Optimal Tag Suppression for Privacy Protection in the Semantic Web[☆]

Javier Parra-Arnau*, David Rebollo-Monedero, Jordi Forné, Jose L. Muñoz, Oscar Esparza

Dept. of Telematics Engineering, Universitat Politècnica de Catalunya (UPC)
C. Jordi Girona 1-3, E-08034 Barcelona, Spain

Abstract

We present an architecture that protects user privacy in the semantic Web via tag suppression. Based on the principle of data minimization, tag suppression users may wish to tag some resources and refrain from tagging some others in order to hinder privacy attackers in their efforts to profile users' interests. Our strategy, however, poses an inherent trade-off between privacy and suppression. In this paper, we fundamentally investigate this trade-off in a mathematically systematic fashion and provide a detailed theoretical analysis. We measure the user privacy as the entropy of the user's tag distribution after the suppression of certain tags. According to this metric, we provide a close-form solution to the problem of optimal tag suppression.

Keywords: Information security, privacy, semantic Web

1. Introduction

The World Wide Web constitutes the largest repository of information in the world. Since its invention in the nineties, the form in which information is organized has evolved substantially. At the beginning, web content was classified in directories belonging to different areas of interest, manually maintained by experts. These directories provided users with accurate information, but as the Web grew they rapidly became unmanageable. Although they are still available, they have been progressively dominated by the current search engines based on web crawlers, which explore new or updated content in a methodic, automatic manner. However, even though search engines are able to index a large amount of web content, they may provide irrelevant results or fail when terms are not explicitly included in web pages. A query containing the keyword *accommodation*, for instance, would not retrieve web pages with terms such as *hotel* or *apartment* not including that keyword.

Recently, a new form of conceiving the Web, called the *semantic Web* [1], has emerged to address this problem. The semantic Web, envisioned by Tim Berners-Lee in 2001, is expected to provide the web content with a conceptual structure so that information can be interpreted by machines. The semantic Web requires

[☆]The material in this paper has been published in part in the proceedings of the 7th International Conference on Trust, Privacy & Security in Digital Business (TRUSTBUS), Bilbao, Spain, Sept. 2010.

*Corresponding author. Tel.: +34 93 401 7041.

Email addresses: javier.parra@entel.upc.edu (Javier Parra-Arnau), david.rebollo@entel.upc.edu (David Rebollo-Monedero), jforne@entel.upc.edu (Jordi Forné), jose.munoz@entel.upc.edu (Jose L. Muñoz), oesparza@entel.upc.edu (Oscar Esparza)

to explicitly associate meaning with resources on the Web. This process is normally referred to as *semantic tagging*, or simply tagging, and is supposed to play a key role for the semantic Web to become a reality. One of the benefits of associating concepts with web pages is the semantic interoperability in web applications. Furthermore, tagging allows applications to decrease the interaction with users, to obtain some form of semantic distance between web pages and to ultimately process web pages whose content is nowadays only understandable by humans.

Despite the many advantages the semantic Web is bringing to the Web community, the continuous tagging activity prompts serious privacy concerns. More specifically, tags submitted to a web server could be used to derive user's preferences [2] or expertise [3], and thus obtain precise user profiles containing sensitive information such as health, political affiliation, salary or religion. This could be the case of movie recommendation web sites such as *MovieLens* [4] or *Jinni* [5], and the social bookmarking web site *Delicious* [6], where user profiles are normally shown by some kind of histogram or tag cloud, as depicted in Fig. 1.

1.1. Anonymity vs. Tag Suppression

Philosophers, scholars and jurists have endeavored to conceptualize privacy since the *right-to-be-alone* definition given by Samuel Warren and Louis Brandeis in the late nineteenth century [7]. Although many admit that this task is virtually impossible [8], the privacy research literature [9] recognizes the distinction between *hard privacy* and *soft privacy*. Hard privacy, which may be regarded as *data minimization*, relies upon the assumption that users mistrust communicating entities and thus strive to reveal as little private information as possible. On the other hand, soft privacy assumes that users entrust their private data to an entity, which is thereafter responsible for the protection of their data. In the literature, numerous attempts to protect privacy have followed the traditional method of anonymous communications [10, 11, 12, 13], which fundamentally is built upon the assumptions of soft privacy. Unfortunately, anonymous communication systems are not completely effective [14, 15, 16, 17], they normally come at the cost of infrastructure, and assume that users are willing to trust other parties. However, even in those cases where we could trust an entity completely, that entity could eventually be legally enforced to reveal the information they have access to [18]. A discussion on the shortcomings of the main approaches regarding hard and soft privacy is provided in Sec. 2.

In this paper, we propose a strategy called *tag suppression*, which is based on the data minimization principle. Despite the fact that the proposed strategy and the state-of-the-art anonymous communication systems rely upon different assumptions, we would like to emphasize that both alternatives are not mutually exclusive and, more importantly, that users could benefit from the synergy of our approach and other systems providing soft privacy. As a matter of fact, there are examples in the literature in which techniques providing hard privacy may complement anonymous communication systems perfectly. One example of this could be the use of dummy messages in combination with the traditional mix networks proposed in [10]. The study of the impact of a possible application of tag suppression in other privacy protecting systems is out of the scope of the present work.

1.2. Contribution and Plan of this Paper

In this paper, we present an architecture that protects user privacy in the semantic Web via tag suppression. More specifically, users may want to tag some resources and refrain from tagging some others when their

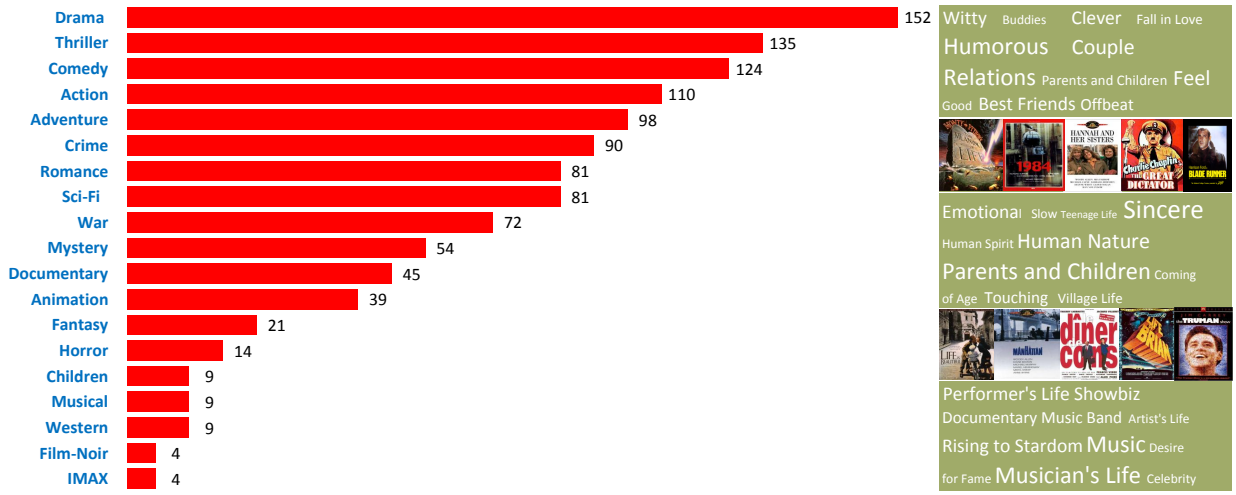

(a) (b)

Figure 1: A histogram (a) and a tag cloud (b) displaying user profiles in *MovieLens* and *Jinni*, respectively.

privacy is being compromised. The proposed architecture helps users decide which tags should be suppressed in order to hinder privacy attackers in their efforts to profile users' interests. Consequently, this approach guarantees user privacy to a certain extent without having to trust an external entity, but at the cost of some processing overhead and, more importantly, the semantic loss incurred by suppressing tags.

The theoretical analysis of this inherent trade-off between privacy and suppression is precisely the main object of this paper. More specifically, we present a mathematical formulation of optimal tag suppression in the semantic Web. We propose an information-theoretic criterion to measure the user privacy, namely, the entropy of the user's tag distribution after the suppression of certain tags, and justify it by the rationale behind entropy maximization methods. Accordingly, we formulate and solve an optimization problem modeling the privacy-suppression trade-off. In our theoretical analysis on tag suppression, we shall encounter several riveting mathematical analogies with our previous work on query forgery [19].

Section 2 explores the basics of the semantic Web and reviews some relevant approaches related to privacy. Section 3 describes a privacy-protecting architecture based on the suppression of tags. In addition, this section presents our privacy measure and the assumed adversarial model. Section 4 introduces a formulation of the trade-off between privacy and suppression in the semantic Web. Section 5 presents a detailed theoretical analysis of the optimization problem characterizing the privacy-suppression trade-off. Section 6 shows a simple but insightful example that illustrates the formulation and theoretical analysis argued in the previous sections. Conclusions are drawn in Sec. 7.

