

A Discussion of Privacy Challenges in User Profiling with Big Data Techniques: The EEXCESS Use Case

Omar Hasan¹, Benjamin Habegger¹, Lionel Brunie¹, Nadia Bennani¹, Ernesto Damiani²

¹ University of Lyon, CNRS, INSA-Lyon, LIRIS, UMR5205, F-69621, France
{omar.hasan, benjamin.habegger, lionel.brunie, nadia.bennani}@insa-lyon.fr

² Department of Computer Technology, University of Milan, Italy
ernesto.damiani@unimi.it

Abstract—User profiling is the process of collecting information about a user in order to construct their profile. The information in a user profile may include various attributes of a user such as geographical location, academic and professional background, membership in groups, interests, preferences, opinions, etc. Big data techniques enable collecting accurate and rich information for user profiles, in particular due to their ability to process unstructured as well as structured information in high volumes from multiple sources. Accurate and rich user profiles are important for applications such as recommender systems, which try to predict elements that a user has not yet considered but may find useful. The information contained in user profiles is personal and thus there are privacy issues related to user profiling. In this position paper, we discuss user profiling with big data techniques and the associated privacy challenges. We also discuss the ongoing EU-funded EEXCESS project as a concrete example of constructing user profiles with big data techniques and the approaches being considered for preserving user privacy.

Keywords—User profiling, recommender systems, big data, privacy, EEXCESS.

I. INTRODUCTION

A user profile is a collection of information that describes the various attributes of a user. These attributes may include geographical location, academic and professional background, membership in groups, interests, preferences, opinions, etc. User profiling is the process of collecting information about a user in order to construct their user profile.

User profiles are utilized by a variety of web based services for different purposes. One of the primary uses of user profiles is for recommendation of items, elements or general information that a user has not yet considered but may find useful. General purpose social networks such as Facebook.com use a user profile to find potential friends based on the existing relationships and group memberships of the user. Professional social networks such as LinkedIn.com exploit the skills and professional background information available in a user profile to recommend potential employees. Search engines such as Google.com use the history of user searches to personalize the current searches of the user.

Big data techniques are a collection of various techniques that can be used to discover knowledge in high volume, highly dynamic, and highly heterogeneous data. Big data techniques offer opportunities for user profiling that can result in very comprehensive user profiles. Big data techniques have two strengths in particular that enable collecting accurate and rich information for user profiles: (1) Big data techniques process

unstructured data as well as structured data. Unstructured data of different varieties generated by users is growing in volume with high velocity and contains lots of useful information about the users. (2) Big data techniques can process high volume data from multiple sources. This enables linking user data from different sources and aggregating them into a single user profile. Moreover, user information from different sources can be correlated to validate or invalidate the information discovered from one source.

On one hand, user profiling with big data techniques is advantageous for providing better services as we have discussed above. On the other hand, user profiling raises a significant threat to user privacy. One can assume that an ethical and trustworthy service would use the information collected in a user profile with the user's explicit consent and only for the benefit of the user. However, services that are less inclined toward protecting user privacy, may use user profiles for a number of purposes which may not be approved by the user and which may result in disclosure of personal information. One example is the utilization of user profile data for targeted advertising [1]. Another example is the selling of personal information in user profiles to third parties for profit. The third parties may then use this private information for commercial or even malicious purposes [2]. Privacy breaches may occur even when a service is willing to protect a user's privacy [3].

The ongoing EU-funded EEXCESS (eexcess.eu) project aims to improve user recommendations by making intensive use of user profiling and therefore collecting detailed information about users. The EEXCESS project has to address various privacy challenges which appear mainly due to the use of big data and related technologies. One of the major challenges is that the EEXCESS architecture is based on a federated recommender system in which future partners may join. The trustworthiness and the intent of these partners are not necessarily known. The information collected and disclosed to recommenders may not, in itself, be sensitive, however, cross-referencing it with external big data sources and analyzing it through big data techniques may create breaches in user privacy. Since, untrustworthy partners may have access to such big data sources and techniques, privacy becomes a clear challenge.

In this position paper, we highlight some of the private content contained in user profiles, the big data techniques that can be used to construct user profiles, and the ongoing research work toward addressing the associated privacy challenges. In particular, we consider the EEXCESS project as a use case.

