EEXCESS

Enhancing Europe’s eXchange in Cultural Educational and Scientific reSources

Deliverable D6.1

Policy Model for Privacy Preservation and Feasibility Report

Identifier:	EEXCESS-D6.1-Policy-Model-for-Privacy-Preservation-and Feasibility-Report-final.odt
Deliverable number:	D6.1
Author(s) and company:	INSA

Internal reviewers:	ZBW
Work package / task:	WP6 / Task 6.1
Document status:	Draft
Confidentiality:	Public
Version	2013-11-26

History

Version	Date	Reason of change
1	2013-10-31	Document created
2	2013-11-26	First draft
3	2013-12-02	Second draft
4	2013-12-06	Internal review
5	2013-12-17	Third draft integrating reviewer comments
6	2013-12-19	Final

Impressum

Full project title: Enhancing Europe’s eXchange in Cultural Educational and Scientific reSources

Grant Agreement No: 600601

Workpackage Leader: Lionel Brunie, INSA Lyon

Project Co-ordinator: Silvia Russegger, Jr-DIG

Scientific Project Leader: Michael Granitzer, Uni-Passau

Acknowledgement: The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 600601.

Disclaimer: This document does not represent the opinion of the European Community, and the European Community is not responsible for any use that might be made of its content.

This document contains material, which is the copyright of certain EEXCESS consortium parties, and may not be reproduced or copied without permission. All EEXCESS consortium parties have agreed to full publication of this document. The commercial use of any information contained in this document may require a license from the proprietor of that information.

Neither the EEXCESS consortium as a whole, nor a certain party of the EEXCESS consortium warrant that the information contained in this document is capable of use, nor that use of the information is free from risk, and does not accept any liability for loss or damage suffered by any person using this information.

Table of Contents

Executive Summary

This document reports the different works lead within the WP6 package throughout first six months of the EEXECESS project. The main objective of this work was to better identify and define the privacy challenges within the EEXCESS project at all levels and to determine work to be done to be able to define a policy model which enables to preserve privacy in the most adapted way.

Many technical challenges arise when it comes to privacy-preservation. Among them, preserving privacy means perturbing user data which may lead to the loss of utility of the recommendations and/or profiling which is done. Therefore a central question is how to limit such utility loss. A linked challenge also rises from the fact that privacy-preservation may depend on how the system is architectured and which of its components are considered trustworthy. Further research is required to settle these questions.

To tackle these questions, we have identified different objectives the privacy policy model should respect to achieve this goal. The system should **empower the user** by providing him control, transparency and feedback on his privacy settings and how his data is being used. It should provide a **balance between privacy and quality** and ensure **privacy for recommendation** but also ensure **privacy for usage-mining**. See chapter Privacy Policy Model Goal and Objectives.

Defining a privacy policy model requires clarifying what is meant by privacy and having a clear understanding of what the privacy preserving system should indeed preserve. Much literature refers to privacy but many different notions have been defined and referred to. Within the literature, some refer to privacy as anonymity while in other situations it refers to preserving a particular link between a known person and a piece of information. This also has an impact when privacy is breached and the attacks leading to such breaches. We have identified mainly two types of attacks we will be considering within the policy model: **protocol attacks** and **statistical attacks**. The latter is challenging as it implies taking into account **external knowledge** an adversary may have access to. Adding to this challenge, privacy must also take the users into account. There are interesting relationships between user information disclosure and privacy. There is therefore also a challenge in helping users **understand both the risks and benefits of disclosure**. All these issues are further discussed in chapter General Privacy Issues.

EEXCESS brings in its own specific issues which need to be taken into account when considering privacy. These specificities and the initial set of architectural choices that have been made (described in more detail in WP1) impose different constraints on the privacy model to be implemented. Among the impacting choices, the direction has been taken to first focus on **query-based recommendation** which will lead initial research on privacy to focus on **privacy preserving search**. In order to not close the personalization capabilities of the federated recommender, it has been decided to send **user profiles to the federator**. Therefore, we will initially consider the **federator as a trusted component** until the divergence of benefits and risks between profile obfuscation and query obfuscation are clarified. We will also investigate possible evolutions, in particular relaxing the trustworthiness hypothesis. These aspects are detailed in chapter EEXCESS Context for Privacy.

Much research has already been lead on privacy in different domains. Differential privacy and k-anonimity appear at the fundamental levels and have particularly well been studied within the domain of data publishing. Much works also exist in the domain of privacy-preserving web search which will be of great interest as relying on query-based recommender systems will be the initial focus. Often privacy relies on finding ways to “hide in the crowd” therefore, much research exists on how to preserve privacy by relying on interest groups. As we will be including users in the privacy process it is interesting to understand works which have a privacy-by-policy approach. Finally, formal privacy policy models exist within the literature and allow ensuring that the policy is provably respected even within complex protocols. All these aspects of related work are further explained in chapter Related Work in Privacy.

Privacy is a transversal issue and therefore there exist many links to other work packages. The main links are described in chapter Links with other EEXCESS Work Packages.

Finally, our work has already lead to different results within the WP6 work package on privacy. We have implemented an initial privacy preserving proxy stub and associated privacy preserving plugin with the objective to better understand the technical impacts of privacy preservation in the EEXCESS context. This has also allowed us to expose an initial version of the EEXCESS privacy proxy API. Furthermore, how trustful the different components of EEXCESS can be considered changes the way privacy is considered. To better understand these impacts we have distinguished the different cases which can be considered. Also, there are different ways privacy preservation measures can be considered in a query-based federated recommendation system. We have modeled different possible strategies which will allow us to better qualify these impacts. All these results are detailed in chapter Current Results of WP6.

Introduction

1 Purpose of this Document

Policy Model for Privacy Preservation and Feasibility Report: The deliverable contains the detail on the policy model to ensure privacy preservation and a feasibility analysis on the implementation of these models.

2 Scope of this Document

This document presents the work having been done within the Work Package 6 (privacy preservation). It gives an overview of the work which has been lead in the direction of defining a policy model for privacy preservation and feasibility. To this effect, it extends the state of the art of building privacy preserving systems and their impacts on the different EEXCESS components.

It does not cover the specificities of each of the components. In particular, it does not cover how recommendation and federated recommendations will be implemented, which is left as work in WP3. Similarly, it does not cover how usage mining will be implemented which is left as work in WP5. Also, the architecture representations used to discuss privacy issues within this document are a simplification of the general EEXCESS architecture discussed within WP1, which remains the official and complete EEXCESS architecture.

However, this document does discuss impacts of privacy on all of these work packages.

More specifically, this document presents the general policy model for privacy of the EEXCESS system. It gives the general objectives the implemented privacy protocols should attain. It translates these objectives into system and protocol requirements with respect to the current state of the art and envisioned future improvements. It also gives a prospective ordering in which the requirements derived from these objectives should be tackled.

3 Status of this Document

This is the final version of D6.1

4 Related Documents

For readability we have tried to keep this document as self-sufficient as possible. It may however be useful to the reader to be familiar with the following other deliverables of the EEXCESS project:

- D1.1 – for the global architectural design
- D3.1 – for a full understanding of how recommendation is planned
- D5.1 – for the user profile definition

Privacy Policy Model Goal and Objectives

Nowadays, faced with the massive amounts of data available, the increasing strength of computing power and increasing performance of data analysis algorithms, privacy has become a very strong issue in many domains. In particular, privacy has been widely studied in statistical databases, recommender systems (Ramakrishnan, Keller, Mirza, Grama, & Karypis, 2001), social networks (Zheleva, Terzi, & Getoor, 2012), distributed and peer-to-peer systems (Ying-hua, Bing-ru, Dan-yang, & Nan, 2011). Building a general system architecture and related exchange protocols which have a concern for privacy at their core can get quite complex. Indeed, privacy is an intrinsically horizontal topic. Every aspect of the system must guarantee that privacy leaks are limited and provide means for end-users to understand what their particular risks are in using the system.

This document gives an overview of the topic, gives insight on how privacy-safe architectures and protocols can be achieved, how privacy impacts the EEXCESS system and defines a policy allowing privacy-related options to be taken by users, by the content-providers, and by the EEXCESS system as a whole.

Ideally the user would like to obtain quality and personalized recommendations without revealing anything about himself. Attaining such an objective means ensuring that a user remains anonymous with respect to the peers he considers non-trustworthy. Works have shown that in some restricted cases anonymization is possible (Canny, 2002; Dwork, 2006, 2008). This however often comes at the cost of quality or utility of the disclosed information (Brenner & Nissim, 2010; Haeberlen & Pierce, 2011; Li, Qardaji, & Su, 2012; Narayanan & Shmatikov, 2008; Sarathy & Muralidhar, 2011). It may also be the case that users do not necessarily require anonymity (for example, in social networks), but rather have control over what is disclosed or not disclosed.

Privacy in EEXCESS raises multiple challenges:

- adapting privacy to user preferences
- limiting the impacts of privacy on result quality
- providing feedback to other components in a privacy preserving way
- conserving privacy in the context of user profiling

1 Empowering the user

Within the EEXCESS context it has been chosen, in a first step, to favor user control over strict absolute privacy. The system will give the user settings which the user has control on but without being able to completely define his own privacy policy. Therefore, the general goal of the policy model described in this document is to ensure EEXCESS users have the appropriate control of the disclosure of their private information throughout the EEXCESS system and to the content providers. This general goal can be split into three main objectives.

- **transparency**: users need to be able to understand what information about them the system and its different components have access to.
- **control**: users need to be able to express their private information disclosure preferences. They should be able to specify what is disclosed and how detailed the disclosure should be as well as to whom it is disclosed.
- **feedback**: users need to understand the impacts of disclosing or not disclosing certain pieces of information.

- **usability**: whether or not users take advantage of the previous three objectives highly depends on how the are implemented and how user-friendly they are.

2 Balancing Privacy and Quality Recommendations

Given the users preferences it is then necessary to have a system capable of optimizing the use of the disclosed information. In EEXCESS, this means that the quality of the recommendations should be maintained as close as possible to that of those the user could have expected with a more detailed profile. It is not clear how much, if at all, obfuscated profiles will impact recommendation quality. Within the the context of data publishing, privacy preserving measures, such as output perturbation, have shown to reduce the utility of the data (Brenner & Nissim, 2010). Work remains to be done on measuring recommendation quality and the impact of privacy measures on this quality.

Furthermore, providing recommendation quality will also rely on user profiling. Such deep user profiling entails many privacy concerns. Indeed, while users are likely to be interested in having very precise recommendations, they may not, at the same time, be willing that a third-party collects private information about them. Therefore providing a policy model for privacy should give insight to the following questions:

- How to collect and store user profile data in a privacy preserving way ?
- How to exchange user-profile data in a privacy preserving way ?
- What guarantees can we give (if any) on sensitive information which does leek to potentially untrusted peers ?

3 Privacy Preserving Profile Usage Mining

To better adapt recommendations to users, usage mining will come in handy to better understand each user and therefore provide them with personalized recommendation. On the other hand, through inference, usage mining may reveal information users have expressed as being private. Also, users may not be aware of what the system has inferred or could infer from the collected usage information.

This raises different challenges to be coped with within the EEXCESS project:

- Inform users about what is being done from a usage-mining perspective
  - What data is collected
  - What is done with the collected data
  - What is done to preserve privacy
- Provide users with tool to understand the risks of disclosure
- Adapt usage-mining in away to minimize disclosure risks
- Give users control on what level of usage-mining if any should be done using his data

General Privacy Issues

As privacy is a broad term, in this chapter we clarify some of the aspects surrounding this notion. In this section, we describe the bounds privacy will have in the EEXCESS context. First, we describe two forms of privacy we will be considering. Secondly, the different types of attacks which will be taken into account. Thirdly, we will describe the constraints on the privacy preserving mechanisms linked to the specificities of the EEXCESS context. Finally, we describe the impacts of allowing users to be full actors of their privacy.

1 What Privacy Should Preserve

Within existing work it is not always clear what is meant by "privacy". Depending on the application and the targeted privacy requirement we can have different levels of information disclosure. Taking as example privacy preserving reputation systems we can have 5 different levels for privacy depending of whether identities, votes and aggregated reputation score are disclosed and linked or not. For example, in the context of calculating the reputation of a user Alice by three other users Bob, Carol and David which respectively have vote +1, +1 and -1, the reputation system may disclose to Alice:

- **Full disclosure**. All tuples (Bob,+1), (Carol,+1), (David,-1) as well as the aggregated score (+1 if sum is used) are known by Alice
- **Permuted disclosure**. All voters Bob, Carol, David are known by Alice as well as the scores but permuted so Alice cannot determine who voted what.
- **Identity disclosure**. All voters Bob, Carol, David known by Alice, however individual votes are hidden and only the aggregated score is known by Alice
- **Vote disclosure**. All votes are known by Alice but votees are hidden.
- **Result disclosure**. No details are disclosed except the aggregated score.
- **No disclosure** An aggregated score for Alice is calculated but she does not have access to it.

These levels of disclosure generalize to other privacy settings such as recommendation.

Also and more generally, these can be subdivided into two privacy objectives.

1.1 User Anonymity

A first objective is preserving user anonymity. In this setting, untrusted peers should not be able to link real users to the requests they receive. For example, if John is navigating on the topic of looms, any request a content provider receive should not be linkable to the real user John. Information such as his IP address, e-mail address, userids on specific platforms (e.g. Gmail identifiers stored in cookies), or any other such information which may help identify John should not be made available.

1.2 Disclosure of Private Information About Known Users

A second objective is preventing the disclosure of private information. If we take the same example of John searching for looms, and John wanting his age to be kept private but doesn’t mind that the peers know he is at the origin of the query. In this case, privacy preservation does not necessarily require anonymity but rather provide guarantees that John’s age will not be disclosed.

2 How Privacy can be Breached

2.1 Protocol Attacks

Protocol attacks are those relying on the fact that since a user wants to obtain an information from a peer, then the peer will have to be contacted by some means. For example, a user wanting to access a web page on looms will have his browser making a request to the hosting server. Having been contacted the server has a trace of the user's IP and knows that this IP has requested the page on looms. Protecting from such attacks can be obtained by using proxies but this just moves the problem of trust from the content provider to the proxy provider. It is then the proxy which must be trusted. This very basic example gives an initial intuition on the fact that protecting from protocol attacks can get complex. Much research has been done on protecting anonymity from such protocol attacks (see chapter Related Work in Privacy for details).

2.2 Statistical Attacks

Statistical attacks are those relying on the information which flows into a given peer. Even if users are protected by a privacy preserving protocol, the data which ends in the hand of a potentially malicious or curious peer may be used to break this anonymity. For example, to be able to find interesting documents for a user, a search engine must be provided with a search query. This query in itself provides information about the user from which it originates (be it only that he is interested in the topic of the query). By correlating information an untrusted peer has collected, it can become possible to de-anonymize the user. Recent work has shown that this is possible, e.g. by using machine learning techniques (see chapter Related Work in Privacy for details).

3 Privacy from a User's Perspective

Privacy is mostly all but a technical issue. Even though most of privacy-related work done within EEXCESS will be related to technical aspects in preserving privacy, it is also important to understand the user's perspective to have adapted solutions. In this section we discuss different issues related to how users perceive privacy and what directions EEXCESS can take.

