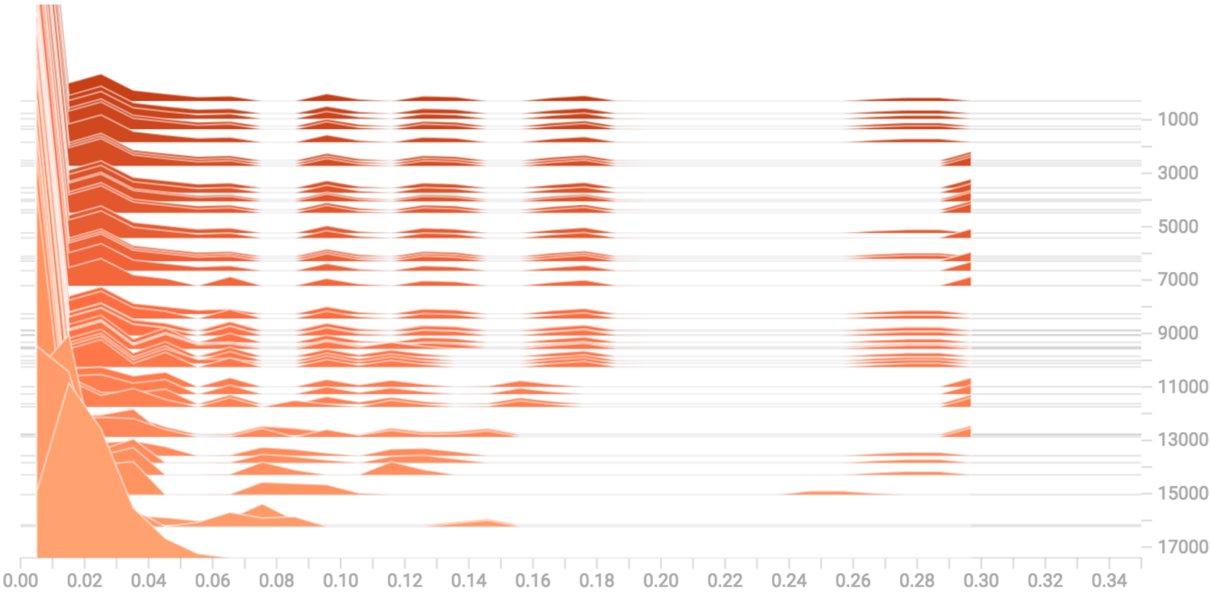

Figure 6.4. An example histogram of user's categories converging towards uniform distribution over 17k iterations of Joint approach. Each slice represent score distribution among categories for the corresponding iteration. In a perfect scenario all scores should be equal to $1/n = 0.02$.

ensuring that the majority of possible followees are followed at the first iteration. Figure 6.4 shows how a user's category distribution changes during the learning process.

6.6 Evaluation

We evaluate our concealing approaches Random, Greedy, and Joint against the user's interests inference pipeline of Section 6.4 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.

6.6.1 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 from a three year period (2014-2016), extracting the names, titles and abstracts of all the papers associated with each author. After gathering the list, the Social Media Toolkit (see Chapter 7) was used to find the corresponding Twitter profile for each author, leveraging the collected textual context. This

Table 6.1. System evaluation against our interests inference pipeline

Approach	Suggestions	System performance (P@N) Top 5	Top 10	Top 15	Score diff to 2nd best	Diff from uniform (KL-div)
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**

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 6.1 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 α parameter. However, we consider the current configuration to be a reasonable tradeoff.

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

Table 6.2. System evaluation against Twitter's Who To Follow

6.6.2 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 (Gupta et al., 2013).[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 (6.4). Then the same metrics used in Section 6.6.1 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 6.2 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 chapter.

[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.

6.7 Conclusions and Future Work

In this chapter, we have presented an application of SocialLink related to the inference and concealment of the passive user's digital footprint on social media. Specifically, we have shown how the high-quality Twitter-DBpedia alignments provided by SocialLink can be used to design a state-of-the-art user interest inference pipelines, and based on the same alignments, we have proposed techniques for concealing the user's interests. We have shown that by using simple techniques together with the SocialLink resource and without degrading user experience, passive Twitter users can prevent the network or a third party system from inferring 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.