We present the proposed EEXCESS architecture, the privacy goals, and the approaches being considered to achieve those goals.

II. USER PROFILE CONTENTS, BIG DATA TECHNIQUES, AND PRESERVATION OF PRIVACY

A. User Profile Contents

The information contained in a user profile can be provided explicitly by the user or alternatively it can be either inferred or mined by the service that manages the profile. Gathering accurate, precise, and rich information is clearly the objective when building a user profile. More and accurate information about a user can indeed help services provide better recommendations.

The most common contents of a user profile include: user interests; user knowledge, user background and skills; user goals; user behavior; user individual characteristics; and user context [4]. We briefly summarize these attributes below. An extended description can be found in [4].

We invite the reader to note while going through the descriptions below that each of these attributes can be considered as a user's private information. In Section II-C, we will further discuss the privacy issues in the context of user profiling.

Interests. The information that can be recorded under this attribute includes a user's professional interests, his interests in hobbies, his interests in entertainment such as music, cinema, books, etc., his interests in sports, as well as his interests in commercial products. If the interests of a user are known, a recommender system can use this information to recommend items that are of the highest interest to the user.

Knowledge, background and skills. This attribute can be used to quantify the knowledge of the user in a given domain. For example, the knowledge that a student has acquired by taking an online course could be measured and recorded. Moreover, professional expertise and skills can also be rated. This information can be used to discover experts in a given domain or conversely to rule out individuals whose knowledge, background, or skills do not correspond to a particular task.

Goals. The goals and intentions of a user represent what he wishes to achieve in a given context. Goals can be classified as short term and long term. For example, a short term goal of a student could be to obtain a high grade in a class whereas a long term goal could be to graduate from college. A recommender system can try to predict the needs of a user given his short term and long term goals and intentions.

Behavior. Users often have repetitive behaviors that can be observed and stored in their user profiles. For example, a user may order pizza online most Tuesdays and purchase an online movie most Fridays. Given this information, a recommender system could suggest pizza deals to the user on Tuesdays and new movies on Fridays. The history of user actions may also be considered under this attribute.

Individual characteristics. The individual characteristics of a user that may be made part of their user profile include

personal information such as age, gender, relationship status, address, etc. Knowledge of demographic information is useful information for a recommender system. For example, attributes such as age, gender, and address can have a strong impact on the movies that a person views and likes and thus on the recommendations that should be made.

Context. The different types of contexts include environmental contexts, personal contexts, social contexts, and spatio-temporal contexts. Entities that are located in the vicinity of the user form his environmental context, e.g., things, services, temperature, light, humidity, noise, and persons. Personal context comprises of physiological contexts, such as weight, pulse, blood pressure, hair color, etc., as well as mental contexts, such as mood, stress level, etc. Social context can comprise of information such as friends, neighbors, co-workers, and relatives. Spatio-temporal information is a combination of time, location, and the direction of movement.

We observe again that each of the attributes contained in a user profile described above can be considered as private information. For example, a person may not wish to share information regarding his whereabouts at all times with everyone. Similarly, it may be detrimental for a person to reveal her interests, goals, behavior etc. and thus she may not wish to divulge this information.

B. Big Data Techniques for User Profiling

We list below some of the big data techniques that can be used for collecting information about a user and building a user profile. An extended list of big data techniques that can be used for user profiling can be found in [5].

It can be noted that many big data techniques are in fact adapted from artificial intelligence and graph theory. However, they take into consideration the added constraints implied by big data: the massive amount of data that often requires distribution over multiple servers or clusters, and the diversity of such data. Many existing big data implementations are algorithms adapted for distributed computation platforms such as Hadoop (hadoop.apache.org).

Network analysis. Network analysis algorithms are used to discover relationships between the nodes in a graph or a network. Network analysis is particularly useful in the context of social networks where important information about the user such as his friends, co-workers, relatives, etc. can be discovered. Social network analysis can also reveal central users in the network, i.e., users who exert the most influence over other users. This information can be used to populate the attributes of social and environmental contexts, individual characteristics, etc. in a user profile.