2. State of the Art

This section describes the fundamentals of the semantic Web and includes some relevant contributions to privacy within this context.

We mentioned in Sec. 1 that the semantic Web requires to explicitly associate meaning with resources on the Web. In order to achieve this meaningful structure, the conceptual description of resources must be described formally. For this purpose, the World Wide Web Consortium (W3C) proposes to use the resource

description format (RDF), which is a general-purpose language for representing information on the Web. In RDF, the meaning is encoded by a triple consisting of a *subject*, a *predicate* and an *object*. According to this format, a resource on a web page (subject) is associated with a property (predicate), to which a value (object) is assigned. For instance, in the statement "1984 was written by George Orwell", "1984" would be the subject, "was written by" the predicate, and "George Orwell" the object.

Although RDF provides the technology to describe meaning, the semantic Web requires also that concepts and terms share a common definition. Ontologies, which are defined in [20] as "a formal, explicit specification of a shared conceptualization", arise with this aim. In the semantic web context, an ontology is a set of statements where terminology is defined using a specific language. Several languages such as RDF schemas (RDF-S) [21] or ontology web language (OWL) [22] are used to express ontologies.

A number of approaches have been suggested to preserve user privacy in the semantic Web, most of them focused on privacy policies. In the traditional Web, the majority of web sites interact with users to provide them with privacy policies, and allowing them to find out how their private information will be managed. Unfortunately, users do not frequently understand [23] or even read [24] privacy policies. The platform for privacy preferences (P3P) is created to deal with this situation and provides a framework with informed online interactions. More specifically, when a web site supports the P3P, it establishes a set of policies to define how user's private information will be used. Users, in turn, set their own privacy policies to determine what kind of personal information they are willing to disclose to the web sites they browse. Accordingly, when a user browses a web site, P3P compares both the web site's and the user's privacy policies. If they do not match, P3P informs the user about this situation and consequently they decide how to proceed. In the semantic Web, this process is intended to be carried out by autonomous agents. In this context, several policy languages to define privacy and security requirements have been proposed. In [25], the authors suggest a new semantic policy language based on RDF-S to express access control requirements over concepts defined in ontologies. In [26], privacy and authentication policies are incorporated into the descriptions of an ontology called OWL *for services* (OWL-S). Furthermore, the authors implement algorithms for the requester to verify the provider's adherence to policies.

In the context of private information retrieval (PIR), users send general-purpose queries to an information service provider. An example would be a user sending the query:"What was George Orwell's real name?". In this scenario, query forgery, which consists in accompanying genuine with false queries, appears as an approach to guarantee user privacy to a certain extent at the cost of traffic and processing overhead. Precisely, in [19] we investigate the trade-off between privacy and the additional traffic overhead in a mathematically systematic fashion. Building on the simple principle of query forgery, several PIR protocols, mainly heuristic, have been proposed and implemented. In [27, 28], a solution is presented, aimed to preserve the privacy of a group of users sharing an access point to the Web while surfing the Internet. The authors propose the generation of fake transactions, i.e., accesses to a web page to hinder eavesdroppers in their efforts to profile the group. Privacy is measured as the similarity between the actual profile of a group of users and that observed by privacy attackers [27]. Specifically, the authors use the cosine measure, as frequently used in information retrieval [29], to capture the similarity between the group genuine profile and the group apparent profile. Based on this model, some experiments are conducted to study the impact of the construction of user profiles on the performance [30]. In line with this, recent surveys with a greater focus on anonymous Internet search include [31, 32]. Further, some simple, heuristic implementations in the form of add-ons for popular browsers have recently started to appear [33, 34].

Yet another strategy for anonymous tagging could be built upon the principle of user collaboration, not unlike the protocols for k-anonymous location-based services (LBSs) [35] and for other forms of anonymity through collaboration [36, 37]. Additionally, we could conceive the adoption of trusted third parties (TTPs), or even digital credentials [38, 39, 40], in order to enable anonymous and pseudonymous tagging. More specifically, we could make use of the anonymous communication systems already introduced in Sec. 1.1. In this sense, a number of approaches have been proposed during the last two decades. Most of them are based on Chaum's mix networks [10], which aimed to address *traffic analysis*, i.e., the process of intercepting and examining messages in order to infer any information from patterns in communication. Essentially, a TTP called mix collects messages from a number of senders and forwards them to their intended receivers, possibly other mixes, rearranging them with the express purpose of hiding the correspondence between inputs and outputs. Messages sent to mixes are encrypted using public-key cryptography, in a layered fashion when several mixes are involved. There are several anonymous communication proposals based on the idea of mix networks. They can be roughly classified into high-latency and low-latency systems. The former introduce significant delay to attain a high degree of anonymity against traffic analysis [41, 42]. Naturally, their main drawback is that they are hardly applicable to real-time interactive tasks such as tagging, web browsing or online chat.

In an attempt to address this limitation, low-latency systems were proposed. To attain a higher degree of anonymity, rather than increasing latency disproportionately, these systems simply benefit from the networking of a combination of several mixes frequently accessed by a significantly large population of users. Quoting [43], "All mix systems require that messages to be anonymized should be relayed through a sequence of trusted intermediate nodes". Some of these systems are based on peer-to-peer communications [44, 45], under the assumption of a very large interconnected population of trusted users who know how to reach each other, and who also act as mixes. The most popular approaches are *onion routing* [11, 12] and its second-generation version TOR [13]. Onion routing uses a single data structure encrypted in a layered fashion to build an anonymous circuit. Alternatively, TOR uses an incremental path-building design, where a client who wishes to communicate with a server negotiates session keys with each successive hop in the circuit.

Although these anonymous communication systems have contributed to protect user privacy to a certain extent, none of them have proved to be completely effective. More specifically, it is widely acknowledged that systems based on mix networks such as the ones described above are vulnerable to *timing attacks.* Namely, an attacker controlling the first and last mix on a path may link a sender to a receiver based on the observation of the period of time in which they exchange messages. This and other attacks have been studied in depth in [14, 15, 16, 17].

Lastly, we would like to remark that, despite the simplicity of query forgery, an analogous *tag forgery* would clearly not be convenient for the semantic Web, which is the motivating application of our work. Submitting a tag implies the construction of conceptual relations, a much more complex process than just sending a simple query to a service provider. Therefore, users might not be willing to manually tag web content they are not interested in. However, even though automatic mechanisms for autonomous tag forgery might be considered, they might lead to qualitative or quantitative semantic distortion.

3. An Architecture for Privacy Protection in the Semantic Web

This section presents a privacy-protecting architecture in the semantic Web via tag suppression. More specifically, Sec. 3.1 provides some insight into the construction of user profiles. The adversarial model is described in Sec. 3.2 and our privacy metric is presented and justified in Sec. 3.3. Next, we examine the proposed architecture from a global point of view and go into the details of its internal functional blocks. The specification of one of these blocks will be given in Sec. 4.

3.1. User Profile Construction

Our architecture contemplates that the profile of a user is directly obtained from specific modules integrated into the user's system. Before giving any details on the construction of user profiles, we will first explore how this information could be represented.

Section 1 already mentioned that some web sites commonly use some kind of histogram to show a user profile, as in the case of *MovieLens*, or tag clouds, as in *Jinni* or *Delicious*. Bearing in mind these examples, we propose a first-approximation, mathematically-tractable model of user profile as a probability mass function (PMF), that is, a histogram of relative frequencies in which each frequency is the proportion of times a particular tag has been submitted by the user with regard to the total number of tags this user has generated.

According to this model, we suggest two alternatives to represent a user profile. Our first proposal entails certain information loss, as it uses categories into which tags are mapped. On the one hand, this could be difficult to carry out, as tags would have to be classified into categories, but on the other, the main advantage would undoubtedly be the simplification of user profiles. Our second alternative represents a user profile by means of tags, which do not necessarily have to coincide with the semantic tags in the RDF format discussed in Sec. 2. Hence, this approach could provide a much more accurate description of user profiles, although at the expense of a higher complexity.