3.1 Privacy and Information Disclosure

There are important links between how a system handles privacy (and not specifically the technical aspects) and how users disclose information. (Knijnenburg & Kobsa, 2013) have shown that users globally tend to disclose less information when faced with a system explicitly talking about privacy. The interpretation given is that when privacy issues are put in the focus, users tend to become more suspicious and therefore leak less information. This is quite of a paradox as system willing to be transparent about privacy finds itself disavantaged in face of one not mentioning privacy at all. However, the same work studies how to improve disclosure (compared to a system not mentioning privacy issues). Giving the same explanations to everyone will lead to the tendency of users disclosing less because of the invocation of privacy. However, adapting explanations to the users can allow to improve disclosure. For example, within the test groups of (Knijnenburg & Kobsa, 2013), giving an explanation what the data will be used for to men, and giving a high percentage of users having disclosed the information to women, tended to globally improve disclosure.

3.2 Impacts

Flexibility and simplicity

Therefore a system aiming to successfully have its users disclose information willingly and at the same time respect their privacy must have solutions which adapt to them. Furthermore, giving high and precise control to users can on one hand show a will of transparency from the service provider but, on the other hand, may make the system look to complex. Therefore, users should be provided with a system allowing them to set their privacy settings simply but without losing flexibility.

To this effect, users should be able to specify their privacy concerns at a high level, but also be allowed more fine grained settings. To reconcile both, one direction which will be investigated is relying on usage mining done WP5 to infer adapted defaults for users based on their similarity with other users.

Helping understand the risks and benefits

Another important aspect to consider is providing users with elements to understand the effects of disclosing information. As discussed previously, this involves providing the appropriate explanations to the appropriate users. In the specific case of EEXCESS, the objective of user information disclosure is mainly to improve the quality of the recommendations for each user. This can be obtained through a tool allowing to compare results using different privacy settings. We have started working on this aspect as explained in section 1 of chapter Current Results of WP6.

EEXCESS Context for Privacy

1 Specific Constraints for Privacy

In the previous chapter we have seen the different difficulties related to privacy in general. The EEXCESS context brings in different specific privacy related constraints and requirements. In this section, we discuss these constraints and their impacts on privacy.

1.1 Providers as Blackbox Recommenders

Currently, among the content providers, all provide an access to their context in the form of a standard search. Only one of them, namely Mendely, provides collaborative filtering. Therefore, in a first step, the focus has been put on recommendation through search. In the future other forms of recommendation (collaborative filtering and hybrid recommendation) will be envisaged. However, in any case, the content-providers will be considered as black boxes in that it will not be known how recommendation is performed. This has an impact on privacy since the privacy preserving mechanisms in place will have to be general enough to work with different recommendation solutions and cannot be limited to one form.

1.2 Providers with Existing Profiles

Some of the content providers already have user-bases for which they may already have pre-calculated recommendations available. If privacy is limited to anonymization, then, through EEXCESS, users will loose access to those recommendations as the recommenders of these providers will not be aware of the user they are sending recommendations for. Therefore, the privacy solutions studied should go beyond simple anonymity and allow users to at a minimum, specify which providers they trust and are willing to share their information with.

1.3 Providers Needing Feedback to Improve Quality

An important objective of recommender systems is to continuously improve themselves through user feedback. To this effect, it will be necessary for them to have access to such user feedback. However, this feedback should not be a source of privacy leaks. This is a challenge privacy-wise as many attempts to anonymizing recommendation feedback data have failed in that the data could be de-anonymized (Narayanan & Shmatikov, 2008).

2 Current Architectural Choices

The general EEXCESS architectural design and requirements are summarized in deliverable D1.1. In addition to the general architectural choices, a set of initial restrictions have been fixed within the consortium allowing to start working within a common simplified setting. The idea is initiating research within this restricted setting and, once satisfactory results obtained, identify which restrictions could then be lifted. Here we recall the restrictions which have an impact on privacy and discuss their impacts.

2.1 Query-based Recommendation

As most providers already have a search-based access to their content, it has been decided to first focus on search-based recommendation. The idea is to map a user's context and profile into a search query sent to the different recommender systems. Therefore, initial work on privacy preserving recommendations will focus on privacy preserving search. However, as the target focus is sending profiles to the providers, it is not clear whether the privacy-preservation should apply to search queries or directly to user profiles. Discussions and initial results on this topic have already taken place and are related further in this document in section 4 of chapter “Current Results of WP6“.

2.2 Profile Anonymization

As a first focus, it has been decided to initially work on anonymized profiles. Within this setting the role of privacy-preservation is to prevent providers from identifying the users at the origin of the query. As currently, EEXCESS components the current architecture are considered trusted leaving as the only potentially untrusted components the content providers. In such an architecture, securing prototols is straightforward. The remaining risk of privacy breaches therefore comes from statistical attacks which malicious content-providers could attempt ton identify the user behind the requests they receive (see section 2.2 of chapter General Privacy Issues and section 2 of chapter Related Work in Privacy).

2.3 Trusted Federator Does the Profile-to-query Mapping

In this query-based context, and as will be discussed further in section 4 of chapter “Current Results of WP6“, multiple privacy-preservation strategies can be defined. One strategy consists in doing the profile to search-query mapping at the proxy-level while the other consists in doing the same mapping at the federator level. The initial focus will be this latter strategy named strategy S2 in section 4 of chapter “Current Results of WP6“. If the information leaked by profiles reveals to be greater than that revealed by queries, then sending profiles could imply that the federator must be trusted or otherwise require more complex mechanisms to implement privacy preservation. For example, this may imply distributing recommendation and/or using cryptographic techniques.

2.4 Privacy Proxy Hosted on a Single Service

As the privacy proxy is considered a trusted component, it is acceptable that it is deployed on a single host. It could however be considered as too strong constraint. This could be lifted by considering a privacy-proxy distributed among the EEXCESS clients, some of which could be considered untrusted.

2.5 Trusted Usage-mining

The usage mining component of EEXCESS will hold large amounts of data about users. Therefore the acquired reasoning knowledge it will have about EEXCESS users could imply many privacy breaches if it were in the wrong hands. Privacy has impacts on many components and it still remains necessary to clarify what information usage mining will be collecting and generating. In order for this work to mature, the current focus for privacy concerns privacy preserving recommendation. This implies that in the current state, usage mining is considered a trusted component. However, lifting this constraint will be necessary at some point thus will require looking into privacy preserving usage mining techniques.

3 Privacy-impacting Possible Evolutions of EEXCESS

Depending on the results and hypotheses which could be lifted during the EEXCESS project, the architecture may evolve. In this section we describe possible evolutions of the architecture and the impacts these may have from the perspective of privacy.

3.1 Other Types of Recommendation

In the literature, there exist different approaches to recommendation the main being collaborative filtering, content-based (or search-based) recommendation and hybrid recommendation combing both. Furthermore different information about users may be used to make recommendations, explicit interests or preferences, demographic information, object ratings, etc. Most existing privacy preserving recommendation systems handle privacy in a highly integrated way and have important impacts on the general system architecture. Allowing for other types of recommendation will either require architectural changes or new privacy-preservation solutions will have to be investigated.

3.2 Where Profile-to-query Occurs

As will be discussed later within this document, in a search-based recommendation context, where profile-to-query mapping occurs may influence the general level of privacy users can expect to have. Therefore choices on where profile-to-query mapping occurs will influence how privacy-preservation should be done, the attainable levels of privacy and the quality of the recommendations.

3.3 Distributed or Single Host Components

With the current architectural choices, components are considered as single hosts. However, distribution often combined with cryptographic techniques are used in many privacy-preservation recommendation systems as a means to provide better privacy. However, this implies more complex architectures and may impose requirements on the content provider systems that they are not willing to take.

3.4 Trustworthiness of the Components

Which components are considered trusted or not of course has a high impact on privacy. The current setting makes a strong hypothesis on the trustworthiness of the components. This is reasonable, at least in a first step, as the sensitivity of the information manipulated within the EEXCESS context (cultural, scientific and educational content) is limited. However, even within the EEXCESS context, there are situations which privacy could require relaxing the trustworthiness of the components. For example, in an economical context, strategical information could be leaked through the interests of the users. Also, in a medical context, private health status information could required higher levels of privacy.

Related Work in Privacy

Privacy is a very active research topic. In this section we summarize the main works on privacy. We start with the works in data publishing which has been very active, namely in the context of publishing medical data where privacy is a core issue. These works are interesting as they pinpoint the difficulties in having reliable privacy preserving systems. We then focus on privacy preservation in web search which is in phase with the initial working context of EEXCESS. In many contexts, privacy is obtained by mechanisms allowing to "hide in the crowd", therefore, we present works focusing on building interests groups in which user can "hide" to gain in privacy. Finally, we present works introducing systems which have policy-based privacy allowing to include users in the process.

1 Privacy in Data Publishing

1.1 Differential Privacy

In the domain of statistical databases a major shifts occurred with the work of Dwork and the introduction of differential privacy (Dwork, 2006, 2008; Nissim, 2008). Through a theoretical framework the authors demonstrate that, as soon as we consider external knowledge, privacy breaches can occur even for people which do not participate in a statistical database. This has introduced a shift in the way to perceive privacy. The objective is no longer to preserve privacy in an absolute manner, but rather limit the risk of increasing the privacy breach for a participant of a statistical database. To this effect, differentially private mechanisms are those that ensure that the statistical outputs of two databases which are only different by a single participant return similar statistical results. This most often consists in adding sufficient noise to the outputs. Even though there are situations in which differential privacy is attainable, in particular count queries, there are many constraints imposed by differential privacy (Brenner & Nissim, 2010; Sarathy & Muralidhar, 2011). In particular, in situations which should allow multiple queries noise must be augmented proportionally to the number of queries to prevent noise reduction techniques to be applied. However, adding too much noises can deprive the outputs of the system of any utility. Therefore much research is ongoing to evaluate the trade-offs between privacy and utility (Sankar, Rajagopalan, & Poor, 2013). Another commonly cited privacy framework is k-anonimity, which we discuss hereafter. Interestingly, recent work has shown it can be linked with differential privacy certain circumstances (Li et al., 2012).

1.2 K-anonymity

Recommenders need to massively gather past user interactions and their ratings about objects that they were concerned with. This allows them to propose a selection of predicted objects to a current user, based on profile similarity analysis with the current user, using techniques like collaborative filtering. While this allows having good recommendation quality, it also creates user privacy concerns. K-anonymity is one of the well-known techniques to preserve user privacy. The recommender in this case should ensure that each selected object has been selected by at least K users and that each object has been rated similarly by at least K users. This allows avoiding structured-based and label-based attacks respectively (Chang, Thompson, Wang, & Yao, 2010). Several methods have been proposed to ensure k-anonymity among them, we can cite (Ardagna, Livraga, & Samarati, 2012; Chang et al., 2010; Luo, Chen, & Li, 2013; Sweeney, 2002). Many solutions are aimed at resolving k-anonymity problem in databases (Ardagna et al., 2012; Sweeney, 2002). (Chang et al., 2010; Luo et al., 2013) both proposed using k-anonimity for privacy preserving recommenders. In both, past user ratings are represented using a bi-partite graph, where nodes are subdivided into user nodes and object nodes. A graph edge represents the rated selection of an object by a user. Projecting the graph on a single user gives the knowledge that the system has about that user rate and selections. The k-anonymity is obtained then by padding the graph cleverly so that a user clustering with less recommendation accuracy could be obtained. Whereas most solutions proposed for recommenders are based on a centralized gathering of user rates, (Luo et al., 2013) propose a user-centric distributed and anonymous solution to gather useful information to make recommendations.

2 Privacy Preserving Web Search

Nowadays, a common way to find information on the Internet is by using web search engines (e.g. Google, Yahoo, Bing, ...). Consequently these companies can learn a lot of information about their users by knowing their queries. In many cases it is not a problem for a user. For instance, a user looking for "chocolate cake recipe" in Google, will likely not care if Google infers that he is going to cook this cake. However, if he is looking for "Liver cancer consequences for a 60 year old woman", he may be bothered by Google knowing this information. Here we relate the existing research we have studied and whose goal is hiding sensitive information from search engines.

A first idea to protect the user privacy is making it impossible to link users to their queries. (D. Chaum, 1981) suggested a routing protocol (Mix network) to anonymize a connection between two servers. A chain of proxy servers allows to shuffle a message among these servers and consequently the user's IP address is hidden to the search engine. This anonymous communications can be broken only if all servers collude. Today the onion router ("Tor Project," n.d.) improves this protocol and can guarantee the anonymity of the user (protocol-wise). The user encrypts the message multiple times (creating different layers) and sends the result to the destination through a randomly selected Tor relay. Each relay decrypts its own layer to know the next relay in the anonymous network. In the end, the final relay does not know where the message came from. To send the response back to the user, each relay returns the response to the incoming caller.

While these solutions can hide user IP addresses, other information can be leaked by the messages themselves. For example, such information can be found in the data sent by the users web browser (cookies, HTTP headers, ...). An HTTP filtering tool such as ("Privoxy," n.d.) or browser plugging such as ("Fox Tor," n.d.) and ("TorButton," n.d.) can remove this sensitive data before sending a request to the search engine. However, these web proxies do not include sufficient filtering to protect user privacy. Different works (Felten & Schneider, 2000; Jackson, Bortz, Boneh, & Mitchell, 2006) have shown techniques allowing to reveal user identities even with such filtering (cache-timing attacks, visited link tracking, cooperative tracking, ...). A solution to these attacks is given by (Saint-Jean, Johnson, Boneh, & Feigenbaum, 2007). This paper suggests filtering HTML answers in order to remove any component that may provide feedback to the search engine. These schemes (anonymous network + HTTP/HTML filtering) have a huge cost in term of latency. According to (Castellà-Roca, Viejo, & Herrera-Joancomartí, 2009), with such filtering, a query to Google is about 10 seconds on average (e.g. 33 times slower than a direct connection which lasts 0.3 second). Other papers (Castellà-Roca et al., 2009; Lindell & Waisbard, 2010) try designing protocols with a lower computational cost based on Mix Network (D. L. Chaum, 1981) to guarantee user privacy. The scheme proposed by (Castellà-Roca et al., 2009) consists in sending queries from other members of a given group. Indeed, a group of n users exchange their queries in an encrypted way and each peer deciphers the query and sends it to the search engine. The search engine answer is then broadcasted to all group members. (Lindell & Waisbard, 2010) improved the security with a decentralized protocol that they called private shuffle. An advantage of this protocol is that the communication remains secure in the presence of malicious adversaries.

A recent study on anonymous networks (Peddinti & Saxena, 2011) shows that a machine learning process can be used by a search engine to re-identify queries sent by a single user among all the queries coming from the same TOR exit node. This is done by using the content of the messages (rather than who they are coming from) and using classification techniques to re-identify their likely origin. This suggests that anonymous networks or query shuffling to guaranty unlinkability between users and their requests may not be enough.