Sentiment analysis. Sentiment analysis is a natural language processing technique that aims to determine the opinion and subjectivity of reviewers. The Internet is replete with reviews, comments and ratings due to the growing popularity of web sites such as Amazon.com, Ebay.com, and Epinion.com where users provide their opinion on others users and items. Moreover, micro-blogging sites

such as Twitter.com and social network sites such as Facebook.com also hold a large amount of user opinions. The goal of sentiment analysis is to classify user opinions. This classification may be a simple polarity classification, i.e., negative or positive, or a more complex one, e.g., multiple ratings. Sentiment analysis can be used to process unstructured text written by a user to discover their interests, opinions, preferences, etc. to be included into their profile.

Trust and reputation management. Trust and reputation management is a set of algorithms and protocols for determining the trustworthiness of a previously unknown user in the context of his reliability in performing some action. For example, a reputation management system could be used for computing the trustworthiness of an online vendor who may or may not deliver the promised product once he receives payment. The reputation of a user is computed as an aggregate of the feedback provided by other users in the system. Trust and reputation information can be an important part of a user profile. It can convey the user's trust in other users as well as his own reputation in various contexts. This information can be subsequently used as a basis for recommending trustworthy users and avoiding those who are untrustworthy. Trust and reputation management systems can function in conjunction with sentiment analysis for obtaining user opinions and then computing trustworthiness and reputation.

Machine learning. Machine learning is a sub-field of artificial intelligence that aims to build algorithms that can make decisions not based on explicit programming but instead based on historical empirical data. An example often cited is the algorithmic classification of email into spam and non-spam messages without user intervention. In the context of user profiling, machine learning can be used for learning user behavior by identifying patterns. Topics in machine learning include: supervised learning approaches, e.g., neural networks, parametric/non-parametric algorithms, support vector machines, etc.; and unsupervised learning approaches, e.g., cluster analysis, reduction of dimensionality, etc.

Cluster analysis. Cluster analysis is the process of classifying users (or any other objects) into smaller subgroups called clusters given a large single set of users. The clusters are formed based on the similarity of the users in that cluster in some aspect. Cluster analysis can be applied for discovering communities, learning membership of users in groups, etc. Cluster analysis can be considered as a sub-topic of machine learning.

C. Preservation of Privacy in Big Data Techniques

Big data techniques offer excellent opportunities for more accurate and richer user profiling. However, privacy is an issue that can hinder acceptance by users of user profiling with big data techniques. Therefore, there is a need to develop big data techniques that can collect information for user profiles while respecting the privacy of the users. Such privacy preserving big data techniques for user profiling would raise the confidence of users toward collection of their personal information.

There is a significant amount of research currently in progress to achieve the goal of preserving user privacy while collecting personal information. As an example, we cite the field of privacy preserving reputation management. A privacy preserving reputation management system operates such that the opinions used to compute a reputation score remain private and only the reputation score is made public. This approach allows users to give frank opinions about other users without the fear of rendering their opinions public or the fear of retaliation from the target user. Privacy preserving reputation management systems for distributed environments have been investigated since long [6], however, they pose scalability problems as they require large-scale handling of rapidly changing pseudonyms.

Privacy preserving reputation management systems for centralized environments include those by Kerschbaum [7] and by Bethencourt et al. [8]. The system by Kerschbaum introduces the requirement of authorizability, which implies that only the users who have had a transaction with a ratee are allowed to rate him even though rating is done anonymously. Bethencourt's system lets a user verify that the reputation of a target user is composed of feedback provided by distinct feedback providers (implying no collusion) even when users are anonymous. Hasan et al. [9], [10] propose privacy preserving reputation management systems for environments where the existence of centralized entities and trusted third parties cannot be assumed. Current privacy preserving reputation management systems still face a number of open issues. These include attacks such as self-promotion and slandering, in which a user either submits unjustified good opinions about himself or unwarranted bad opinions about a competitor.

Differential privacy, introduced by Dwork et al. [11], is a recent approach to preserving privacy that has received significant attention. It provides a mathematical process for adding randomness to statistical queries with a quantifiable degree of privacy for individuals joining a database. The framework offers guarantees on the risk of joining a statistical database. However, in practice, differential privacy can render some subsets of the randomized data less useful while poorly preserving the privacy of specific individuals. This has been demonstrated for instance in [12]. Thus, privacy preserving techniques still have much to achieve in order to render personal information of users truly private.