Once we have described our proposals to model a user profile, we will now focus on how to extract this information from a user tag activity. For simplicity, we shall hereafter assume that user profiles are modeled by categories, although all considerations also apply to tag-based profiles. A possible approach would first classify the submitted tags into categories, would keep a histogram of all of them, and finally would calculate the relative frequency of each category. In principle, this categorization process could be accomplished by exploring contextual information from the web page the user is tagging. Specifically, this could be done by using the vector space model [46], as normally done in information retrieval, to represent web pages as tuples containing their most representative terms. Namely, the term frequency-inverse document frequency (TF-IDF) would be applied to calculate the weights of each term appearing in a web page. Afterwards, the most weighted terms could be combined with the semantic tag submitted by the user in order to obtain the category associated with that tag. As an example, consider a user browsing a web page and submitting the tag "A conference was held in Copenhagen". Instead of using this tag to update the user profile, the system would first extract contextual information from the web page as described above, and later, the category "climate change"in the user profile would be updated.

3.2. Adversarial Model

Our proposal is built on the simple principle of tag suppression. More specifically, a user may wish to tag some resources and refrain from tagging some others to enable the resulting user profile, as observed from the outside, approach the uniform profile. We shall refer to this resulting user profile as the user's *apparent* profile.

Bearing in mind the above considerations, we assume a rudimentary adversarial model in which users submitting tags are observed by a passive attacker who is able to ascertain which tags are associated to which resources. Namely, this could be the case of the web server storing the tags submitted by users, or any privacy attacker able to crawl through this information. In addition, we may contemplate the definition of the profile of a user tagging across several of these web servers. In this case, we may also suppose that an attacker has the ability to link several profiles across different servers. However, for the sake of simplicity, we consider a user interacting with a single server, although the architecture proposed in Sec. 3.4 could be easily extended to the more general case in which a user tagging activity spans a number of servers.

Last but not least, we also assume that the privacy attacker is unable to discern whether a particular user is adhered to the proposed privacy strategy, and therefore can not estimate their tag suppression rate.

3.3. Privacy Metric

We use an information-theoretic quantity to reflect the intuition that an attacker will be able to compromise user privacy as long as the user's apparent profile diverges from the uniform profile. Specifically, we measure privacy as the entropy of the user's apparent distribution, which may be interpreted as a measure of uncertainty of that probability distribution, and also regarded as a special case of Kullback-Leibler (KL) divergence [47]. More specifically, the KL divergence between two probability distributions s and p, that is, $\mathrm{D}(s \,\|\, p)$, is essentially equivalent to the entropy of s in the special case when p becomes the uniform distribution. According to this, we may establish a connection between our privacy criterion and the privacy metric proposed in [19], in which s represents the user's apparent distribution and p the population distribution. In other words, the criterion in [19] is a slight generalization of the criterion proposed in this paper. However, our privacy measure significantly simplifies the architecture of a practical implementation since the population distribution, which in principle could be difficult to estimate, is not required to evaluate the user privacy. In any case, the mathematical analysis in Sec. 5 can be readily extended to divergence.

Another interpretation of entropy stems from the observation that a privacy attacker will have actually gained some information about a user whenever their interests are significantly concentrated on a subset of categories. In other words, a user without any apparent interest in any category hides their preferences from an attacker.

Having defined our measure of privacy, Sec. 4 later considers the problem of maximizing the entropy of the user's apparent distribution and thus the user privacy for a given tag suppression rate. Precisely, our privacy criterion is justified by the rationale behind entropy maximization methods [48, 49]. Namely, some of the arguments in favor of these methods are related to the highest number of permutations with repeated elements associated with an empirical distribution [48], or more generally, the method of types and large deviation theory [47, §11].

In addition, we would like to emphasize that, although our privacy criterion is based on a fundamental quantity in information theory, the convergence of these two fields is by no means new. As a matter of fact,

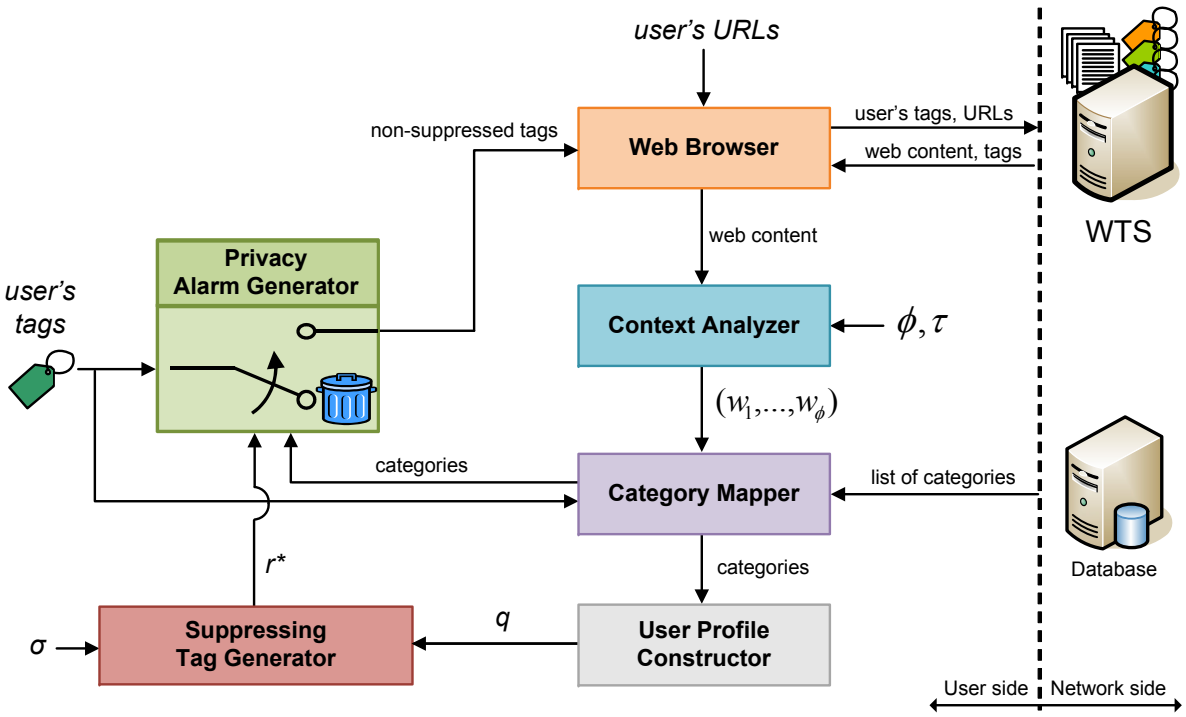

Figure 2: Internal components of the user-side architecture.

Shannon's work in the fifties introduced the concept of *equivocation* as the conditional entropy of a private message given an observed cryptogram [50], later used in the formulation of the problem of the wiretap channel [51, 52] as a measure of confidentiality. In addition, recent studies [53, 54] reassert the suitability and applicability of the concept of entropy as a measure of privacy. Specifically, the authors propose to measure the degree of anonymity observable by an attacker as the entropy of the probability distribution of possible senders of a given message.

3.4. Architecture

We now describe an architecture that helps users decide which tags should be suppressed in order to hinder privacy attackers in their efforts to construct a user profile too different from the uniform profile. Specifically, this section examines the internal components of the proposed architecture and goes into the details of a practical implementation.

On a side note, we shall assume that only a small number of users adopt the tag suppression strategy, in contrast to the large population of Internet users. In that case, the global detriment in semantic functionality is small.

From a global perspective, the main component of this architecture is the web and tag server (WTS), a single entity in which web pages and their semantic tags are stored. Users browsing the Web would retrieve those data from the WTSs. The web browser would represent this information so that it could be understood by users. Afterwards, users would generate their own semantic tags and would submit them to the WTSs.

Figure 2 depicts the proposed architecture from the user's point view. As it can be seen in this figure, the user-side architecture is composed by a number of modules, each of them performing a specific task. Next, we provide a functional description of all of their components.

Web Browser. This module is essentially responsible for the communication with the WTS. Specifically, it downloads both the web content and the semantic tags that the user specifies by means of a URL. Afterwards, the web content is delivered to the *context analyzer*, which extracts contextual information from the web page. Last but not least, the web browser is also in charge of submitting the tags proposed by the user to the WTS.

Context Analyzer. This module is aimed to process the web content that is requested by the user. Particularly, it performs this task by using the vector space model and the TF-IDF weights commented on in Sec. 3.1. As a result, a tuple of weighted terms is internally generated for each web page. Later, the context analyzer takes a number of the most weighted terms of each tuple, and sends them to the *category mapper* module. The selection of these terms could be done according to these two possible alternatives: a user could choose either a fixed number of terms ϕ, or those terms with weights above a threshold τ. This selection poses a compromise between accuracy and computational overhead, regardless the alternative chosen. The higher the resulting number of terms, the higher the accuracy in the categorization of the semantic tag, but the higher the computational processing performed by the category mapper.

Category Mapper. This component maps the tags submitted by the user into a set of fixed, predefined categories. This set of categories could be obtained by querying databases with this kind of information, or directly from an application such as *Google Insight*. The categorization process performed by this module uses both the semantic tag and the contextual information given by the context analyzer. The resulting categories are delivered to the modules *user profile constructor* and *privacy alarm generator*.