Another approach to protect user privacy is based on obfuscating user profiles. One of the most famous approaches is TrackMeNot (Howe & Nissenbaum, 2009). This Mozilla Firefox plugin attempts to hide users queries generating additional search queries which simulate the user's behavior. This approach can be considered as a way to get k-Anonymity (Sweeney, 2002). However, this basic idea has been broken with machine learning classifiers (Peddinti & Saxena, 2010) and thus TMN authors released a new version of TrackMeNot (Howe & Nissenbaum, 2009) including more feeds and random clicks on query suggestion.

While TrackMeNot provides a good trade-off for the user, this solution has collateral effects. It degrades : (1) the accuracy of search engines by creating distorted profiling and (2) the performance of the network by increasing the traffic. A similar way to provide privacy in web search is to use Private Information Retrieval (PIR) (Murugesan, 2010). This technique hides the user query among k plausible queries pre-computed on the client side. The query Q=(q1 OR q2 … OR qk) is sent to the search engine and thus, an adversary cannot distinguish which of the k queries is the user's one. To keep a high level of accuracy, results are filtered to discard non relevant answers.

A way to decrease the high computation cost of the previous solution is to forge false queries with an existing database. GooDIR (Domingo-Ferrer, Solanas, & Castellà-Roca, 2009) uses thesaurus to find similar keywords and construct a masked query similarly to TMN. This masked query is sent to the search engine and then the results are re-ranked according to the initial query. To improve the accuracy of this method, the authors of (Sánchez, Castellà-Roca, & Viejo, 2013) suggest constructing the masked query using the semantics of the initial query instead of finding similar words. This scheme enables a higher privacy control because the user's privacy can be adjusted according to the degree of distortion introduced in the masked query.

Finally, (Pang, Ding, & Xiao, 2010) have introduced an approach using similar ideas called query embellishment. It relies on homomorphic cryptography techniques to calculate the scores of queries with the obfuscated queries. The client selects extra terms, called decoys, to be added to the query using Wordnet. Terms are sent accompanied by an encrypted 0 or 1 value if they are respectively a decoy or not. Scores are calculated using these encrypted values. Decrypting the final score, the client has access to the effective score of his query without decoys. However, this approach has the heavy drawback of requiring a modification of the search engine itself and the introduction of the extra level of computation implied by using cryptographic calculation.

3 Privacy by Building Interest Groups

Hiding users in groups of people with similar behavioral characteristics and using them as a way for doing recommendations at the level of such “interest groups“ instead of at an individual level has recently gained much attention in the research community.

The idea of ensuring individual user privacy by aggregating users in groups with common disposition – serving as an “anonymity set“ in which subjects are not identifiable from one another – has first been proposed by (Canny, 2002). He describes a collaborative filtering approach where each community of users computes a public aggregate of the matrix of their ratings, modeled as a partial SVD (Singular Value Decomposition) of their data. This aggregate is calculated iteratively by adding vectors of users' data. Homomorphic encryption is used to allow sums of encrypted vectors to be computed and decrypted without exposing individual data. Based on the group's aggregated profile, personalized recommendations can then be calculated both by the members of the community (intra-group recommendation) and by outsiders (inter-group recommendation).

Canny's idea has been seized by (Nandi, Aghasaryan, & Bouzid, 2011) and (Shang, Hui, Hui, Cuff, & Kulkarni, 2013) which both are presenting solutions of aggregating users' data without the need for computationally expensive cryptography. The privacy preserving method defined by (Nandi et al., 2011) specifies a local profile computation as a first step. Thereby, each client locally collects and analyzes user's activities in order to create a detailed profile. User's preferences are then represented as <key, value>-pairs denoting item categories and ratings that indicate the user's interest level in the corresponding topic. On the basis of this personal information, the interest group the user belongs to is determined, e.g. via clustering or hashing. Subsequently, all the user profiles of each community are aggregated anonymously to form a group member profile. Therefore, each client slices its local profile into several segments of <key, value>-pairs and sends them - including the user's cluster-ID – over an anonymization network like TOR to different profile slice collectors which themselves forward the received segments to the corresponding group-wise aggregator with the aid of a DHT. The aggregator then computes the top-k popular items within the group by concocting all the <key, value>-elements. Thus, a intra-group preference aggregation is performed. The top-k recommendations are finally returned to the requesting client which filters out the list by removing all the items the user has already rated. (Shang et al., 2013) follows a similar approach. The process also starts with a local profile computation where the client locally collects user's traces of activity. The so gathered preferences are then represented in a pairwise (item x item) comparison matrix in which each entry compares two items and signifies with a value of “0” or “1”, respectively, user's each favored item. Afterwards, the client determines the user's interest group. Before uploading his profile to the recommendation service provider, the user exchanges some of his preference information with other group members in a P2P-fashion in order to prevent the central server from learning his individual profile information. Hence, the server receives not a full rating profile, but only perturbed user data. By computing the Kemeny ranking, the server then determines heuristically the top-k popular items within the group. Besides this intra-group recommendation, Shang's approach also proposes a way to include recommendations from other groups (inter-group recommendation). Therefore, a recommendation graph is constructed on which a random walk based algorithm is performed. Finally, a local recommendation personalization is required. Thus, the server responds with a list consisting of the top-k recommendations which is then filtered out by the client to remove all the items the user has already rated.

The key challenge in doing recommendations on the level of interest groups is to determine an appropriate group size. One the one hand, in order to hide individual user's preferences within a community by computing an aggregate of their data and thus maintaining their anonymity, the size of the group must not be too small. But on the other hand, if it is too big, the aggregated profile is just a vague representation of user's preferences resulting in a loss of recommendation quality. As far as we know, this problem has not been studied yet.

4 Policy-based Privacy

In addition to a variety of systems establishing **privacy-by-architecture** – by trying to create systems that minimize the amount of personal data being collected and processed for the well-functioning of the service -, several **privacy-by-policy** approaches have been developed during the past decade. These systems are aimed at implementing notice and choice principles of fair information practices (Toch, Wang, & Cranor, 2012), i.e. making users aware of the currently visited web site's or service’s privacy practices and giving them different options to choose from regarding the use and dissemination of information collected from and about them.

The Platform for Privacy Preferences Project (**P3P**) which has been published as a standard by the World Wide Web Consortium (W3C) in 2002 can be considered as a pioneer in the area of systems mitigating privacy risks by following a privacy-by-policy approach. P3P enables web sites to express their privacy practices in a standard format that can be retrieved automatically and interpreted by user agents. Client-side agents then not only inform users about the visited site’s privacy practices (in both machine- and human-readable formats) and warn them when its stated policies deviate from their previously specified preferences (Wang & Kobsa, 2007) but also allow automate decision-making based on these practices. Thus, transparency is ensured without users having to read the privacy policies on each and every site they visit (“W3C - Platform for Privacy Preferences (P3P) Project,” n.d.).

Users’ specification of privacy preferences is also the basis for the **privacy- enhancing user modeling framework** described in (Wang & Kobsa, 2007) In order to respect users’ individual privacy constraints, the user modeling system dynamically reconfigures itself by only selecting and using personalization methods that are in compliance with users’ current privacy settings. Therefore, the framework’s architectural structure relies on a product line architecture- based approach consisting of the following components:

- A User Modeling Server (UMS) including the Directory Component which acts as a repository of user models, thus storing users’ characteristics, behavior and individual privacy constraints, and a pool of User Modeling Components (UMCs) encapsulating user modeling methods for making inferences about users based on already existing data.
- A Selector that verifies for every UMC whether it may operate under the privacy constraints applying to the specific user and subsequently creates an architectural instance containing solely those UMCs that are permissible under the given privacy specification.

In this way, the user modeling framework allows personalization services adapted to and controlled by each user’s potentially different privacy constraints.

“**Scrutability**” represents another approach for ensuring users’ control over their personal data by enabling them to permit individual services to access selective so-called “personas”, i. e. partial user models. Hence, the user is given the opportunity to provide one persona to one application and a different persona to another. Therefore, the system consists of a set of components, each of which models a single aspect of user’s knowledge, beliefs and preferences. These pieces of information make up a scrutable representation of the user model allowing the user to examine the value of its components and inference mechanisms operating on them. Whenever a new application's persona is to be created, the application first proposes the components desired in the model among which the user then selects the ones he agrees with being part of this specific persona while he restricts the others. In this manner, the user exercises control over which aspects of his user model are available in each persona and thus, a certain application has access to (Kay, Kummerfeld, & Lauder, 2006).

The idea of leaving it to the user to decide which parties may access which parts of his personal data also forms the basic concept of the browser extension named "**RePriv**" as described in (Fredrikson & Livshits, 2011). Based on the user's browsing behavior, the browser infers information about his interests in order to form a user interest profile kept locally inside the browser. Each time a site requests information about the user, the latter will be asked by the browser for a permission to send his interests to the petitioning site. That way, the user stays in control of what information is released by the browser.

As a last example of an architectural approach establishing privacy-by-policy, "**Do Not Track**" shall be mentioned. It offers an HTTP header-based mechanism enabling users to express their preferences about third-party web tracking, including to opt-out of tracking some or all of the time. When processing a request including an opt-out header, a server must not perform third-party tracking, whereas in case of an opt-in header, tracking may be executed (Mayer, Narayanan, & Stamm, 2011).

Aforementioned preference signaling mechanisms implementing policy-based privacy basically rely on **explicit user consent** which means that a user not only accepts the choice he makes, but also is likely to understand both the scope and the consequences of approving or denying, respectively, the collection, usage and dissemination of his personal information in a specific context. Though, the problem posed by this architectural approach is that users' stated preference settings can easily be ignored by bad actors. Therefore, privacy-by-policy designs are primarily **targeted at fair players** as a means to aid them honoring a user's specified privacy preferences (Mayer et al., 2011).

5 Formal Policy Models

In this section, we will present some articles in which, their authors talk about privacy policy matters and privacy policy technologies. We will also see privacy policy formalization, meaning describing it using mathematical theories. Finally, we will also speak about some application fields of it, such as access control management and Privacy Preserving Internet Transfer (PPIT)/ Access control management is how a system makes a decision to grant or reject an access request from a subject, based on what the subject is authorized to access. It aims to assure, knowing the credentials and authorizations of subjects, information isn't accessed by unauthorized people.

PPIT on the other hand, is when a client needs to retrieve sensitive information held by another party (for example a server) such that:

- The former only gets the information for which it is authorized
- The latter doesn't learn about a client's authorizations

In (Fischer-Hübner and Ott 1998), the authors describe the concept of a formal security model that directly enforces basic legal privacy requirements, such as purpose binding or necessity of data processing. They informally described the privacy policy which is to be enforced by this model as follows:

"A user may only have access to personal data, if this access is necessary to perform his/her current task and only, if the user is authorized to perform this task. The user may only access data in a controlled manner by performing a (certified) transformation procedure for which the user’s current task is authorized. Besides, the purpose of his/her current task must correspond to the purposes for which the personal data were obtained or there has to be consent by the data subjects”

This task-based privacy model is characterized in their paper as a state machine model. This privacy model has been formally defined using the following elements: state variables, invariants (they define conditions for a system state to meet specific privacy principles), constraints (properties of sequences of states) and state transitions functions (model rules which describe all possible changes of state variables).

In another work (Ashley et al. 2002), the Platform for Enterprise Privacy Practices (E-P3P), a fine-grained privacy policy model is defined. A Chief Privacy Officer can use E-P3P to formalize the desired enterprise-internal handling of collected data. A particular data user is then allowed to use certain collected data for a given purpose if and only if the E-P3P authorization engine allows this request based on the applicable E-P3P policy.

By enforcing such formalized privacy practices, E-P3P enables enterprises to keep their promises and prevent accidental privacy violations. E-P3P privacy policies define the purpose for which collected data can be used, model the consent a data subject can give, and may impose obligations onto the enterprise. An E-P3P privacy policy defines what data users can perform what action for what purposes on which data categories. In their paper, the authors define the syntax of an E-P3P privacy policy and the semantic of processing simple and compound requests (requesting access to data belonging to multiple categories such as multiple column of a database).

Concerning access control management, (Fischer-Hübner and Ott 1998) specify how the privacy policy can be enforced according to the Generalized Framework for Access Control (GFAC) approach in Unix System V. GFAC is a framework for expressing and integrating multiple policy components to make it feasible to configure and to extend a system with security policies. It consists of an access enforcement facility (AEF) and an access decision facility (ADF). ADF enforces the system’s mandatory security policies and a metapolicy to decide whether processes’ requests satisfy those security policies. AEF uses the ADF-decisions to implement the operations of system call function. For the GFAC implementation, the access control system of the system kernel is divided into the AEF and ADF components and the ACI-module which administrates Access Control Information (ACI, e.g. security attributes).

The E-P3P (Ashley et al. 2002) authorization engine receives access requests, which consist of a single data user, a single data category, a single purpose, a single action, and context data as defined in the policy. It outputs a ruling (‘allow’, ‘deny’, ‘none’, or ‘error’), the rule identifier that mandates the ruling, and the list of obligations that are specified in the rule. In case the privacy policy mandates the default ruling, both the rule identifier and the obligation list are empty.

In conclusion, the authorization engines of access control management systems receive access requests containing certain information about it (see above). As a consequence, it needs that an enterprise deploys an authentication system that identifies these elements on top of these engine.

Also another issue that can be difficult in practice is determining the purpose of an access. Privacy-enabled application can provide the purpose and will respect the decision of the authorization engine, but if application and storage system interact using a pre-defined interface such as SQL, the system have to identify purposes based on other accessible attributes of the context of the request. So, the resolution of purpose may be coarser than desired.

In (Cristofaro et al. 2009), the authors introduce Privacy preserving Policy-based Information Transfer (PPIT). It is applicable to any scenario with a need to transfer information – and, more generally, perform some data-centric task – between parties who:

1. Are willing and/or obligated to transfer information in an accountable and policy-guided (authorized) manner.
2. Need to ensure privacy of server's data by preventing unauthorized access.
3. Need to ensure privacy of client's authorization(s) which grant it access to server's data.

So, to address this scenario, in this paper, the authors introduce and explore the concept of Privacy preserving Policy-base Information Transfer (PPIT) (with descriptions and mathematical definitions).They also construct three PPIT schemes based, respectively, on: RSA, Schnorr and IBE techniques. They have then tested theses schemes, investigate various performance improvements and demonstrate the practicality of proposed PPIT schemes.

Finally, in the authors demonstrate Mask, the first system addressing the problem of how to selectively share contents among a group of users based on access control policies expressed as conditions against the identity attributes of these users while at the same time assuring the privacy of these identity attributes from the content publisher.

Mask consists of three entities:

- Content Publisher: publish content on which selective access control is enforced
- Subscribers: consume published content based on their credentials
- Identity Provider: issue certified identity attributes

So, in conclusion, these two articles talked about a kind of a specific access control management in which we want the privacy policy to be preserved and the identity attributes of the content publisher or the client not to be revealed. But, to do so the use of cryptographic techniques are required and this can lower the system performance and also make it more complex.

Links with other EEXCESS Work Packages

1 Privacy and Architecture (WP1)

Privacy has many implications on the architecture of the EEXCESS system as a whole. These are more thoroughly discussed in this deliverable in section 3 of chapter Current Results of WP6.