Another well-known approach in privacy preservation of published data is *k-anonymity* [13]. It relies on the distinction of quasi-identifiers and sensitive attributes. Quasi-identifiers are the attributes allowing to determine the identity of the individuals referred to by a record (e.g. age, gender, city). The sensitive attributes (e.g. a disease) are those which should not be linkable to the individuals. A set of records V is said to satisfy *k-anonymity* for the set A_q of quasi-identifiers if for every tuple $t \in V$ there exists $k - 1$ distinct records v_i ($i \in [1, k - 1]$) such that $\forall i \in [1, k - 1] \pi_{A_q}(t) = \pi_{A_q}(v_i)$ (where $\pi_A(r)$ denotes the projection of record r on the attribute set A).

III. THE EEXCESS PROJECT

EEXCESS (Enhancing Europe's eXchange in Cultural Educational and Scientific resources) (eexcess.eu) is a European Union FP7 research project that commenced in February 2013. The project consortium comprises of INSA Lyon

(insa-lyon.fr), Joanneum Research (joanneum.at), University of Passau (uni-passau.de), Know-Center (know-center.tugraz.at), ZBW (zbw.eu), Bit media (bit.at), Archäologie und Museum Baselland (archaeologie.bl.ch), Collections Trust (collection-trust.org.uk), Mendeley (mendeley.com), and Wissenmedia (wissenmedia.de). In this section we present the EEXCESS project to illustrate how user profiling can benefit recommender systems particularly with the use of big data techniques. We also discuss the associated privacy issues and the approaches currently being considered in the project for tackling the privacy problem.

The main objective of EEXCESS is promoting the content of existing rich data sources available throughout Europe. While user context is more and more present, the current response of web search engines and recommendation engines to the massive amount of data found on the web has been to order query results based on some form of popularity. It is evident that the introduction of PageRank [14] in search engines has changed the landscape of online searching. However, this has led to the effect of having large quantities of valuable content remaining simply unaccessed due to low levels of global popularity but at the same time being of high interest for a particular user. This unseen data is sometimes referred to as “long-tail content” in reference to the long-tail of a power-law distribution which in many cases characterizes the distribution of user interest in particular content.

It is this type of long-tail content that some of the EEXCESS partners are providing. This includes precise and rich content such as museum object descriptions, scientific articles, business articles, etc. Currently, this very specific content has trouble finding appropriate visibility, even though they would be invaluable in the appropriate contexts where fine-grained and precise information is sought for.

The aim of EEXCESS is to push such content made available by its partners to users when appropriate for them. However, this relies on having a precise understanding of a given user’s interests and their current context. Different levels of user profiling can help to characterize a user’s interests. In EEXCESS, precise user profiles will allow recommending the appropriate content found in multiple data sources.

A. Architecture

Figure 1 gives a sketch of the currently envisioned architecture for the EEXCESS project from a privacy perspective. From this perspective, EEXCESS is made of four components: (1) A plugin added to the user’s client whose role is to collect and transfer the user’s context, trigger recommendation requests and render them through rich visualizations, (2) a privacy proxy which collects the user’s privacy policy and ensures that it is respected, (3) a usage mining component allowing to identify common usage patterns and enrich user profiles accordingly, and (4) a federated recommender service composed of individual data-sources hosting a specific data collection. The circled numbers on the figure give the information flow when content is being recommended.

As suggested by the presence of a privacy-proxy, one major goal in EEXCESS is to respect its users’ privacy. In particular, no information about a user profile data should leak out of the system without the user’s consent. As will be