User Profile Constructor. It is responsible for the estimation of the user's category profile. Specifically, this module receives the categories corresponding to the tags submitted by the user, and accordingly, updates the user profile.

Suppressing Tag Generator. This module is the core of the user-side architecture as it is directly responsible for the user privacy. First, this component is provided with the user profile and a *tag suppression* rate σ, which is a parameter reflecting the proportion of tags that the user is willing to suppress. Next, this module computes the optimum tuple of suppressing tags r^*, which contains information about the tags that should be suppressed. Finally, this tuple is given to the *privacy alarm generator* module. The suppressing tag generator block is specified in Sec. 4 by means of a mathematical formulation of the trade-off between privacy and suppression, whereas Sec. 5 presents a theoretical analysis that investigates the optimization problem characterizing this trade-off.

Privacy Alarm Generator. The functionality of this module is to warn the user when their privacy is being compromised. When the user submits a tag to the system, this module waits for the category mapper block to send the category corresponding to that semantic tag. Additionally, this module receives the tuple r^* and proceeds as follows: if the category associated with that tag is included in r^*, a privacy alarm is generated to warn the user, and it is then for the user to decide whether to eliminate the tag or not. However, if that category is not contained in the tuple, the system is not aware of any privacy threat and then sends the tag to the web browser.

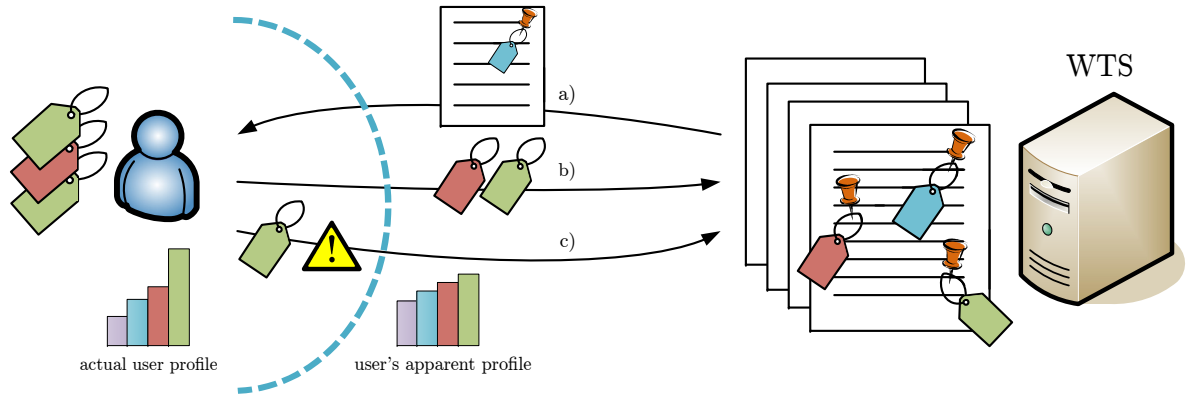

Figure 3: A user retrieves a web page and the tags submitted by the other users from a particular WTS (a). Later, the user submits their own tags to that server (b). Afterwards, the user receives a privacy alarm when trying to submit a new tag (c).

Having examined each individual component, we shall next describe how this system would work by means of the example depicted in Fig. 3. Initially, the user would specify a URL and would retrieve both the web content and the other user's tags from a particular WTS (Fig. 3a). Afterwards, the user would submit some tags to that particular WTS (Fig. 3b). The contextual information derived by the context analyzer would be used to transform these tags into categories, and then construct the user profile. The user profile would be used to calculate the tuple r^* every time this profile was updated. At a certain point, the user could receive a privacy alarm when trying to submit a tag that would contribute to make the user profile significantly different from the uniform profile. If this was the case, the user would have to decide whether to eliminate the tag or not. Finally, if this tag was suppressed, the user's apparent profile would diverge from the actual user profile (Fig. 3c).

4. Formulation of the Trade-Off between Privacy and Suppression

This section presents a formulation of the trade-off between privacy and suppression in the semantic Web, which is used to specify one of the functional blocks in Sec. 3.4.

Section 3.1 explained how certain web sites show user profiles. In particular, we mentioned that this information is normally displayed using histograms or tag clouds. Now, we provide a more formal approach to describe user profiles. Specifically, we model user *tags* as random variables (r.v.'s) taking on values on a common finite alphabet of categories or topics, namely the set $\{1, \ldots, n\}$ for some $n \in \mathbb{Z}^+$. This model allows us to describe user profiles by means of a PMF, leading to a similar representation than that shown in Fig. 1a. Accordingly, we define q as the probability distribution of the tags of a particular *user* and, in line with Sec. 3.4, we introduce a *tag suppression* rate $\sigma \in [0, 1)$, which is the ratio of suppressed tags to total tags. Thus, we define the user's *apparent* tag distribution s as $\frac{q-r}{1-\sigma}$ for some suppression strategy $r = (r_1, \ldots, r_n)$ satisfying $0 \leqslant r_i \leqslant q_i$ and $\sum r_i = \sigma$ for $i = 1, \ldots, n$. Specifically, the user's apparent tag distribution may be interpreted as the result of, on the one hand, the suppression of certain tags from the actual user profile, that is, $q - r$, and one the other, the subsequent normalization by $\frac{1}{1-\sigma}$ so that $\sum_i s_i = 1$. The information about which tags should be suppressed is encoded in the tag suppression strategy r. Namely, the component r_i is the relative frequency of tags that the proposed architecture suggests eliminating in the category i.

According to the justification provided in Sec. 3.3, we use Shannon's entropy to measure user privacy. In particular, our privacy metric is the entropy of the user's apparent tag distribution. Consistently with this measure, now we define the *privacy-suppression* function

$$\mathcal{P}(\sigma) = \max_{\substack{r \\ 0 \leqslant r_i \leqslant q_i \\ \sum r_i = \sigma}} \mathrm{H}\left(\frac{q-r}{1-\sigma}\right), \tag{1}$$

which characterizes the optimal trade-off between privacy and suppression, and formally expresses the intuitive reasoning behind tag suppression: the higher the tag suppression rate σ, the higher the uncertainty in terms of the entropy of the apparent distribution, and the higher the user privacy. In addition, this formulation allows us to describe the functional block *suppressing tag generator* in Sec. 3.4. Namely, this module will be responsible for solving the optimization problem in (1), which will be addressed in Sec. 5.

5. Optimal Tag Suppression

In this section, we shall analyze the fundamental properties of the privacy-suppression function (1) defined in Sec. 4, and present a closed-form solution to the inherent maximization problem. Our theoretical analysis only considers the case when all given probabilities are strictly positive:

$$q_i > 0 \text{ for all } i = 1, \ldots, n. \tag{2}$$

This assumption will be properly justified in Sec. 5.2. We shall suppose further, now without loss of generality, that

$$q_1 \leqslant \cdots \leqslant q_n. \tag{3}$$

Before proceeding with the mathematical analysis, it is immediate from the definition of the privacy-suppression function that its initial value is $\mathcal{P}(0) = \mathrm{H}(q)$. The behavior of $\mathcal{P}(\sigma)$ for $0 < \sigma < 1$ is characterized by the theorems presented in this section.

5.1. Monotonicity and Quasiconcavity

Theorem 5.1. *The privacy-suppression function $\mathcal{P}(\sigma)$ is nondecreasing and quasiconcave.*

Proof First, let $0 \leqslant \sigma < \sigma' \leqslant 1$. Based on the solution r to the maximization problem corresponding to $\mathcal{P}(\sigma)$, consider the tag suppression strategy r' given by the equation

$$\frac{q-r'}{1-\sigma'} = \frac{q-r}{1-\sigma}.$$

The feasibility of r' may be checked, on the one hand, by observing that the constraints $0 \leqslant r'_i \leqslant q_i$ are equivalent to $0 \leqslant \frac{q_i - r'_i}{1-\sigma'} \leqslant \frac{q_i}{1-\sigma'}$ for $i = 1, \ldots, n$. According to the implicit definition of r', we may rewrite these constraints as $0 \leqslant \frac{q_i - r_i}{1-\sigma} \leqslant \frac{q_i}{1-\sigma'}$. Given that r is feasible, the left-hand inequality is satisfied. The right-hand inequality is also verified by simply noting that $\frac{q_i}{1-\sigma} < \frac{q_i}{1-\sigma'}$. On the other hand, it is immediate to check that $\sum_i r'_i = \sigma$.