2 Privacy and Federated Recommendation (WP3)

The role of WP3 is in providing recommendation in a federated way. Of course, privacy issues have many impacts on how this can be done. Among those are:

- How to provide the recommendation system sufficient information about each user and his context to obtain quality recommendations but at the same time respecting the user's privacy constraints ?
- How can we provide the recommender systems the necessary feedback in a privacy preserving way which will allow the system to learn from the previous recommendations and improve future ones ?
- How can sending profile and feed back be done by respecting users with different privacy constraints.

3 Privacy and Usage Mining (WP5)

The main role of the WP5 is to analyze user navigation traces with the aim to learn more about the user and then propose her the most accurate recommendations later.

Potentially the interaction with the WP5 in the project will increase progressively during the project:

- Assisting the user in choosing the most adapted disclosure policy. Indeed, according to (Knijnenburg & Kobsa, 2013), the willing of the user concerning what she decides to disclose and concrete choice have been observed to be sometimes contradictory. The usage mining could hint the user to choose a particular disclosure policy that corresponds to his analyzed behavior.
- Assisting the user in choosing a privacy policy based on other users having some similarities in their profiles and/or their behaviour. In other word, to recommend to the user a certain privacy policy. This alleviates the user task in selecting the appropriate privacy policy to adopt. In this case, thank to the mining, a set of disclosure policy template could be created and proposed to the user.

Advanced interaction with the privacy proxy. Indeed, in the first version of the proposed architecture, the usage mining module in the architecture is assumed to be trustable. In further versions of our work, we plan to consider the case where the usage mining is malicious. In this case, data exchange between the privacy proxy will change to take into account this case.

Current Results of WP6

1 Privacy Proxy Prototype

To better clarify the notions of transparency, control and feedback, and be able to discuss privacy related issues within the EEXCESS consortium, we implemented an initial prototype of a subset of the EEXCESS system focused on privacy. This prototype introduces the three previously discussed aspects:

- **transparency** through a detailed view of the data collected about the user
- **control** through a privacy settings page giving users control on the degree of disclosure of profile attributes
- **feedback** through a privacy sandbox, allowing users to see the impacts of their privacy settings within the same concrete context

1.1 Architecture of the Prototype

The prototype was implemented with a simplified version of the general EEXCESS architecture described in WP1. Figure 1 gives the architecture of what appears in the prototype. It is composed of 3 components:

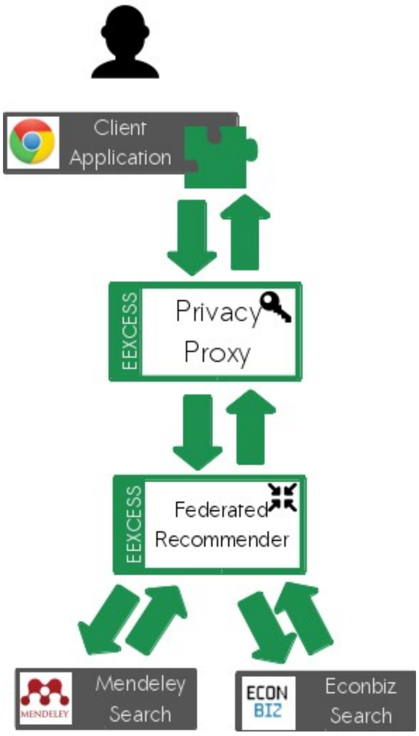

Figure 1: Simplified architecture

- A Chrome plugin to interact with the user
- An stub implementation of a privacy proxy
- A simplified version of a federated recommender

The plugin (further detailed hereafter) provides the interface between the user and the privacy proxy through which recommendations flow. The proxy collects and stores user profile and context data and relays the recommendation requests which it obtains from the plugin to the federated recommender. This latter component then forwards the requests down to each individual content provider. The endpoints integrated within this prototype are the search endpoints provided by Mendely and Econbiz.

The Mendely API was chosen as we wanted to test integration of the plugin through OAuth and Mendely was the only partner which, to date, provides such an API. To enforce the federated aspect of the recommender, we chose to include a second partner API. Econbiz was chosen both because they have similar content as those provided by Mendely on common domains (economics). As our focus is neither federation nor integration we also favored Econbiz for the simple technical reason that its search API is very similar to that of Mendely.

Even though many aspects were simplified, and the internals of the components are limited to the strict minimum we have an initial functional prototype of the EEXCESS system which we can use as a live technical testing environment. It is planned that the demo of this prototype be put online shortly. Having such a prototype also helps clarifing where the difficulties and challenges, in particular privacy challenges, are found within the EEXCESS system.

1.2 Integrating partner user profiles

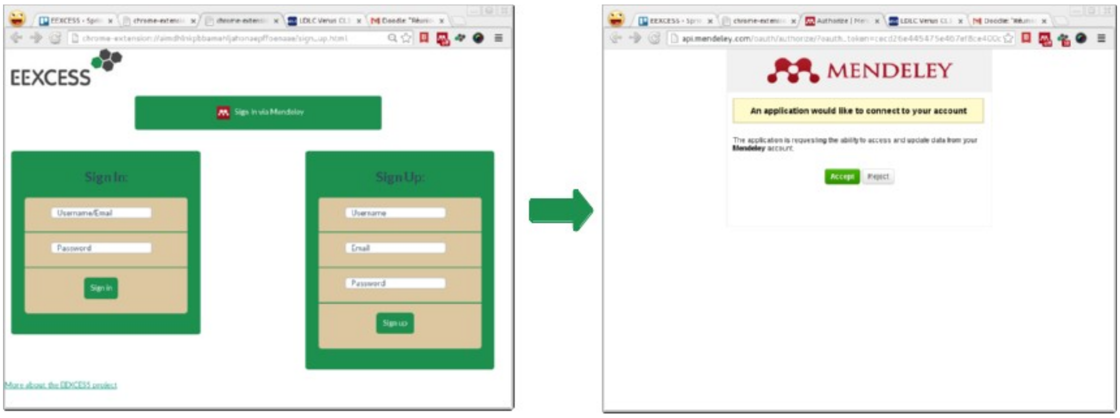

An application would like to connect to your account

The application is requesting the ability to access and update data from your **Mendeley** account.

Figure 2: OAuth connection through Mendeley API

One of the aspects of the EEXCESS project is collecting user profile information. Some partners already have profile information which could be useful to integrate. Within our prototype we have implemented connecting to Mendeley’s API through OAuth in order for the EEXCESS profile to be completed with such information (see Figure 2). At the same time, users connecting their profile with a given partner service also enacts their trust with the given service. This is a means by which user can express their trust of a partner. Having such connections between EEXCESS profiles and partner profiles will have to be taken into account for as they may have impacts on how to manage privacy. In particular, this could impose privacy requirements when transmitting profile data to different peers. For example, consider the case a user is willing that Mendeley , a peer he trusts, knowns he is at the origin of a recommendation request, but for whatever reason, he wishes to remain anonymous among the other peers.

1.3 User Interfaces

With the prototype, we have user interfaces which allow to illustrate and serve as the basis of discussion on how user can be empower and the impacts this may have within the system. Below are described the different user interfaces and how they respond to the objectives (transparency, control, and feedback) previously fixed. Within the prototype these are provided as three interface components.

Transparency

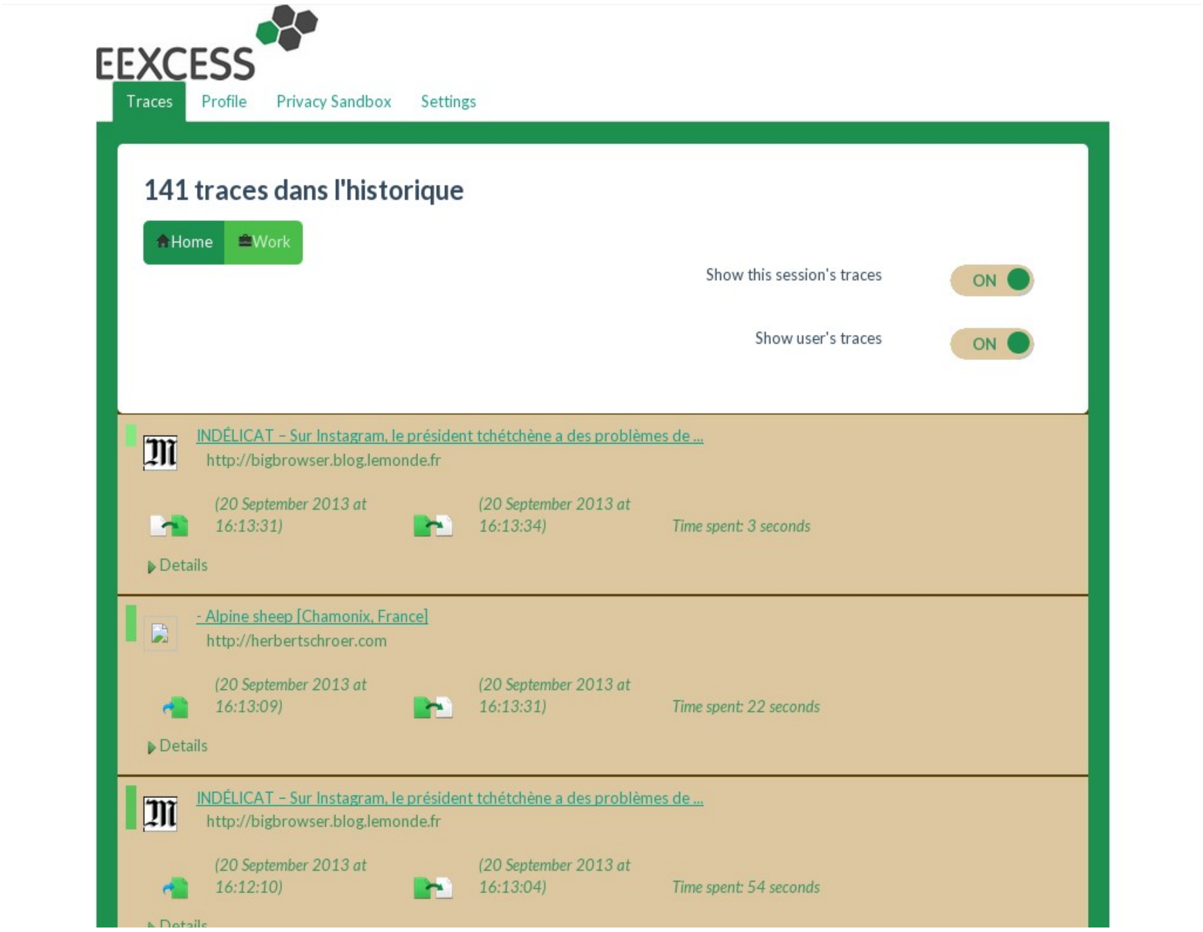

141 traces dans l'historique

Show this session's traces ON

Show user's traces ON

INDÉLICAT – Sur Instagram, le président tchétchène a des problèmes de ...
http://bigbrowser.blog.lemonde.fr
(20 September 2013 at 16:13:31) (20 September 2013 at 16:13:34) Time spent: 3 seconds
▸ Details

- Alpine sheep [Chamonix, France]
http://herbertschroer.com
(20 September 2013 at 16:13:09) (20 September 2013 at 16:13:31) Time spent: 22 seconds
▸ Details

INDÉLICAT – Sur Instagram, le président tchétchène a des problèmes de ...
http://bigbrowser.blog.lemonde.fr
(20 September 2013 at 16:12:10) (20 September 2013 at 16:13:04) Time spent: 54 seconds
▸ Details

Figure 3: Giving users transparency by showing collected data

Transparency is given by allowing the user to navigate through the information the system has about his browsing history and profile. This information is provided visually as complete as possible so that even non-technical users may easily understand what data is concerned (see Figure 3). For more technical users, the detailed internal representation (JSON in this case) is also made available to provide complete transparency.

Control

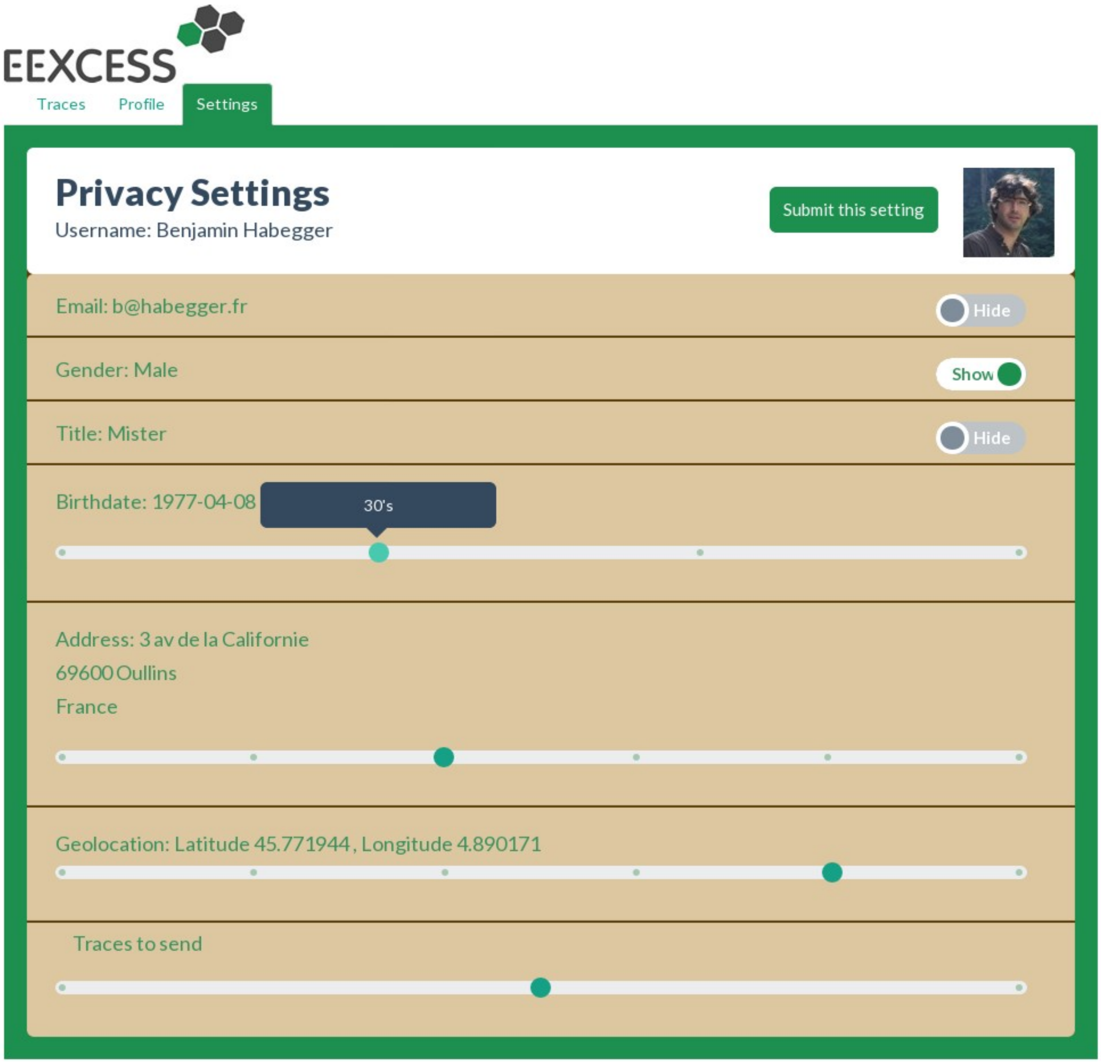

Privacy Settings

Username: Benjamin Habegger

Email: b@habegger.fr — Hide

Gender: Male — Show

Title: Mister — Hide

Birthdate: 1977-04-08 — 30's

Address: 3 av de la Californie
69600 Oullins
France

Geolocation: Latitude 45.771944, Longitude 4.890171

Traces to send

Figure 4: Giving users control over their privacy

Control is given to users by enabling them to select which information is disclosed. Furthermore, users are provided with different levels of disclosure, allowing them to disclose a particular piece of information (for example his age), but only partly. Figure 4 shows a screenshot of the privacy settings page. Currently, it mostly consists of demographic attributes. However, it show how control can be given by providing different levels of disclosure for each attribute. The “gender”, “title”, “email” attributes each have a toggle allow users to control their disclosure. The “birthday”, “address” and “geolocation” attributes extend this option by providing intermediate levels of disclosure. For example, in the screenshot below, the user is only disclosing the decade of his age rather than his exact age or even his exact birth date. The exact value and disclosed values are shown to the user.