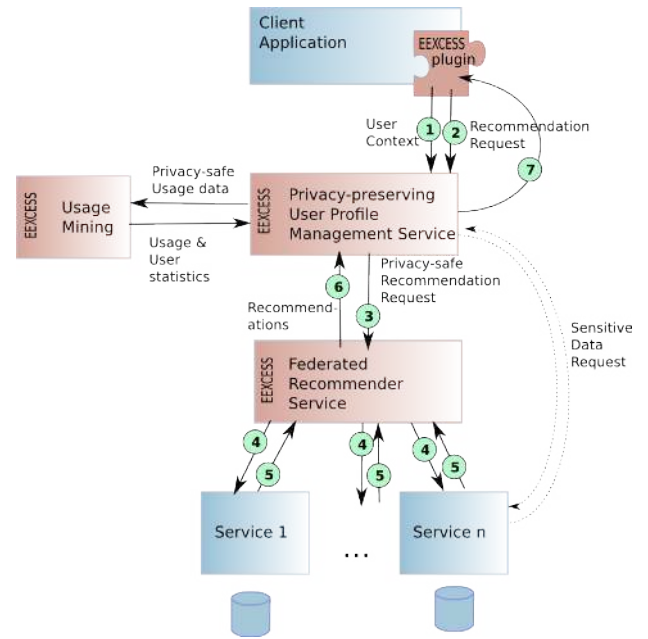


Fig. 1. EEXCESS architecture from a privacy perspective

discussed later, the project is faced with a conflicting situation in which disclosing more information will allow to improve recommendation quality but will also augment the risk if privacy leaks. The exact internals of the privacy proxy are among the works to be completed during the project’s time span. For simplicity, we consider the proxy-service as a single peer in this paper.

Let us consider a typical EEXCESS user scenario. Alice is an economist employed by a consulting firm. She is currently working on a business plan for one of her customers on a market which is new to her. As usual she uses her favorite search engine to investigate on the different actors of the market and in particular the potential competitors for her client. Fortunately, EEXCESS is connected to an economic database, and starts pushing to Alice relevant content from this database, which includes detailed descriptions of companies found in the target market of her client and strategic economic data. Alice requires that a high level of privacy is ensured by the system. In fact, she is legally-tied by a non-disclosure policy with her customer. In particular, it should not be learned that Alice’s customer is taking a move toward the new market.

B. User Profiling

One of the major objectives of EEXCESS is providing its users with quality recommendations. To this extent, fine-grained user-profiling will be an important part of the project and will consist of collecting sensitive data about the user. Many of the attributes discussed in section II-A will be collected or enriched using big data techniques described in section II-B.

Of course, the user’s *individual characteristics* will be part of his profile. An EEXCESS user’s *interests* will either be interactively collected and/or completed using big data techniques implemented particularly by the usage mining service. User actions will be tracked by the EEXCESS plugin allowing

to keep track of a user's *behavior*. Among the partners of EEXCESS, Bit Media is an e-learning platform. In this case, it is clear that the user's learning *goals* and current *knowledge* (e.g. in the form of courses already taken) will be part of the user's profile. In EEXCESS, the user's *context* will consist of information such as his current geo-location, the document or web page (both URL and content) he is working on, his browsing history, the navigation page which lead to the current page, etc.

To capture an even better understanding of the user, different big data techniques will be applied to further enrich his profile. For example, usage mining will try to identify usage trends, as well as information about the user's unexpressed goals and knowledge. On-the-fly analysis of user interests, context, and expectations is also planned. Clustering techniques may be used to identify communities within EEXCESS users. This profiling and better understanding of the user has a unique goal in EEXCESS of providing the user a personalized experience of the system and in particular *personalized recommendations*. Indeed, the content of the EEXCESS partners being very specific (i.e. being in the long-tail of documents when ordered by popularity), having a fine-grained understanding of EEXCESS user's is essential to link the correct users to the correct content.

In our example, the EEXCESS system will have collected significant information about Alice: her interests (economic information), some comprehension of her goal (writing a business plan), her knowledge (expert in economics), her context of work (information about her customer, the target market, the information she has already collected, etc.). Knowing as much as possible about Alice and her customer will allow the EEXCESS system to provide her with adapted recommendations. For example, instead of presenting general-purpose information about the market, the system will propose more detailed technical data which Alice needs and understands.

C. Privacy

Providing users with quality recommendations is a seemingly conflicting objective with the equally important goal of privacy preservation. Even a small amount of personal information may lead to identifying a user with high probability in the presence of side channel external data [3].

Returning to our example, it would be unacceptable to Alice that any information about herself or her customer leak out of the system. Alice's project may even be so sensitive that even the fact that *someone* (without particularly knowing who) is setting up a business plan on the target market may be an unacceptable leak because it could lead to competitors taking strategic moves. This emphasizes the fact that preserving only anonymity may not be sufficient in some cases.

Therefore, for EEXCESS to be a success, many privacy-related challenges will have to be addressed.