Once we have confirmed that r' is feasible, we now turn to prove the first part of the theorem. Since the feasibility of r' does not necessarily imply that r' be a maximizer of the problem corresponding to $\mathcal{P}(\sigma')$, it follows that $\mathcal{P}(\sigma') \geqslant \mathrm{H}\left(\frac{q-r'}{1-\sigma'}\right) = \mathcal{P}(\sigma)$, and consequently, that the privacy-suppression function is nondecreasing.

Finally, the quasiconcavity of the privacy-suppression function is directly proved by the fact that $\mathcal{P}(\sigma)$ is a nondecreasing function of σ. □

The quasiconcavity of the privacy-suppression function (1) guarantees its continuity on the interior of its domain, namely $(0, 1)$, but it is fairly straightforward to verify, directly from the definition of $\mathcal{P}(\sigma)$ and under the positivity assumption (2), that continuity also holds at the interval endpoint 0.

5.2. Critical Suppression

The following theorem will confirm the intuition that there must exist a tag suppression rate beyond which maximum privacy is achievable, in the sense that the privacy-suppression function attains its maximum value, that is, $\mathcal{P}(\sigma) = \ln n$. Precisely, this *critical suppression* is

$$\sigma_{\text{crit}} = 1 - n \min_i q_i = 1 - n\, q_1,$$

according to the labeling assumption (3). From the above, it is interesting to note that σ_{crit} becomes worse (closer to one) with worse (smaller) ratio $\frac{q_1}{u_1} = n\, q_1$.

Theorem 5.2 (Critical suppression). *Let u be the uniform distribution on $\{1, \ldots, n\}$, that is, $u_i = 1/n$. For all $\sigma \in [0, 1)$, if $\sigma \geqslant \sigma_{\text{crit}}$, then $\mathcal{P}(\sigma) = \mathrm{H}(u) = \ln n$. In addition, the optimal tag suppression strategy is $r^* = q - u\,(1 - \sigma)$, for which the user's apparent distribution and the uniform's match. Conversely, if $\sigma < \sigma_{\text{crit}}$, then $\mathcal{P}(\sigma) < \ln n$.*

Proof We consider only the nontrivial case when $q \neq u$, which implies that $q_1 < 1/n$ and, consequently, $\sigma_{\text{crit}} > 0$. To confirm this implication, assume $q \neq u$ and suppose now that $q_1 \geqslant 1/n$. Taking into account the labeling assumption (3) and the fact that q is a probability distribution in the sense that $\sum_i q_i = 1$, we arrive at the contradiction that q must be the uniform distribution. Given that $q_1 < 1/n$, it immediately follows that $\sigma_{\text{crit}} > 0$. The converse, that is, $\sigma_{\text{crit}} > 0$ implies $q \neq u$, is easily checked by noting that when $q_1 < 1/n$, q cannot be, by definition, the uniform distribution. On the other hand, the positivity assumption (2) ensures that $\sigma_{\text{crit}} < 1$.

Once we have determined the interval of values in which σ_{crit} is defined, we now proceed to confirm the feasibility of r^*. It is clear from its form that $\sum_i r_i^* = \sigma$, thus it suffices to verify that $0 \leqslant r_i^* \leqslant q_i$. First, observe that the right-hand inequality is satisfied for all i as $\sigma < 1$. Secondly, note that requiring that $r_i^* = q_i - \frac{1}{n}(1 - \sigma) \geqslant 0$ for all i is equivalent to $\sigma \geqslant 1 - n\, q_i$, and finally to

$$\sigma \geqslant \max_i 1 - n\, q_i = 1 - n \min_i q_i,$$

as assumed in the theorem. Interestingly, observe that the expression for the critical suppression is independent from the privacy criterion assumed. To complete the first part of the proof, it is immediate to check that the proposed r^* maximizes the user privacy, since the uniform distribution maximizes entropy.

Now it remains to prove that $\mathcal{P}(\sigma) < \ln n$ when $\sigma < \sigma_{\text{crit}}$. To this end, note that the KL divergence between the user's apparent distribution and the uniform's is

$$\mathrm{D}(s \| u) = \sum_i s_i \ln \frac{s_i}{u_i} = \ln n - \mathrm{H}(s),$$

as informally argued in Sec. 3.3. But the information inequality [47] asserts that $\mathrm{D}(s\|u) \geqslant 0$, with equality if, and only if, $s = u$ for all i. Hence, when $\sigma < \sigma_{\text{crit}}$, the solution s to the optimization problem corresponding to $\mathcal{P}(\sigma)$ satisfies that $s \neq u$, and therefore $\mathcal{P}(\sigma) = \mathrm{H}(s) < \ln n$. □

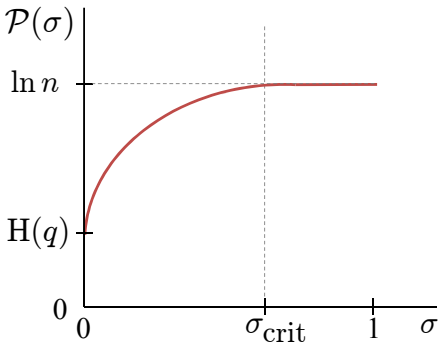

Figure 4: Conceptual plot of the privacy-suppression function.

After routine manipulation, we may write the optimal solution at exactly the critical suppression as

$$r_i^* = q_i - q_1,$$

equal to zero if, and only if, $q = u$. Owing to the fact that we are dealing with relative rather than absolute frequencies, it is not surprising that $r_1^* = 0$ at $\sigma = \sigma_{\mathrm{crit}}$. More generally, in accordance with the labeling assumption (3) observe that only the first components of r^* may vanish. Figure 4 conceptually illustrates the results stated by Theorems 5.1 and 5.2.

Before proceeding further with our theoretical analysis, we would like to remark that our assumption about the strictly positiveness of q is conveniently made, albeit not without loss of generality, to guarantee that maximum privacy be attainable for a suppression $\sigma < 1$, as proved in Theorem 5.2.

5.3. Closed-Form Solution

Our last theorem, Theorem 5.4, will provide a closed-form solution to the maximization problem involved in the definition of the privacy-suppression function (1). This solution will be obtained from a resource allocation lemma, namely Lemma 5.3, which addresses an extension of the usual water filling problem. Even though Lemma 5.3 provides a parametric-form solution, fortunately, we will be able to proceed towards an explicit closed-form solution, albeit piecewise.

More specifically, this lemma considers the allocation of resources $x_1, \ldots, x_n$ minimizing the sum $\sum_i f_i(x_i)$ of convex cost functions on the individual resources. Resources are assumed to be nonnegative, upper bounded by positive thresholds b_i, and to amount to a total of $\sum_i x_i = t$, for some $t > 0$. The well-known water-filling problem [55, §5.5] may be regarded as the special case when resources are not upper bounded and $f_i(x_i) = -\ln(\alpha_i + x_i)$, for $\alpha_i > 0$.

Lemma 5.3 (Resource allocation). *For all $i = 1, \ldots, n$, let $f_i : [0, b_i] \to \mathbb{R}$ be twice differentiable on $[0, b_i)$, with $f_i'' > 0$, and hence strictly convex. Additionally, assume that $\lim_{x_i \to b_i^-} f_i'(x_i) = \infty$. Because $f_i'' > 0$, f_i' is strictly increasing, and, interpreted as a function from $[0, b_i)$ to $f_i'([0, b_i))$, invertible. Denote the inverse by $f_i'^{-1}$. Consider the following optimization problem in the variables $x_1, \ldots, x_n$:*

$$\begin{aligned} &\textit{minimize} && \sum_{i=1}^{n} f_i(x_i) \\ &\textit{subject to} && 0 \leqslant x_i \leqslant b_i \ \textit{for all } i, \\ & && \textit{and } \sum_{i=1}^{n} x_i = t, \ \textit{for some } t > 0. \end{aligned}$$

i. *The solution to the problem exists, is unique and of the form* $x_i^* = \max\left\{0, {f_i'}^{-1}(\nu)\right\}$, *for some* $\nu \in \mathbb{R}$ *such that* $\sum_i x_i^* = t$.