Feedback

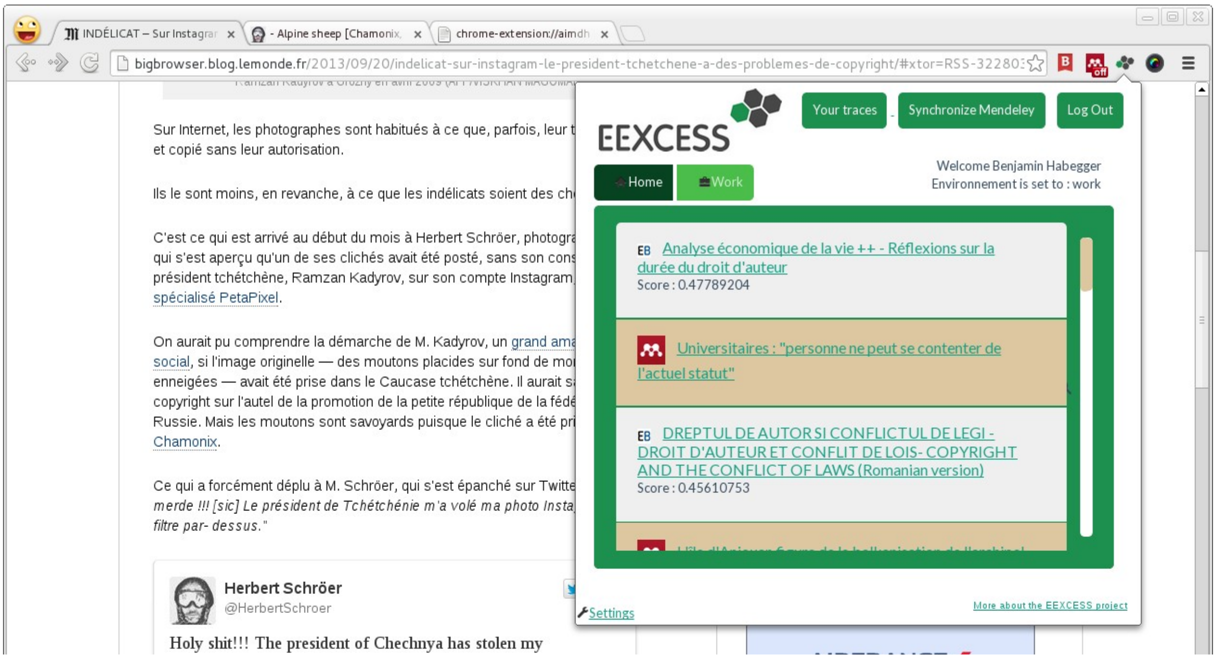
INDÉLICAT – Sur Instagrar
- Alpine sheep [Chamonix,
chrome-extension://aimdh
bigbrowser.blog.lemonde.fr/2013/09/20/indelicat-sur-instagram-le-president-tchetchene-a-des-problemes-de-copyright/#xtor=RSS-32280:
off
EEXCESS
Your traces
Synchronize Mendeley
Log Out
Home
Work
Settings
More about the EEXCESS project

Sur Internet, les photographes sont habitués à ce que, parfois, leur t
et copié sans leur autorisation.

Ils le sont moins, en revanche, à ce que les indélicats soient des ch

C'est ce qui est arrivé au début du mois à Herbert Schröer, photogra
qui s'est aperçu qu'un de ses clichés avait été posté, sans son cons
président tchétchène, Ramzan Kadyrov, sur son compte Instagram,
spécialisé PetaPixel.

On aurait pu comprendre la démarche de M. Kadyrov, un grand ama
social, si l'image originelle — des moutons placides sur fond de mo
enneigées — avait été prise dans le Caucase tchétchène. Il aurait s
copyright sur l'autel de la promotion de la petite république de la fédé
Russie. Mais les moutons sont savoyards puisque le cliché a été pri
Chamonix.

Ce qui a forcément déplu à M. Schröer, qui s'est épanché sur Twitte
*merde !!! [sic] Le président de Tchétchénie m'a volé ma photo Insta
filtre par- dessus."*

Herbert Schröer
@HerbertSchroer

Holy shit!!! The president of Chechnya has stolen my

Welcome Benjamin Habegger
Environnement is set to : work

EB Analyse économique de la vie ++ - Réflexions sur la durée du droit d'auteur
Score : 0.47789204

Universitaires : "personne ne peut se contenter de l'actuel statut"

EB DREPTUL DE AUTOR SI CONFLICTUL DE LEGI - DROIT D'AUTEUR ET CONFLIT DE LOIS- COPYRIGHT AND THE CONFLICT OF LAWS (Romanian version)
Score : 0.45610753

Figure 5: Example context and associated recommendations

The prototype also gives insight on how feedback could be given to users. In the screenshot below, users are offered a tool in which users can get a “feel” of the impacts of their privacy settings on the recommendations they would obtain. The tool is split in two parts as visible in figures 6 and 7. The part to the left, entitled “Privacy Settings”, offers the same toggles as the previous screenshot. The part to the right, entitled “Recommendations”, allows the user to select one of the past ten most recent browsing contexts (i.e. each green point in the slider representing a particular context). In the figures, the recommendations proposed are those he would have obtained on the last page he was browsing. The context just to its left would correspond to the next-to-last page he was on. Figure 5 shows an example browsing context page outside of the sandbox with its recommendations as they are obtained in a normal browsing context.

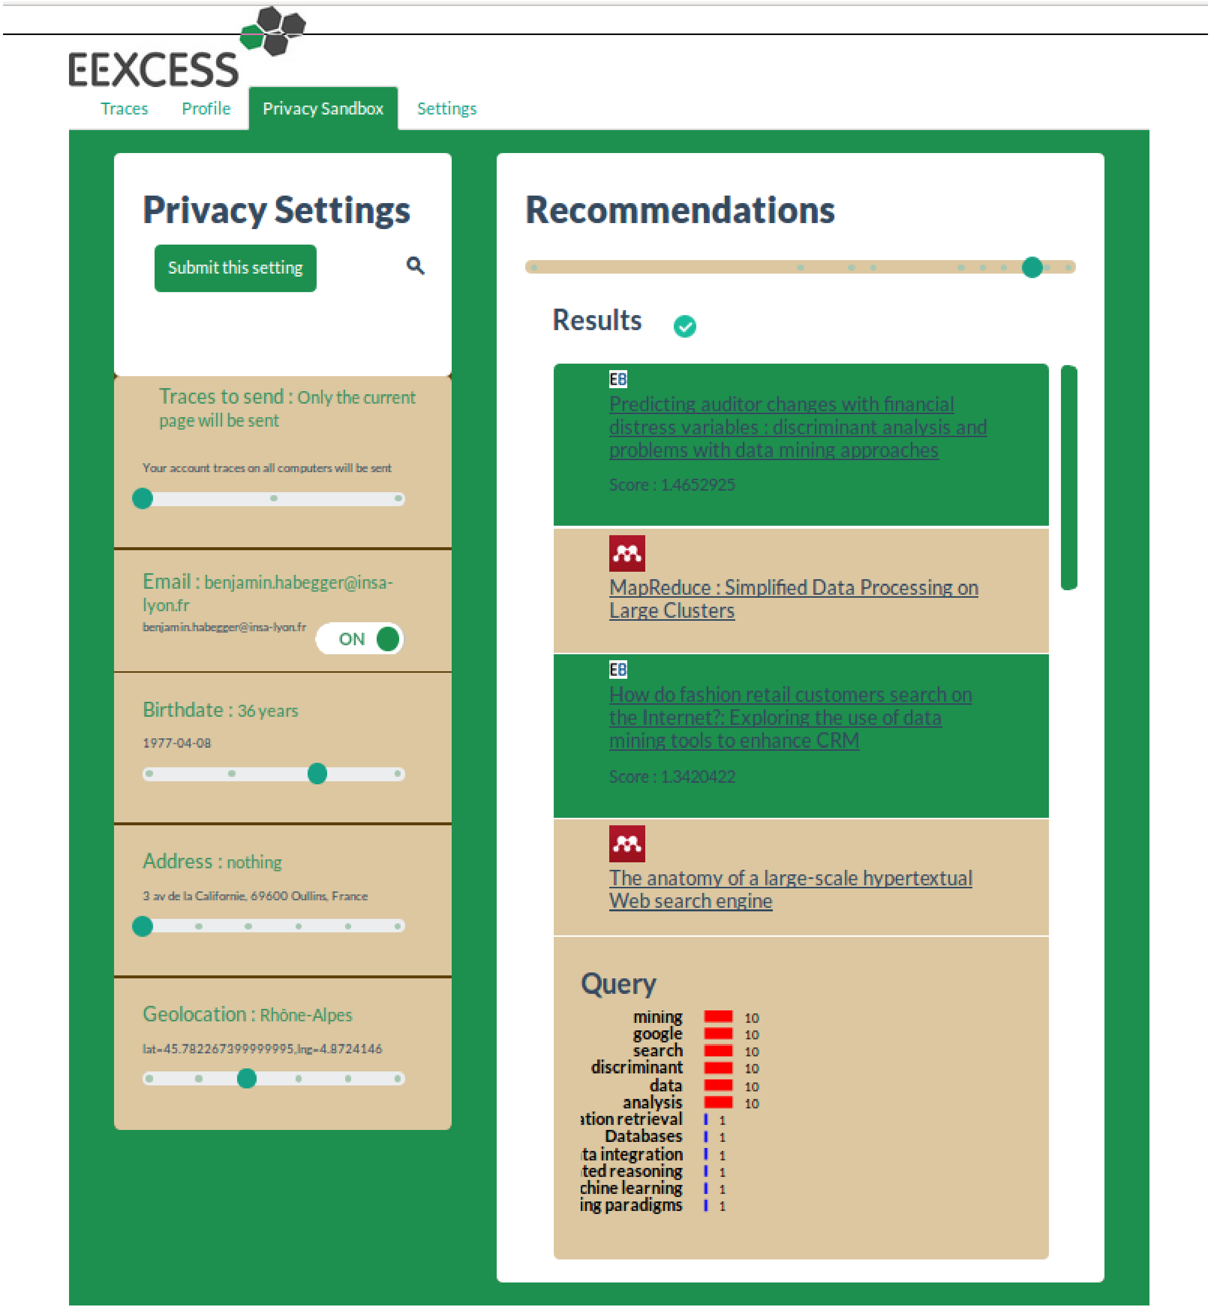

Privacy Settings

Traces to send : Only the current page will be sent

Your account traces on all computers will be sent

Email : benjamin.habegger@insa-lyon.fr

benjamin.habegger@insa-lyon.fr ON

Birthdate : 36 years

1977-04-08

Address : nothing

3 av de la Californie, 69600 Oullins, France

Geolocation : Rhône-Alpes

lat=45.782267399999995,lng=4.8724146

Recommendations

Results ✓

Predicting auditor changes with financial distress variables : discriminant analysis and problems with data mining approaches

Score : 1.4652925

MapReduce : Simplified Data Processing on Large Clusters

How do fashion retail customers search on the Internet?: Exploring the use of data mining tools to enhance CRM

Score : 1.3420422

The anatomy of a large-scale hypertextual Web search engine

Query

mining	10
google	10
search	10
discriminant	10
data	10
analysis	10
ation retrieval	1
Databases	1
ta integration	1
ted reasoning	1
chine learning	1
ing paradigms	1

Figure 6: Privacy sandbox (limited traces)

Within the sandbox, once a context has been selected, the user can play with the toggles of his privacy settings and directly see their impact on the recommendations that would have been made in a context he already knows.

Figures 6 and 7 show screenshots in which the user has selected the same context but different settings on the number of traces (i.e. the size of the browsing history) to be taken into account. The impacts of this choice are immediately visible on the query generated by the profile to query generator. Our prototype weighs terms in function of both their closeness in time to their current context and the time stayed on the pages they come from. Therefore, taking only one page gives the same weight to all terms which are those of the title of the only page of the context. In the second figure, terms from previously seen pages also come into play and influence the weights of the query terms.