Providing privacy guarantees. At all levels within the system user privacy guarantees must be given. This is most likely one of the hardest tasks. Indeed, as soon as information flows out of a system, sensitive information leaks become a risk. Solutions which may seem trivial, such as anonymization have been shown to be inefficient. A

well known example showing that simple anonymization is insufficient to protect privacy is the de-anonymization of the data of the Netflix contest [3]. Furthermore, Dwork [11] has shown that the published results of a statistical database may lead to privacy breaches even for users who are not originally part of the database. These examples show the difficulties which will have to be overcome in order to provide a privacy-safe system. Furthermore, these works show that research on privacy has shifted from totally preventing privacy breaches to minimizing privacy risks. One of the difficulties to overcome in the EEXCESS project, is to ensure that the collection of information flowing out of the system to potentially malicious peers, limits the risks in breaching any of the users' policies. It goes without saying that the attackers themselves very likely have access to big data techniques and that this aspect should be taken into account.

Flexible privacy policies. Users are different, in particular with respect to privacy. Some may not have any privacy concerns at all where as others may not want to disclose a single piece of information about themselves. For example, in one hypothesis, our user Alice may simply wish to remain anonymous. In another hypothesis, Alice may not be concerned by her identity being revealed, but wish that some information about her be kept private (e.g. she may wish to keep private that she is affected by a particular disease). One big challenge will be to define a policy model which allows for such flexibility and at the same time allows to ensure the policy is respected. Preventing direct disclosure of information marked private is quite straight forward. However, a real challenge is preventing the disclosure of the same information *indirectly*. Indeed, leaking other non-private information of a user's profile can lead, through inference, to unwanted disclosures.

Evaluating trust and reputation. What user profile information is disclosed, or at which granularity it is disclosed, may depend on the trust (with respect to privacy concerns) that the user and/or the EEXCESS system has in the content provider. Calculating a content provider's reputation and trustworthiness in a privacy preserving manner is thus an issue.

Let us consider the case of a user wishing to remain anonymous to all the recommenders. In this case, the attacker could be one of the content-providers trying to collect information about the user that it receives queries from. The EEXCESS privacy requirements for such a user would include:

Content anonymity. To guarantee privacy, the attacker should not be able to identify the user from the provided data. Therefore, the system should ensure that an attacker cannot deduce from the content of a request who it originated from.

Request unlinkability. If multiple queries can be linked together, even while having content-anonymity for each individual query, the combination of the two could reveal information about the user. Therefore, it should be required that the protocols guarantee that two independent requests originating from the same user are unlinkable.

Origin unlinkability. This should be feasible by anonymizing

the origin of the request but under the condition that the origin is not revealed by the application level protocols. Therefore, we also need to guarantee that the application level protocols are privacy-preserving (i.e. an attacker cannot link a given request to the requesting user).

Respecting these three constraints is an ideal goal which requires limiting the information transmitted in each request. Such limitations have a high impact on the utility of the profile information disclosed. Thus the challenge is more to find a balance between privacy and utility than to ensure complete privacy.

In information systems (such as recommender systems, statistical databases, anonymized datasets), the main goal of privacy preservation is to not reveal sensitive information about a single entity within the underlying data. This has been shown to be a difficult goal [11], [3]. In a survey on privacy in social networks, Zheleva and Getoor [15] describe some of the common approaches for preserving privacy: *differential privacy* and *k-anonymity*. In the context of recommender systems using collaborative filtering, an approach is to use big data techniques such as clustering to group users together in order to provide privacy [16], [17], [18] with the theory of *k-anonymity*.

In our particular setting, we are faced with a federated recommender system in which trusted and untrusted peers may exchange information. This requires that both the protocols for exchanging information and the content disclosed are privacy-safe. Furthermore, recommendations may not always be limited to a single recommendation technique among the peers. Each content source may wish to use its own approach. In the context of EEXCESS, few hypotheses can be made on the computational capacities or the background knowledge that an untrusted peer may have access to.

Our work in the EEXCESS project will include developing mechanisms for the definition of flexible user privacy policies, guarantees based on the user privacy policies for non-disclosure of private information, quantification of the risk of disclosing private information, mechanisms for exchange of information based on the reputation and trustworthiness of partners, as well as the definition of the relationship between the amount of information revealed and the quality of recommendations.