ii. *Suppose further, albeit without loss of generality, that* $f_n'(0) \leqslant \cdots \leqslant f_1'(0)$. *Then, either* $f_i'(0) < \nu \leqslant f_{i-1}'(0)$ *for* $i = 2, \ldots, n$, *or* $f_i'(0) < \nu$ *for* $i = 1$, *and for the corresponding index* i,

$$x_j^* = \begin{cases} {f_j'}^{-1}(\nu), & j = i, \ldots, n \\ 0 \qquad\quad , & j = 1, \ldots, i-1 \end{cases},$$

and

$$\sum_{j=1}^{n} x_j^* = \sum_{j=i}^{n} {f_j'}^{-1}(\nu) = t.$$

Proof The existence and uniqueness of the solution is a consequence of the fact that we minimize a strictly convex function over a compact set. Systematic application of the Karush-Kuhn-Tucker (KKT) conditions [55] leads to the Lagrangian cost

$$\mathcal{L} = \sum f_i(x_i) - \sum \lambda_i x_i + \sum \mu_i (x_i - b_i) - \nu \left(\sum x_i - t \right),$$

which must satisfy $\frac{\partial \mathcal{L}}{\partial x_i} = 0$, and finally to the conditions

$$\begin{aligned} 0 \leqslant x_i \leqslant q_i, \sum x_i = t \quad & \text{(primal feasibility)}, \\ \lambda_i, \mu_i \geqslant 0 \quad & \text{(dual feasibility)}, \\ \lambda_i x_i = 0, \mu_i (x_i - b_i) = 0 \quad & \text{(complementary slackness)}, \\ f_i'(x_i) - \lambda_i + \mu_i - \nu = 0 \quad & \text{(dual optimality)}. \end{aligned}$$

Since $\lim_{x_i \to b_i^-} f_i'(x_i) = \infty$, it follows from the dual optimality condition that $x_i < b_i$. But then, the complementary slackness condition implies that $\mu_i = 0$, and consequently, we may rewrite the dual optimality condition as $f_i'(x_i) = \lambda_i + \nu$. By eliminating the slack variables λ_i, we finally obtain the simplified condition $f_i'(x_i) \geqslant \nu$. In addition, observe that since $f_i'(x_i) = \lambda_i + \nu$, the complementary slackness condition implies that $(f_i'(x_i) - \nu)\, x_i = 0$. In short, we may rewrite the dual optimality and the complementary slackness conditions equivalently as

$$\begin{aligned} f_i'(x_i) \geqslant \nu \quad & \text{(dual optimality)}, \\ (f_i'(x_i) - \nu)\, x_i = 0 \quad & \text{(complementary slackness)}. \end{aligned}$$

Now, we proceed to directly solve these equations. To this end, recall that, since $f_i'' > 0$, f_i' is strictly increasing. Consider, first, the case when $f_i'(0) \geqslant \nu$, or equivalently, ${f_i'}^{-1}(\nu) \leqslant 0$. Suppose that $x_i > 0$, so that by complementary slackness, $f_i'(x_i) = \nu \leqslant f_i'(0)$, contradicting the fact that f_i' is strictly increasing. Consequently, $x_i = 0$.

Consider now the opposite case, that is, when $f_i'(0) < \nu$, or equivalently ${f_i'}^{-1}(\nu) > 0$. In this case, the only conclusion consistent with the dual optimality condition is $x_i > 0$. But then, it follows from the complementary slackness condition that $f_i'(x_i) = \nu$, or equivalently, $x_i = {f_i'}^{-1}(\nu)$. This could be interpreted as a Pareto equilibrium. Specifically, for all positive resource $x_i > 0$, the marginal ratios of improvements $f_i'(x_i)$ must all be the same. Otherwise, minor allocation adjustments on the resources could improve the overall objective. In summary,

$$x_i = \max\{0, {f_i'}^{-1}(\nu)\},$$

which proves claim (i) in the lemma.

In order to verify (ii), observe that whenever $\nu \leqslant f'_{i-1}(0) \leqslant \cdots \leqslant f'_1(0)$ holds for some $i = 2, \ldots, n$, then $f'^{-1}_{i-1}(\nu), \ldots, f'^{-1}_1(\nu) \leqslant 0$, and thus $x_{i-1} = \cdots = x_1 = 0$. Note that the index $i = n+1$ is not permitted, since the zero solution, that is, $x_i = 0$ for all $i = 1, \ldots, n$, contradicts the primal feasibility condition $\sum_i x_i = t$. □

Next, we shall provide a close-form solution for the privacy-suppression function. However, before presenting the theorem in question, we shall introduce some notation. Let $\bar{Q}_i = \sum_{j=i+1}^{n} q_j$ denote the complementary cumulative distribution function. In addition, define

$$\sigma_i = \bar{Q}_i - q_i(n-i),$$

for $i = 1, \ldots, n$, and, conveniently, define $\sigma_0 = 1$. Note that $\sigma_n = 0$, that $\sigma_1 = 1 - n\,q_1 = \sigma_{\text{crit}}$, and consistently with Theorem 5.2, the solution in this theorem at $\sigma = \sigma_{\text{crit}}$ becomes $\frac{q_j - r^*_j}{1-\sigma} = \frac{1}{n}$, for $j = 1, \ldots, n$. Further, define

$$\tilde{q} = (q_1, \ldots, q_{i-1}, \tfrac{\bar{Q}_{i-1}}{n-i+1}, \ldots, \tfrac{\bar{Q}_{i-1}}{n-i+1}),$$
$$\tilde{r} = (0, \ldots, 0, \tfrac{\sigma}{n-i+1}, \ldots, \tfrac{\sigma}{n-i+1}),$$

a distribution in the probability simplex in $\mathbb{R}^n$, and an n-tuple representing a tag suppression strategy, respectively.

Theorem 5.4. *For any $j = 2, \ldots, n$, $\sigma_i \leqslant \sigma_{i-1}$, with equality if, and only if, $q_i = q_{i-1}$. For any $i = 1, \ldots, n$ and any $\sigma \in [\sigma_i, \sigma_{i-1}]$, the optimal suppression strategy is*

$$r^*_j = \begin{cases} 0 & , \quad j = 1, \ldots, i-1 \\ q_j - \frac{\bar{Q}_{i-1} - \sigma}{n-i+1}, & j = i, \ldots, n \end{cases},$$

and, consequently, the corresponding optimal user's apparent tag distribution is

$$s^*_j = \begin{cases} \frac{q_j}{1-\sigma} & , \quad j = 1, \ldots, i-1 \\ \frac{\bar{Q}_{i-1} - \sigma}{(1-\sigma)(n-i+1)}, & j = i, \ldots, n \end{cases}$$

Accordingly, the corresponding, maximum entropy yields the privacy-suppression function

$$\mathcal{P}(\sigma) = \mathrm{H}\left(\frac{\tilde{q} - \tilde{r}}{1-\sigma}\right).$$

Proof From the definition of σ_i and under the labeling assumption (3), it is immediate to check the monotonicity of these suppression thresholds.

Now, we proceed to prove the rest of the theorem for the nontrivial case $\sigma \in (0, 1)$. Using the definition of entropy, we may write the objective function in the (original) optimization problem (1) as $-\mathrm{H}(s) = \sum_i s_i \ln s_i$, with $s_i = \frac{q_i - r_i}{1-\sigma}$, since the maximization of entropy is equivalent to the minimization of negative entropy. Recall that r is optimal for the original problem if, and only if, r is optimal for the scaled problem. After this convenient, straightforward transformation, the objective function exposes the structure of the privacy-suppression optimization problem as a special case of the resource allocation lemma, Lemma 5.3. Specifically, the functions $f_i(r_i) = s_i \ln s_i$ of r_i are twice differentiable on $[0, q_i)$, and satisfy $f''_i > 0$ and $\lim_{r_i \to q_i^-} f'_i(r_i) = \infty$. Further, the equality constraint in (1) becomes $\sum_i r_i = \sigma$. In this special case, $f'_i(r_i) = -\frac{1}{1-\sigma}\left(\ln \frac{q_i - r_i}{1-\sigma} + 1\right)$ and

$$f'^{-1}_i(\nu) = q_i - (1-\sigma)\, e^{-(1-\sigma)\nu - 1},$$

the solution for r_i when $r_i > 0$.

The labeling assumption (3) is equivalent to the assumption that $f_n'(0) \leqslant \cdots \leqslant f_1'(0)$ in the lemma, since $f_i'(0) = -\frac{1}{1-\sigma}\left(\ln \frac{q_i}{1-\sigma} + 1\right)$ is a strictly decreasing function of q_i. From the second part of the lemma,

$$\sigma = \sum_{j=i}^{n} {f_j'}^{-1}(\nu) = \bar{Q}_{i-1} - (n-i+1)(1-\sigma)\, e^{-(1-\sigma)\nu - 1},$$

and hence,

$$\nu = -\frac{1}{1-\sigma}\left(\ln \frac{\bar{Q}_{i-1} - \sigma}{(1-\sigma)(n-i+1)} + 1\right).$$

Now it suffices to substitute ν into $f_i'(\nu)$ in order to obtain the expression for the non-zero optimal suppression strategy r_j in the theorem. The optimal user's apparent tag distribution s is easily derived from this expression.