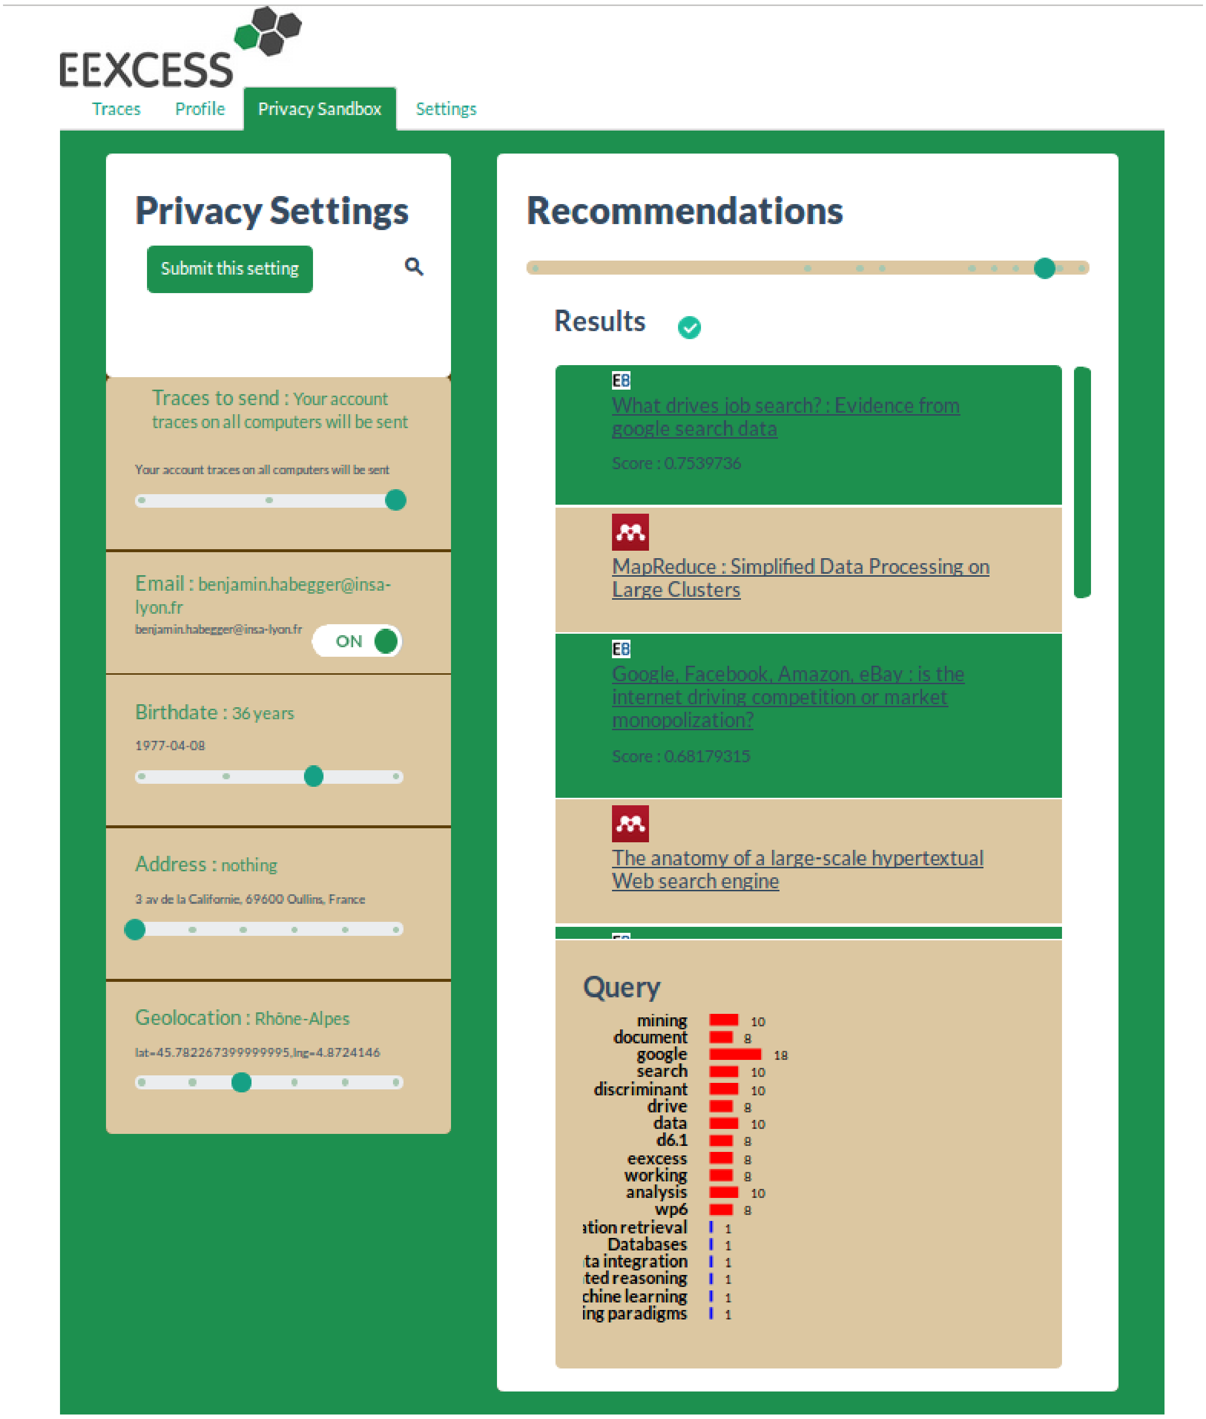

Privacy Settings

Submit this setting

Traces to send : Your account traces on all computers will be sent

Your account traces on all computers will be sent

Email : benjamin.habegger@insa-lyon.fr

benjamin.habegger@insa-lyon.fr ON

Birthdate : 36 years

1977-04-08

Address : nothing

3 av de la Californie, 69600 Oullins, France

Geolocation : Rhône-Alpes

lat=45.782267399999995,lng=4.8724146

Recommendations

Results ✓

What drives job search? : Evidence from google search data

Score : 0.7539736

MapReduce : Simplified Data Processing on Large Clusters

Google, Facebook, Amazon, eBay : is the internet driving competition or market monopolization?

Score : 0.68179315

The anatomy of a large-scale hypertextual Web search engine

Query

mining	10
document	8
google	18
search	10
discriminant	10
drive	8
data	10
d6.1	8
eexcess	8
working	8
analysis	10
wp6	8
ation retrieval	1
Databases	1
ta integration	1
ted reasoning	1
chine learning	1
ing paradigms	1

Figure 7: Privacy sandbox (extended traces)

2 Proxy REST API

This sections describes the API of the privacy proxy as it is implemented in our initial prototype. The details of the API can be found within on github[1]. This is an initial version of fitting the needs of the current prototype. Future version of the API will be augmented with the needs of the project following an agile approach.

2.1 Data Formats

Below are shortly described the data formats used by the current version of the privacy proxy prototype. The user profile format used within the prototype essentially uses a subset of the attributes of those defined in WP5. JSON has been favored within the prototype for technical simplicity as the components used within the prototype natively support it.

trace

A trace contains all the information about the user's current activity context. In the current state of our privacy proxy a trace has the following attributes (organized in groups):

[1] https://github.com/EEXCESS/privacy-proxy/blob/v0/API.md

user		
	user_id	An internal id for the user within the proxy
	environment	The user’s current environment tag: one of “home” or “work”
plugin		
	version	Version of the plugin
	uuid	Unique id of the plugin installation
temporal		
	begin	trace episode start time
	end	trace episode end time
events		
	begin	trace episode start trigger event
	end	trace episode end trigger event
document		
	url	URL of the document of the current context
	title	Title of the document of the current context
geolocation		
	country	Country the user has been geolocated in
	region	Region the user has been geolocated in
	district	District the user has been geolocated in
	place	Place the user has been geolocated in
	coord	Geolocation coordinates in the format “lat=<X>,long=<Y>”

Table 1: Trace attributes

Below is given an example JSON representation of a trace using these attributes.

```
{
   "user": {
       "user_id": "G7-Eh2EHRj6trSUa7rbbJw",
       "environnement": "home"
   },
   "plugin": {
       "version": "1.00",
       "uuid": "8437409A-B015-4F86-A3FA-05D911CC58F2"
   },
   "temporal": {
       "begin": "2013-09-11T17:00:35Z",
       "end": "2013-09-11T17:00:37Z"
   },
   "events": {
       "begin": "focus",
       "end": "unload"
   },
   "document": {
       "url": "http:fr.wikipedia.org/wiki/Xkcd",
       "title": "xkcd - Wikipédia"
   },
   "geolocation": {
       "country": "FR",
       "region": "Rhone-Alpes",
       "district": "Arrondissement de Lyon",
       "place": "Villeurbanne",
       "coord": "lat=45.771944,lng=4.890171"
   }
}
```

Profile

In the context of the prototype, a profile contains all the static information about a user.

username	Human readable username
email	User email
password	User encrypted password
title	User title among (“Miser”, “Miss”, “Mrs”)
lastname	User last name
firstname	User first name
gender	User gender
birthdate	User birthdate (in the ISO date format YYYY-MM-DD)
address	User address
street	
postalcode	
city	
region	
district	
country	
topics	List of topics
label	Label of the topic
env	Environment in which topic is applicable (either “Home” or “Work”)
source	The source of the topic (currently “eexcess” or “mendeley”)

Table 2: Profile attributes

Below is given an example profile using these attributes

```
{
   "username": "johndoe",
   "email": "john.doe@email.com",
   "password": "223252646cffg9dac77d03e16802078c",
   "privacy": {
       "email": "1",
       "gender": "0",
       "title": "0",
       "traces": "2",
       "geoloc": "0",
       "birthdate": "2",
       "address": "2",
   },
   "title": "Mister",
   "lastname": "Doe",
   "firstname": "John",
   "gender": "Male",
   "birthdate": "1990-04-06",
   "address": {
       "street": "13 rue du Chene",
       "postalcode": "69100",
       "city": "Villeurbanne",
       "country": "France",
       "region": "Rhone-Alpes",
       "district": "Arrondissement de Lyon"
   },
   topics: [
      {
          "label": "privacy preservation",
          "env": "all",
          "source": "eexcess"
      },
      {
          "label: artificial intelligence",
          "env: work",
          "source: mendeley"
      },
      {
          "label": "cookies",
          "env": "home",
          "source": "eexcess"
      }
   ]
}
```

2.2 Recommendation

api/v0/recommend/fetch

Given a current user context, this endpoint triggers a recommendation request and returns a scored list of recommendation. Internally the proxy, completes the user context with profile information, applies the user's privacy preserving settings and relays the request to the federated recommender.

In the current version, the obtained recommendations are not reranked and are returned as is.

api/v0/recommend/rewrite

This endpoint is the same as the previous one but requests only that the recommender returns the internally built recommendation query. This is used for feedback purposes in the privacy sandbox allowing users to visualize how their context has been mapped to a query.

2.3 Profiling

api/v0/privacy/trace

This endpoint stores a given user context (trace) in the user's activity log.

2.4 User

api/v0/user/profile

This endpoint allows to retrieve a user profile given an user's internal id.

api/v0/user/traces

This endpoint allows to retrieve the traces for a given context pattern. The pattern may include the target userId and/or the target pluginId completed by an optional target environment (among "work" or "home").

api/v0/user/data

This endpoint allows to store a user profile with the proxy.

api/v0/user/authenticate

This endpoint allows to authenticate a given user.

2.5 Connection

api/v0/connect/mendeley/init

This endpoint initiates an OAuth authentication through Mendeley and redirects the calling user to Mendeleys authorization page.

api/v0/connect/mendeley/finalize

This endpoint finalizes an OAuth authentication through Mendeley. It imports the user's Mendeley profile and merges it with and eventually existing EEXCESS user profile. Within the proxy, multiple raw profiles (one for EEXCESS and one for Mendely) are physically stored. User edits are saved to the EEXCESS raw profile. Any time a raw profile is updated a combined profile is constructed by giving preference to the raw EEXCESS profile for scalar attributes. Topic attributes are unified by creating a multi-labeled list of topics. Those coming from EEXCESS are marked EEXCESS and those coming from Mendeley are marked Mendeley. If the same topic comes from both it has both marks.

3 Impacts of Different Trust Scenarios

In the privacy literature, peers are often considered at three different levels of trust. "Trusted" peers are those which are entirely trusted. They will be respectful of the protocols and data they have been provided. Such data will only serve the purposes they were intended to (e.g. in the EEXCESS case, data transmitted for recommendation purposes will only be used for recommendation purposes). "Honest but curious" peers are those which will respect the protocols but as they are curious may peek at the data they are given access to and use the obtained information for other reasons than those that they were provided for. (e.g. using data originally given to obtain recommendations but also used to make a database of user interests used for other purposes than recommendation). “Malicious” peers are those which intentionally misuse the system (e.g. by submitting false data, hijacking the systems protocols) to obtain information they normally would not have access to (e.g. a fake partner joining as a content-provider with the intent of receiving recommendations requests and associated user profiles to build a spamming database).

3.1 All EEXCESS Components Trusted, Untrusted Providers

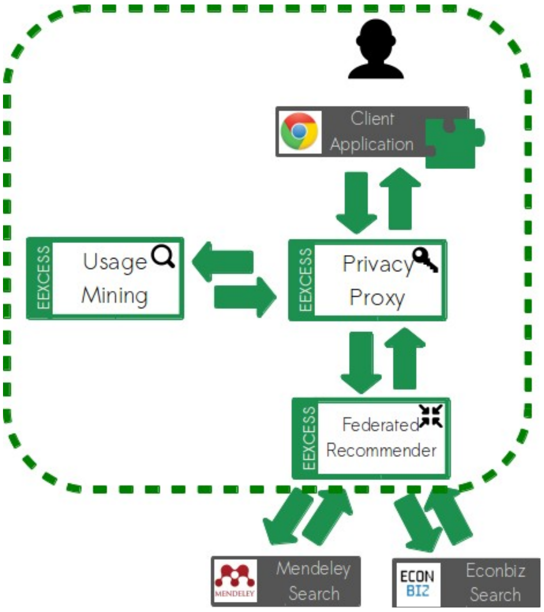

Figure 8: All trusted scenario

In this setting, all EEXCESS components (usage mining, privacy proxy, federated recommender) are trusted. The architectural split is only to distinguish different functional components of the EEXCESS system. In this case, the privacy proxy needs only to protect the EEXCESS users from untrusted (malicious or honest but curious) content partners. In figure 8, the circle of trust (ie. the components considered as trusted) are those found within the dashed green line.

This scenario is the currently active one within the EEXCESS consortium. The following scenarios will be studied as possible evolutions of the EEXCESS system. However, some anticipation of the following scenario already occurs by only giving the federated recommender anonymized profiles. How sufficient the anonymization is remains an open question to be settled before being able to consider that the federated recommender be a potentially untrusted peer (see section 4).

In this setting, protocols being straightforward, the main focus is on protecting user anonymity from statistical attacks. In the previous figure, this means protecting users from being re-identified by an adversary, which, in this case, would always be a content partner. Therefore, what is to be protected is the content flowing from the trusted parts of the EEXCESS system into the potentially untrusted peers. We are currently studying how to adapt the previously presented techniques on query obfuscation in the EEXCESS context.

Different factors of choice enter the game of choosing a privacy preserving approach. Among them are (1) the impact of privacy-preservation on the quality of the recommendations, (2) the level of privacy provided by the method, (3) the extra computational costs on the different peers implied by privacy preservation and (4) the extra networking costs implied by privacy preservation. Within the EEXCESS project in

WP5 a test dataset is planned to be created. This dataset will be useful to evaluate these different aspects.

3.2 Untrusted Federated Recommender

In this setting, the federated recommender is also left out of the circle of trust. This requires that the privacy settings users have specified also apply to this component. Otherwise said, the sum of information made available to the federator should not allow to de-anonymize users at the origin of the requests. This may have different impacts on the way information is anonymized. Further in this chapter (see section 4), we discuss the possible issues related to where obfuscation takes place in a search-based recommendation context. Depending on the outcomes of this research, stronger profile anonymization strategies may need to be considered. This may also have an impact on the works done in WP3 in that inputs to personalization may different depending on the found outcomes.