IV. CONCLUSION

In this paper, we discussed the challenges raised when building systems which require at the same time a deep level of user-profiling and a high level of user privacy. Building and disclosing fine-grained user profiles can be highly effective in providing quality recommendations. This is particularly true when recommending long-tail data. Big data techniques play an important role in making these profiles even more specific. On the other hand, this raises the issue of respecting a given user's privacy. Big data may even increase this risk by providing attackers the means of circumventing privacy-protective actions. We illustrated these issues by introducing the challenges raised by EEXCESS, a concrete project aiming both to provide high quality recommendations and to respect user privacy.

ACKNOWLEDGMENT

The presented work was developed within the EEXCESS project funded by the EU Seventh Framework Program, grant agreement number 600601.

REFERENCES

- [1] P. Jessup, "Big data and targeted advertising," <http://www.unleashed-technologies.com/blog/2012/06/28/big-data-and-targeted-advertising>, June 2012.
- [2] J. Yap, "User profiling fears real but paranoia unnecessary," <http://www.zdnet.com/user-profiling-fears-real-but-paranoia-unnecessary-2062302030/>, September 2011.
- [3] A. Narayanan and V. Shmatikov, "Robust De-anonymization of Large Sparse Datasets," in *2008 IEEE Symposium on Security and Privacy (sp 2008)*. IEEE, May 2008, pp. 111–125.
- [4] S. Schiaffino and A. Amandi, "Intelligent user profiling," in *Artificial Intelligence An International Perspective*. Springer, 2009, pp. 193–216.
- [5] J. Manyika, M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A. H. Byers, "Big data: The next frontier for innovation, competition, and productivity," The McKinsey Global Institute, Tech. Rep., May 2011.
- [6] E. Damiani, S. D. C. di Vimercati, S. Paraboschi, P. Samarati, and F. Violante, "A reputation-based approach for choosing reliable resources in peer-to-peer networks," in *Proceedings of the ACM Conference on Computer and Communications Security*, 2002, pp. 207–216.
- [7] F. Kerschbaum, "A verifiable, centralized, coercion-free reputation system," in *Proc. of the 8th ACM workshop on privacy in the e-society (WPES'09)*, 2009, pp. 61–70.
- [8] J. Bethencourt, E. Shi, and D. Song, "Signatures of reputation: Towards trust without identity," in *Proc. of the Intl. Conf. on Financial Cryptography (FC '10)*, 2010, pp. 400–407.
- [9] O. Hasan, L. Brunie, E. Bertino, and N. Shang, "A decentralized privacy preserving reputation protocol for the malicious adversarial model," *IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY*, vol. DOI: 10.1109/TIFS.2013.2258914, 2013.
- [10] O. Hasan, L. Brunie, and E. Bertino, "Preserving privacy of feedback providers in decentralized reputation systems," *Computers & Security*, vol. 31, no. 7, pp. 816 – 826, October 2012, <http://dx.doi.org/10.1016/j.cose.2011.12.003>.
- [11] C. Dwork, "Differential Privacy," vol. 4052, pp. 1–12, 2006.
- [12] R. Sarathy and K. Muralidhar, "Evaluating Laplace Noise Addition to Satisfy Differential Privacy for Numeric Data," *Transactions on Data Privacy*, vol. 4, no. 1, pp. 1–17, Apr. 2011.
- [13] L. Sweeney, "k-ANONYMITY: A MODEL FOR PROTECTING PRIVACY," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, Oct. 2002.
- [14] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Computer networks and ISDN systems*, vol. 30, no. 1, pp. 107–117, 1998.
- [15] E. Zheleva and L. Getoor, "PRIVACY IN SOCIAL NETWORKS: A SURVEY," C. C. Aggarwal, Ed. Boston, MA: Springer US, 2011.
- [16] J. Canny, "Collaborative Filtering with Privacy," p. 45, May 2002. [Online]. Available: <http://dl.acm.org/citation.cfm?id=829514.830525>
- [17] D. Li, Q. Lv, H. Xia, L. Shang, T. Lu, and N. Gu, "Pistis: A Privacy-Preserving Content Recommender System for Online Social Communities," in *2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, vol. 1. IEEE, Aug. 2011, pp. 79–86.
- [18] A. Boutet, D. Frey, A. Jegou, and A.-m. Kermarrec, "Privacy-Preserving Distributed Collaborative Filtering," INRIA, Rennes, Tech. Rep. February, 2013.