Next, we shall confirm the interval of values of σ in which it is defined. To this end, observe that the condition $f_i'(0) < \nu$ in the lemma, is equivalent to

$$-\frac{1}{1-\sigma}\left(\ln \frac{q_i}{1-\sigma} + 1\right) < -\frac{1}{1-\sigma}\left(\ln \frac{\bar{Q}_{i-1} - \sigma}{(1-\sigma)(n-i+1)} + 1\right),$$

and finally, after routine algebraic manipulation, to

$$\sigma > \bar{Q}_i - q_i(n-i).$$

One could proceed to carry out an analogous analysis on the upper bound condition $\nu \leqslant f_{i-1}'(0)$ of the lemma to determine the interval of values of σ in which the solution is defined. However, it is simpler to realize that because a unique solution will exist for each σ, then the intervals resulting from imposing $f_i'(0) < \nu \leqslant f_{i-1}'(0)$ must be contiguous and nonoverlapping, hence, of the form $(\sigma_i, \sigma_{i-1}]$. Further, because $\mathcal{P}(\sigma)$ is continuous on $[0, 1)$, one may write the intervals as $[\sigma_i, \sigma_{i-1}]$ in lieu of $(\sigma_i, \sigma_{i-1}]$.

To complete the proof, we shall express the privacy-suppression function in terms of the optimal user's apparent tag distribution, that is, $\mathcal{P}(\sigma) = -\sum_{j=1}^{n} s_j \ln s_j$. We split the sum into two parts, namely,

$$-\sum_{j=1}^{i-1} \frac{q_j}{1-\sigma} \ln \frac{q_j}{1-\sigma} - \sum_{j=i}^{n} \frac{\bar{Q}_{i-1} - \sigma}{(1-\sigma)(n-i+1)} \ln \frac{\bar{Q}_{i-1} - \sigma}{(1-\sigma)(n-i+1)},$$

where we observe that the terms in the second sum do not depend on j. From this expression, it is straightforward to identify the terms of $\mathcal{P}(\sigma)$ as the entropy of the distribution

$$\left(\frac{q_1}{1-\sigma}, \ldots, \frac{q_{i-1}}{1-\sigma}, \frac{\bar{Q}_{i-1} - \sigma}{(1-\sigma)(n-i+1)}, \ldots, \frac{\bar{Q}_{i-1} - \sigma}{(1-\sigma)(n-i+1)}\right),$$

precisely the distribution $\frac{\tilde{q}-\tilde{r}}{1-\sigma}$, given at the end of the theorem. □

The optimal tag suppression strategy in Theorem 5.4 is interpreted as follows. On the one hand, only tags corresponding to the categories $j = i, \ldots, n$ are suppressed, which is not surprising because, precisely, these are the categories with the highest probabilities, or roughly speaking, with probabilities furthest away from the uniform distribution. On the other, the optimal user's apparent tag distribution within those categories does not depend on j, and hence they all have the same probability. Further, consistently with the fact that we are dealing with relative frequencies, the components of the apparent distribution belonging to the categories $j = 1, \ldots, i-1$ are obtained by normalizing the genuine user distribution. Figure 5 captures this intuitive analysis by illustrating a simple example with $n = 4$ categories. Namely, this figure shows a user with an actual profile q who is willing to accept a tag suppression rate $\sigma \in [\sigma_3, \sigma_2]$, causing that

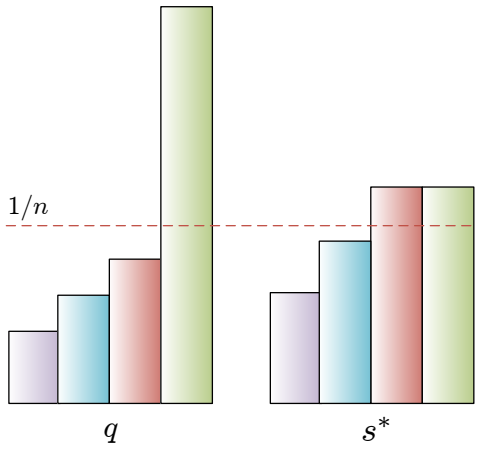

Figure 5: A user's tag distribution q and their corresponding apparent tag distribution s^* after an optimal suppression of tags.

a privacy attacker observe an optimal user's apparent profile s^* significantly different from q, specially in those categories with the highest ratio $\frac{q_j}{u_j} = \frac{q_j}{1/n}$.

A number of conclusions can be drawn from the results obtained in this last theorem. The following two sections will be focused on the analysis of the behavior of the privacy-suppression function at low suppression rates and high privacy.

5.4. Low-Suppression Case

In this section we characterize $\mathcal{P}(\sigma)$ for $\sigma \simeq 0$.

Proposition 5.5 (Low suppression). *In the nontrivial case when $p \neq q$, there exists a positive integer i with suppression thresholds satisfying $0 = \sigma_n = \cdots = \sigma_i < \sigma_{i-1}$. For all $\sigma \in [0, \sigma_{i-1}]$, the optimal tag suppression strategy r^* contains $n - i + 1$ nonzero components, and the slope of the privacy-suppression function at the origin is $\mathcal{P}'(0) = \mathrm{H}(q) + \ln q_n$.*

Proof The hypothesis $p \neq q$ implies that $n > 1$, and the existence of a positive integer i enabling us to rewrite the labeling assumption (3) as

$$q_1 \leqslant \cdots \leqslant q_{i-1} < q_i = \cdots = q_n,$$

and to express q_j as $\frac{\bar{Q}_{i-1}}{n-i+1}$, for $j = i, \ldots, n$. On account of Theorem 5.4,

$$0 = \sigma_n = \cdots = \sigma_i < \sigma_{i-1} \leqslant \cdots \leqslant \sigma_1,$$

and for all $\sigma \in [0, \sigma_{i-1}]$, we have that

$$\mathcal{P}(\sigma) = \mathrm{H}\left(\frac{\tilde{q} - \tilde{r}}{1 - \sigma}\right).$$

It is routine to check that

$$\mathcal{P}'(0) = -\sum_{j=1}^{i-1} q_j \ln q_j - \sum_{j=i}^{n} q_j \ln \tfrac{\bar{Q}_{i-1}}{n-i+1} + \ln \tfrac{\bar{Q}_{i-1}}{n-i+1} = -\sum_{j=1}^{n} q_j \ln q_j + \ln q_n,$$

where the last equality follows from the fact that $q_i = \cdots = q_n$, as shown previously. □

Now we define the *relative increment factor*

$$\delta = \frac{\mathcal{P}'(0)}{\mathcal{P}(0)} = 1 + \frac{\ln q_n}{\mathrm{H}(q)}.$$

The results from Proposition 5.5 allows us to approximate the privacy-suppression function at $\sigma \simeq 0$ as

$$\mathcal{P}(\sigma) \simeq \mathrm{H}(q) + \sigma\,(\mathrm{H}(q) + \ln q_n)$$

or, in terms of the relative increment,

$$\frac{\mathcal{P}(\sigma) - \mathrm{H}(q)}{\mathrm{H}(q)} \simeq \delta\,\sigma. \tag{4}$$

Conceptually speaking, q_n characterizes the privacy gain at low suppression, together with $\mathrm{H}(q)$, in contrast to the fact that the ratio $\frac{q_1}{1/n}$ determines σ_{crit}, the minimum suppression rate for which maximum privacy is achievable, as defined in Sec. 5.2. We mentioned in that section that $q_1 < 1/n$ in the nontrivial case when $q \neq u$. An entirely analogous argument shows that $q_n \geqslant 1/n$, with equality if, and only if, $q = u$, since the opposite, that is, $q_i < 1/n$, leads to the contradiction that $1 = \sum q_i < \sum 1/n = 1$. This result allows us to conclude that $\delta < 1$, unless $q = u$, for which, unsurprisingly, δ becomes zero. In other words, the relative privacy gain (4) is lower than the suppression introduced. Namely, the privacy increment at low suppression rates becomes less noticeable with smaller q_n, for a fixed $\mathrm{H}(q)$.