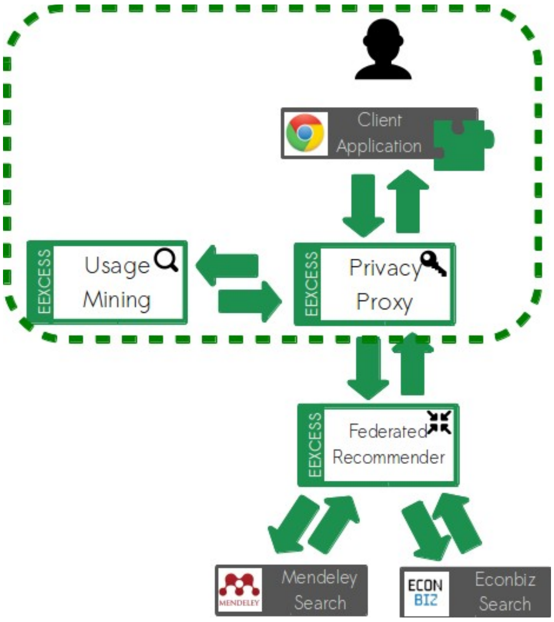

Figure 9: Untrusted federator

Furthermore, considering the federator as untrusted also has impacts on other types of recommendation it might implement. As previously stated, search-based recommendation is the current focus. However, many other forms of recommendation exist such as collaborative filtering or hybrid approaches (see state of the art within deliverable D3.1). It has been shown that anonymized versions of the data collected by recommender systems can be de-anonymized (Narayanan & Shmatikov, 2008; Peddinti & Saxena, 2010). Different recent techniques allow for collaborative filtering based on groups, requiring distributed recommendation and/or using cryptographic techniques.

3.3 Untrusted Usage-Mining

In this setting, it is the usage mining component which is left out the circle of trust. In such a case, the information collected about the users must not allow identifying them but at the same time still provide means to infer knowledge which can be used to enrich their user profiles. For example, this could be done is by collecting usage data in an anonymous way using or adapting techniques such as (Dwork, 2006; Nissim, 2008; Sweeney, 2002) and identify trends in a privacy preserving way. These trends could then be used to do query expansion with the goal to improve recommendation results. In a context where the federated recommender is trusted, this component could be the one using the trends directly and choosing those most adapted to the user.

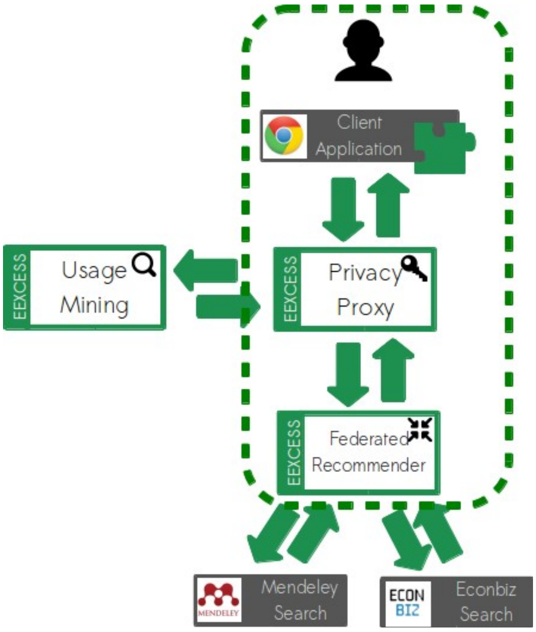

Figure 10: Untrusted usage-mining

3.4 Untrusted Federation and Untrusted Usage-Mining

In this setting, both the usage mining and federated recommender are considered untrusted. Developing solutions working with both constraints will require adapting those developed in the previous sections. As many techniques for privacy preserving data mining and privacy preserving recommendation rely on distributed computing and cryptography, these adaptations might be achieved using similar techniques.

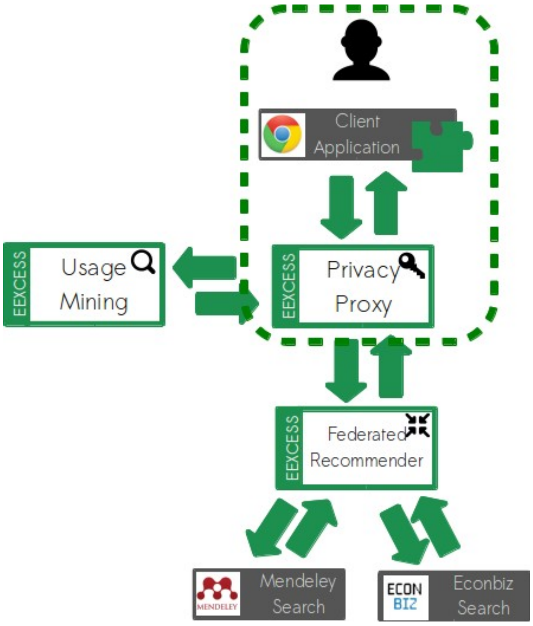

Figure 11: Untrusted federator and usage-mining

3.5 Untrusted Proxy

In this setting, all components are untrusted including the privacy proxy. In this case, the proxy can no longer be a physical component on its own but would require the privacy proxy to be a system distributed among different services or clients. In this case, privacy

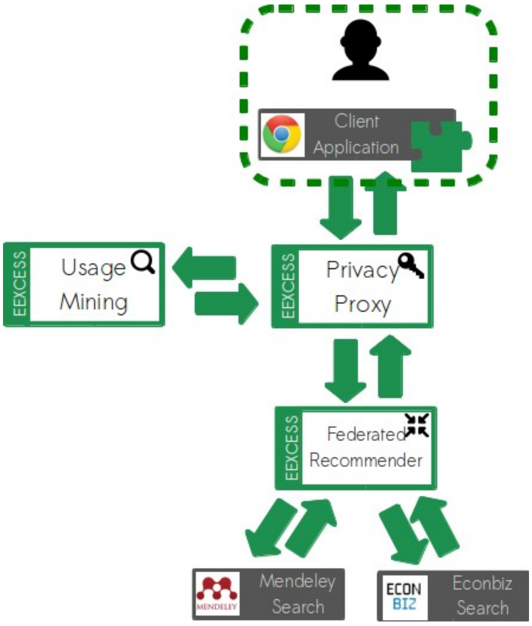

Figure 12: Untrusted proxy, usage-mining and federator

preserving profiles and recommendation groups would be calculated in a distributed manner using cryptographic techniques. Such techniques exist to calculate the reputation of different peers without revealing individual votes (Hasan, Brunie, & Bertino, 2012).

4 Privacy Strategy Model for Search-based Recommendation

Among the different techniques discussed previously, many rely on the general idea of “hiding in the crowd” (see section 3 of chapter “Related Work in Privacy“). This has lead to techniques in which queries or profiles are obfuscated or recommendations are requested per group of users rather than per user. Where obfuscation or grouping occurs in the architecture may have different impacts on:

- the effective privacy users can expect from the system
- the quality of the recommendations obtained
- which components should be trusted for privacy to be ensured

To better understand these impacts, we have modeled in a similar manner three different strategies in which obfuscation occurs at different levels in the EEXCESS architectural stack. These strategies follow the current focus of the EEXCESS project on search-based recommendation.

Our model is based on a fixed vocabulary of functions reused as much as possible in order to have comparable strategies. How these functions are effectively implemented is not the topic of this study. These function may rely on external parameters which are not shown here. For example, the SP function’s role is to secure a profile. What a secure profile looks like and how it is obtained (e.g. through grouping or obfuscation) are out of scope here. For example, it could be an attribute filtering approach removing quasi-identifiers and the profile would have a similar structure. SP might also rely on external parameters. Returning to our attribute filtering example, these parameters might be the attributes considered as quasi-identifier.

The functions in the model are the following:

- SP (secure profile): transforms a profile into a secured form
- MQ,MQ_i (map query): maps a profile into a recommendation query (in case this is done within the content providers, each provider may have its own mapping function distinguished by the indices).
- S_i (search) : runs a search given a query on a providers search index
- M (merge): merges a collection of search results into one
- RS (reorder search): reorders and filters the search results to adapt them to the real profile

4.1 Strategy S1 - Giving Providers Access to (Anonymized) User Profiles

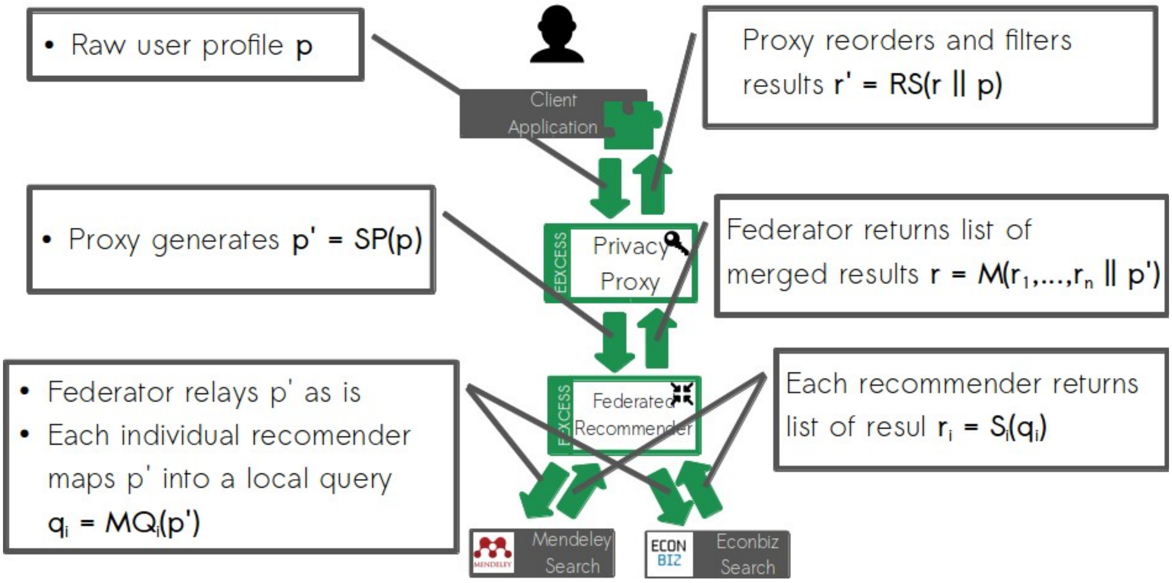
• Raw user profile p
Client Application
Proxy reorders and filters results r' = RS(r || p)
• Proxy generates p' = SP(p)
EEXCESS Privacy Proxy
Federator returns list of merged results r = M(r_1,...,r_n || p')
• Federator relays p' as is
• Each individual recomender maps p' into a local query q_i = MQ_i(p')
EEXCESS Federated Recommender
Each recommender returns list of resul r_i = S_i(q_i)
Mendeley Search
Econbiz Search

Figure 13: Strategy S1 (profile all the way down)

In this first strategy, a privacy-preserved profile is sent all the way to the content-providers. The user’s raw profile received by the privacy proxy (p) is secured in some privacy preserving way into a secured profile (p’) and sent to the federated recommender. The content-providers receive the secured user profile (p’) and internally map it into locally adapted search queries (q_i). The queries are run locally and results (r_i) returned to the federated recommender. The federated recommender merges the results together into one collection of results (r) returned to the proxy. The proxy then re-ranks the results and filters them according the real profile known only by him.

4.2 Strategy S2 - Giving Federated Recommender Access to (Anonymized) User Profiles and Providers (Anonymized) Recommendation Queries

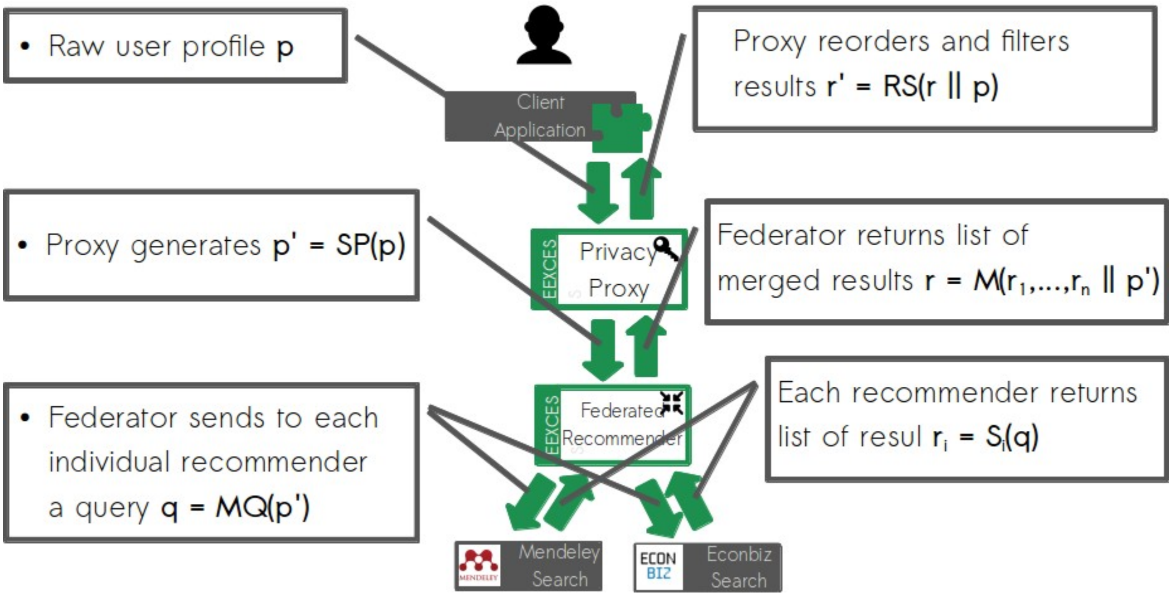
• Raw user profile p
Client Application
Proxy reorders and filters results r' = RS(r || p)
• Proxy generates p' = SP(p)
EEXCESS
Privacy Proxy
Federator returns list of merged results r = M(r₁,...,rₙ || p')
EEXCESS
Federated Recommender
• Federator sends to each individual recommender a query q = MQ(p')
Each recommender returns list of resul rᵢ = Sᵢ(q)
MENDELEY
Mendeley Search
ECON BIZ
Econbiz Search

Figure 14: Strategy S2 (half-way down)

In this second strategy, the privacy-preserved profile reaches only the federated recommender. The content-providers are only provided with a translation of the profile as a search query. The user's raw profile received by the privacy proxy (p) is secured in some privacy preserving way into a secured profile (p') and sent to the federated recommender. The federator map it into a search query (q) dispatched to all the content providers. Each provider runs the query and their results (r_i) are returned to the federated recommender. The federated recommender merges the results together into one collection of results (r) returned to the proxy. The proxy then re-ranks the results and filters them according the real profile known only by him.

4.3 Strategy S3 - Limiting User Profiles Access to Privacy-proxy

In this third strategy, the user profile doesn't even reach the federated recommender itself. The privacy-proxy (p) translates the user-profile into a query (q) representing the user's current information need. By applying some privacy preserving technique (such as query obfuscation), (q) is secured into another query (q') which is sent to the federator which relays it as is to the content-providers. Each provider runs the query and their results (r_i) are returned to the federated recommender. The federated recommender merges the results together into one collection of results (r) returned to the proxy. The proxy then re-ranks the results and filters them according the real profile known only by him into the final result set (r').

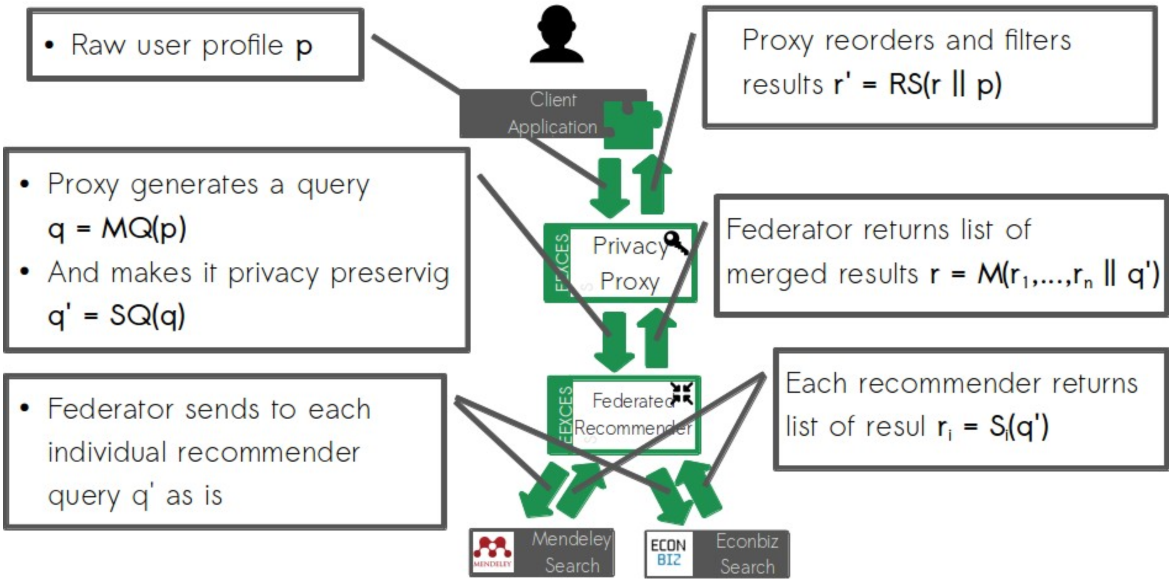
• Raw user profile p
Client Application
Proxy reorders and filters results r' = RS(r || p)
• Proxy generates a query q = MQ(p)
• And makes it privacy preservig q' = SQ(q)
EEXCESS
Privacy Proxy
Federator returns list of merged results r = M(r1,...,rn || q')
EEXCESS
Federated Recommender
Each recommender returns list of resul ri = Si(q')
• Federator sends to each individual recommender query q' as is
Mendeley Search
ECON BIZ
Econbiz Search

Figure 15: Strategy S3 (no profile)

4.4 Discussion

It is currently not settled how these strategies compare in terms of privacy preservation or recommendation quality. From the privacy perspective, one interesting aspect is what information flows into the untrusted peers. Based on the previous descriptions, table 1 gives a summary of how queries are generated in each strategy and of what is disclosed to the federator on one hand and to the content-providers on the other hand.

	Query generated	**Disclosed to federator**	**Disclosed to content-providers**
S1 (full down)	q_i = MQ_i(SP(p))	p’ = SP(p)	p’ = SP(p)
S2 (half down)	q = MQ(SP(p))	p’ = SP(p)	q = MQ(SP(p))
S3 (never down)	q = SQ(MQ(p))	q = SQ(MQ(p))	q = SQ(MQ(p))

Table 3: Impacts of the where privacy-preservation occurs on information disclosure

With this model, it is easier to compare how the strategies differ. By fixing the space P in which the profiles p can take their values and fixing the different functions SP, SQ and MQ, the next step is building estimates on how much information is disclosed by each strategy.

What remains to be determined is how to “measure” the information disclosed by a given privacy preserving function. We are currently investigating on how this could be answered both in theoretical and experimental contexts. From a theoretical point of view, one direction is to determine how maximum, average and minimum information disclosure could be bounded or approximated based on the properties of the profile and query space. These could confirmed or completed experimentally, by looking at how much information effectively disclosed at different points using statistical analysis.

Once these questions will be cleared out, we will have a framework which should allow us to compare which strategies are more appropriate given different profile securing and query securing solutions.

Conclusions

In this document we presented the constraints related to privacy within the EEXCESS projects and which challenges remained to be resolved in order to provide a policy mode preserving user privacy. The general challenges identified can be summarized as follows:

- Preserving privacy during recommendations
- Preserving privacy when sending user recommendation feedback
- Preserving privacy in usage mining
- Implying users in the process

The general policy goal is preventing the unwanted disclosure of information. To this effect it will be necessary find ways to measure in some way the level of information disclosure. A great challenge is that such measures should not only consider raw disclosures but also those which could be obtained through inference. Furthermore, users should have control in what should be disclosed or not and provided feedback on disclosure risks and data used.

In a first step, the first implementations of the privacy policy will focus on the following constraints:

- Trusted EEXCESS peers with anticipation of an untrusted federated recommender
- Protecting user anonymity by securing the search queries sent to content-providers
- Providing transparency, control and feedback to users

Research will be lead to identify how such constraints can be lifted, to better understand and measure the risks of information disclosure and how this relates to recommendation quality. The outcome of this research will orient future implementations.

References

Ardagna, C. A., Livraga, G., & Samarati, P. (2012). Protecting Privacy of User Information in Continuous Location-Based Services. In *2012 IEEE 15th International Conference on Computational Science and Engineering* (pp. 162–169). IEEE. doi:10.1109/ICCSE.2012.31

Brenner, H., & Nissim, K. (2010). Impossibility of Differentially Private Universally Optimal Mechanisms. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010* (pp. 71–80). Cryptography and Security; Computer Science and Game Theory, Las Vegas, Nevada, USA. Retrieved from http://arxiv.org/abs/1008.0256

Canny, J. (2002). Collaborative Filtering with Privacy, 45. Retrieved from http://dl.acm.org/citation.cfm?id=829514.830525

Castellà-Roca, J., Viejo, A., & Herrera-Joancomartí, J. (2009). Preserving user's privacy in web search engines. *Computer Communications*, *32*(13-14), 1541–1551. doi:10.1016/j.comcom.2009.05.009

Chang, C.-C., Thompson, B., Wang, H. (Wendy), & Yao, D. (2010). Towards publishing recommendation data with predictive anonymization. In *Proceedings of the 5th ACM Symposium on Information, Computer and Communications Security - ASIACCS '10* (p. 24). New York, New York, USA: ACM Press. doi:10.1145/1755688.1755693

Chaum, D. (1981). Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, *24*(2), 84–90. doi:10.1145/358549.358563

Chaum, D. L. (1981). Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, *24*(2), 84–90. doi:10.1145/358549.358563

Domingo-Ferrer, J., Solanas, A., & Castellà-Roca, J. (2009). h(k)-private information retrieval from privacy-uncooperative queryable databases. *Online Information Review*, *33*(4), 720–744. doi:10.1108/14684520910985693

Dwork, C. (2006). Differential Privacy, *4052*, 1–12. doi:10.1007/11787006_1

Dwork, C. (2008). Differential privacy: a survey of results, 1–19. Retrieved from http://dl.acm.org/citation.cfm?id=1791834.1791836

Felten, E. W., & Schneider, M. A. (2000). Timing attacks on Web privacy. In *Proceedings of the 7th ACM conference on Computer and communications security - CCS '00* (pp. 25–32). New York, New York, USA: ACM Press. doi:10.1145/352600.352606

Fox Tor. (n.d.).

Fredrikson, M., & Livshits, B. (2011). RePriv: Re-imagining Content Personalization and In-browser Privacy. *2011 IEEE Symposium on Security and Privacy*, 131–146. doi:10.1109/SP.2011.37

Haeberlen, A., & Pierce, B. C. (2011). Differential Privacy Under Fire. *Proceedings of the 20th USENIX conference on Security*, 33–33. Retrieved from http://www.usenix.org/event/sec11/tech/full_papers/Haeberlen.pdf

Hasan, O., Brunie, L., & Bertino, E. (2012). Preserving Privacy of Feedback Providers in Decentralized Reputation Systems. *Computers & Security*, *31*(7), 816–826. doi:10.1016/j.cose.2011.12.003

Howe, D., & Nissenbaum, H. (2009). {TrackMeNot}: Resisting Surveillance in Web Search, 417 – 436. Retrieved from http://www.citeulike.org/user/tnhh/article/7147121

Jackson, C., Bortz, A., Boneh, D., & Mitchell, J. C. (2006). Protecting browser state from web privacy attacks. In *Proceedings of the 15th international conference on World Wide Web - WWW '06* (p. 737). New York, New York, USA: ACM Press. doi:10.1145/1135777.1135884

Kay, J., Kummerfeld, R. J., & Lauder, P. (2006). Managing private user models and shared personas.

Knijnenburg, B. P., & Kobsa, A. (2013). Helping Users with Information Disclosure Decisions : Potential for Adaptation. In *Proceedings of the 2013 ACM international conference on Intelligent User Interfaces* (pp. 407–416). Santa Monica, CA USA: ACM Press.

Li, N., Qardaji, W., & Su, D. (2012). On sampling, anonymization, and differential privacy or, k -anonymization meets differential privacy. In *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security - ASIACCS '12* (p. 32). New York, New York, USA: ACM Press. doi:10.1145/2414456.2414474

Lindell, Y., & Waisbard, E. (2010). Private web search with malicious adversaries, 220–235. Retrieved from http://dl.acm.org/citation.cfm?id=1881151.1881164

Luo, Z., Chen, S., & Li, Y. (2013). A distributed anonymization scheme for privacy-preserving recommendation systems. In *2013 IEEE 4th International Conference on Software Engineering and Service Science* (pp. 491–494). IEEE. doi:10.1109/ICSESS.2013.6615356

Mayer, J., Narayanan, A., & Stamm, S. (2011). Do not track: A universal third-party web tracking opt out. *IETF Request for Comments*, 1–12.

Murugesan, M.-C. C. W. (2010). Privacy through deniable search. Retrieved from http://dl.acm.org/citation.cfm?id=2125739

Nandi, A., Aghasaryan, A., & Bouzid, M. (2011). P3: A privacy preserving personalization middleware for recommendation-based services. *Hot Topics in Privacy …*, 1–12. Retrieved from http://petsymposium.org/2011/papers/hotpets11-final6Nandi.pdf

Narayanan, A., & Shmatikov, V. (2008). Robust De-anonymization of Large Sparse Datasets. In *2008 IEEE Symposium on Security and Privacy (sp 2008)* (pp. 111–125). IEEE. doi:10.1109/SP.2008.33

Nissim, K. (2008). Private Data Analysis via Output Perturbation. In *Privacy-Preserving Data Mining: Models and Algorithms* (pp. 383–405).

Pang, H., Ding, X., & Xiao, X. (2010). Embellishing text search queries to protect user privacy. In *Proceedings of the VLDB Endowment* (Vol. 3, pp. 598–607). Retrieved from http://dl.acm.org/citation.cfm?id=1920841.1920918

Peddinti, S. T., & Saxena, N. (2010). On the privacy of web search based on query obfuscation: a case study of TrackMeNot, 19–37. Retrieved from http://dl.acm.org/citation.cfm?id=1881151.1881153

Peddinti, S. T., & Saxena, N. (2011). On the effectiveness of anonymizing networks for web search privacy. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security - ASIACCS '11* (p. 483). New York, New York, USA: ACM Press. doi:10.1145/1966913.1966984

Privoxy. (n.d.).

Ramakrishnan, N., Keller, B. J. B. J., Mirza, B. J. B. J., Grama, A. Y. A. Y., & Karypis, G. (2001). Privacy Risks in Recommender Systems. *IEEE Internet Computing*, *5*(6), 54–63. doi:10.1109/4236.968832

Saint-Jean, F., Johnson, A., Boneh, D., & Feigenbaum, J. (2007). Private web search. In *Proceedings of the 2007 ACM workshop on Privacy in electronic society - WPES '07* (p. 84). New York, New York, USA: ACM Press. doi:10.1145/1314333.1314351

Sánchez, D., Castellà-Roca, J., & Viejo, A. (2013). Knowledge-based scheme to create privacy-preserving but semantically-related queries for web search engines. *Information Sciences*, *218*, 17–30. doi:10.1016/j.ins.2012.06.025

Sankar, L., Rajagopalan, S. R., & Poor, H. V. (2013). Utility-Privacy Tradeoffs in Databases: An Information-Theoretic Approach. *IEEE Transactions on Information Forensics and Security*, *8*(6), 838–852. doi:10.1109/TIFS.2013.2253320

Sarathy, R., & Muralidhar, K. (2011). Evaluating Laplace Noise Addition to Satisfy Differential Privacy for Numeric Data. *Transactions on Data Privacy*, *4*(1), 1–17. Retrieved from http://dl.acm.org/citation.cfm?id=2019312.2019313

Shang, S., Hui, Y., Hui, P., Cuff, P., & Kulkarni, S. (2013). Privacy Preserving Recommendation System Based on Groups. Information Retrieval. Retrieved from http://arxiv.org/abs/1305.0540

Sweeney, L. (2002). k-ANONYMITY: A MODEL FOR PROTECTING PRIVACY. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, *10*(05), 557–570. doi:10.1142/S0218488502001648

Toch, E., Wang, Y., & Cranor, L. F. (2012). Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction*, *22*(1-2), 203–220. doi:10.1007/s11257-011-9110-z

Tor Project. (n.d.).

TorButton. (n.d.). Retrieved from http://www.torproject.org/torbutton/

W3C - Platform for Privacy Preferences (P3P) Project. (n.d.). Retrieved from http://www.w3.org/P3P/

Wang, Y., & Kobsa, A. (2007). Respecting users' individual privacy constraints in web personalization. In *User Modeling 2007* (pp. 157–166). Retrieved from http://link.springer.com/chapter/10.1007/978-3-540-73078-1_19

Ying-hua, L., Bing-ru, Y., Dan-yang, C., & Nan, M. (2011). State-of-the-art in distributed privacy preserving data mining. In *2011 IEEE 3rd International Conference on Communication Software and Networks* (pp. 545–549). IEEE. doi:10.1109/ICCSN.2011.6014329

Zheleva, E., Terzi, E., & Getoor, L. (2012). Privacy in Social Networks. *Synthesis Lectures on Data Mining and Knowledge Discovery*, *3*(1), 1–85. doi:10.2200/S00408ED1V01Y201203DMK004

Glossary

Terms used within the EEXCESS project.

Partner Acronyms

JR-DIG	JOANNEUM RESEARCH Forschungsgesellschaft mbH, AT
Uni Passau	University of Passau, GE
Know	Know-Center - Kompetenzzentrum für Wissenschaftsbasierte Anwendungen und Systeme Forschungs- und Entwicklungs Center GmbH, AT
INSA	Institut National des Sciences Appliquées (INSA) de Lyon, FR
ZBW	German National Library of Economics, GE
BITM	BitMedia, AT
KBL-AMBL	Kanton Basel Land, CH
CT	Collection Trust, UK
MEN	Mendeley Ltd., UK
WM	wissenmedia, GE

Abbreviations

API	Application Programming Interfaces
EC	European Commission
EEXCESS	Enhancing Europe’s eXchange in Cultural Educational and Scientific resource
GFAC	Generalized Framework for Access Control
HTML	HyperText Markup Language
HTTP	HyperText Transfer Protocol
IP	Internet Protocol
JSON	JavaScript Object Notation
PIR	Personal Information Retrieval
PPIT	Privacy Preserving Internet Transfer
REST	Representational state transfer
RSA	Rivest Shamir Adleman (authors of the RSA cryptosystem)
SQL	Structured Query Language
SVD	Singular Value Decomposition
TMN	Track Me Not
TOR	The Onion Router

Acknowledgement: The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 600601.