5.5. High-Privacy Case

Next, we shall analyze the case when $\sigma \simeq \sigma_{\mathrm{crit}}$ and consequently the privacy-suppression function attains its maximum value. Consider the index $i = 2$ just to confirm that, whenever $\sigma \in [\sigma_2, \sigma_{\mathrm{crit}}]$, for $q \neq u$,

$$\mathcal{P}(\sigma) = \mathrm{H}\left(\frac{\left(q_1, \frac{1-q_1}{n-1}, \ldots, \frac{1-q_1}{n-1}\right) - \left(0, \frac{\sigma}{n-1}, \ldots, \frac{\sigma}{n-1}\right)}{1-\sigma}\right) < \ln n.$$

In addition, we are implicitly assuming that $q_1 \neq q_2$, so that, on account of Theorem 5.4, $\sigma_2 < \sigma_{\mathrm{crit}}$. Consequently, we avoid an empty interval and we may express the privacy-suppression function as

$$\mathcal{P}(\sigma) = -\tfrac{q_1}{1-\sigma} \ln \tfrac{q_1}{1-\sigma} - \tfrac{1-q_1-\sigma}{1-\sigma} \ln \tfrac{1-q_1-\sigma}{(1-\sigma)(n-1)}.$$

From this expression, it is routine to conclude that $\mathcal{P}'(\sigma_{\mathrm{crit}}) = 0$ and $\mathcal{P}''(\sigma_{\mathrm{crit}}) = -\frac{1}{q_1^2 n^2 (n-1)}$, and finally,

$$\mathcal{P}(\sigma) \simeq \ln n + \frac{1}{2}\mathcal{P}''(\sigma_{\mathrm{crit}})(\sigma - \sigma_{\mathrm{crit}})^2.$$

We would like to remark that the fact that $\mathcal{P}(\sigma)$ admits a quadratic approximation for $\sigma \simeq \sigma_{\mathrm{crit}}$, with $\mathcal{P}'(\sigma_{\mathrm{crit}}) = 0$, may be concluded immediately from the fundamental properties of the Fisher information [47]. Recall that for a family of distributions f_θ indexed by a scalar parameter θ, $\mathrm{D}(f_\theta \| f_{\theta'}) \simeq \frac{1}{2}\,\mathrm{I}(\theta')(\theta' - \theta)^2$, where $\mathrm{I}(\theta') = \mathrm{E}\left(\frac{\partial}{\partial \theta'} \ln f_{\theta'}\right)^2$ is the Fisher information. Denote by $s^*_\sigma = \frac{q - r^*}{1-\sigma}$ the family of optimal apparent distributions, indexed by the suppression rate. Theorem 5.2 guarantees that $s^*_{\sigma_{\mathrm{crit}}} = u$, thus we may write $\mathcal{P}(\sigma) = \mathrm{H}(s^*_\sigma) = \ln n - \mathrm{D}(s^*_\sigma \| s^*_{\sigma_{\mathrm{crit}}})$. Under this formulation, it is clear that the Fisher information associated with the suppression is $\mathrm{I}(\sigma_{\mathrm{crit}}) = -\mathcal{P}''(\sigma_{\mathrm{crit}})$.

Finally, the observation at the end of Sec. 5.2 that $r^*_1 = 0$ at $\sigma = \sigma_{\mathrm{crit}}$ is consistent with the fact that σ_{crit} is the endpoint of the interval corresponding to the solution for r^* with $n - 1$ nonzero components in Theorem 5.4.

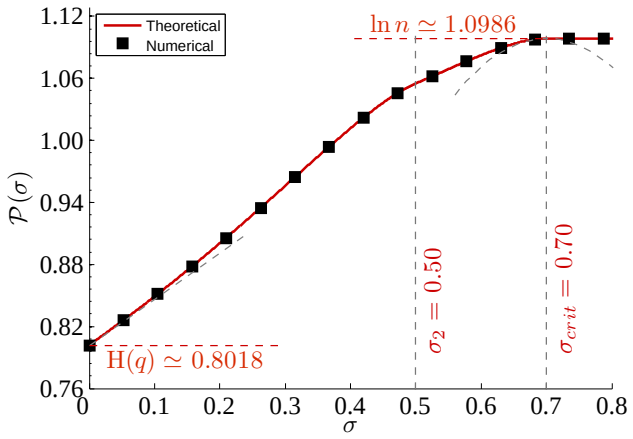

Figure 6: Optimal trade-off curve between privacy and suppression, and the corresponding approximations and suppression thresholds for $q = (0.100, 0.200, 0.700)$.

6. Experimental Results

In this section, we present some numerical results for a simple but insightful example that will illustrate the formulation presented in Sec. 4 and the theoretic analysis argued in Sec. 5. In this practical example, we shall consider three categories and assume that the user's distribution is $q = (0.100, 0.200, 0.700)$, thus fulfilling both the positivity and the labeling assumptions (2,3). On account of Theorem 5.4, the suppression thresholds are $\sigma_3 = 0$, $\sigma_2 = 0.500$ and $\sigma_1 = \sigma_{\text{crit}} = 0.700$. In addition, the initial privacy value is $\mathcal{P}(0) \simeq 0.8018$, which is the privacy level achieved by a user who is not willing to accept the suppression of any tag. Furthermore, Sec. 5.4 and 5.5 allow us to characterize the behavior of the privacy-suppression function for $\sigma = 0$ and $\sigma = \sigma_{\text{crit}}$. More specifically, the first and second order approximations are determined by the quantities $\mathcal{P}'(0) \simeq 0.4451$ and $\mathcal{P}''(\sigma_{\text{crit}}) \simeq -5.56$. All these results are captured in Fig. 6, where the privacy-suppression function $\mathcal{P}(\sigma)$ is represented. Namely, the optimization problem involved in the definition of this function has been computed theoretically, by simply applying Theorem 5.4, and numerically [(1)].

After observing the behavior of the optimal trade-off curve between privacy and suppression, now we turn to examine the optimal apparent tag distribution for a set of suppression rates. To this end, the user's distribution q, the optimal apparent distribution s^* and the uniform distribution u are represented in the probability simplexes shown in Fig. 7. In addition, the contours of the entropy $\mathrm{H}(\,\cdot\,)$ of a distribution in the simplex are depicted. More interestingly, this figure also shows the region, highlighted in dark blue, which corresponds to all the possible apparent tag distributions, not necessarily optimal, for a given suppression rate. Namely, this feasible region results from the intersection of the set $\left\{ s = \frac{q-r}{1-\sigma} \,\middle|\, 0 \leqslant r_i \leqslant q_i, \sum_i r_i = \sigma \right\}$, and the probability simplex.

We now turn our attention to Fig. 7(a), where a suppression $\sigma \in [\sigma_3, \sigma_2]$ has been selected to check that, according to the notation of Theorem 5.4, r^* has $n - i + 1 = 1$ nonzero components. Geometrically, this places the solution s^*, not entirely unexpectedly, at one vertex in the feasible region. In addition, observe that a suppression of 10 % increases the user privacy to a 5.8 % of the original privacy $\mathrm{H}(q)$. This confirms an interesting result obtained in Sec. 5.4, where we concluded that the relative increment factor δ for low-

[(1)] The numerical method chosen is the interior-point optimization algorithm [55] implemented by the Matlab R2009a function `fmincon`.

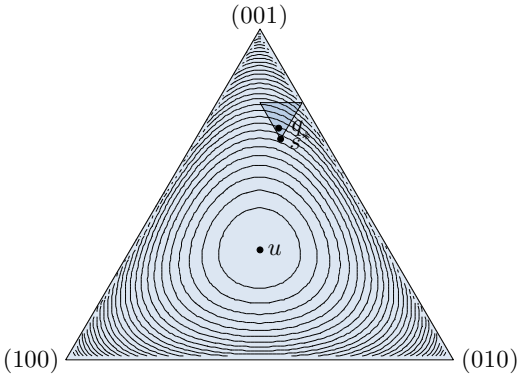

(a) $\sigma = 0.100$, $\sigma/\sigma_{\text{crit}} \simeq 0.143$, $\mathcal{P}(\sigma) \simeq 0.8487$, $\mathcal{P}(\sigma)/\mathcal{P}(0) \simeq 1.0585$, $r^* = (0.000, 0.000, 0.100)$, $s^* \simeq (0.111, 0.222, 0.667)$.

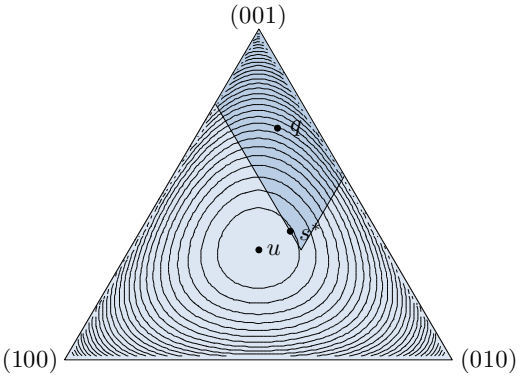

(b) $\sigma = 0.550$, $\sigma/\sigma_{\text{crit}} \simeq 0.786$, $\mathcal{P}(\sigma) \simeq 1.0688$, $\mathcal{P}(\sigma)/\mathcal{P}(0) \simeq 1.3330$, $r^* = (0.000, 0.025, 0.525)$, $s^* \simeq (0.222, 0.389, 0.389)$.

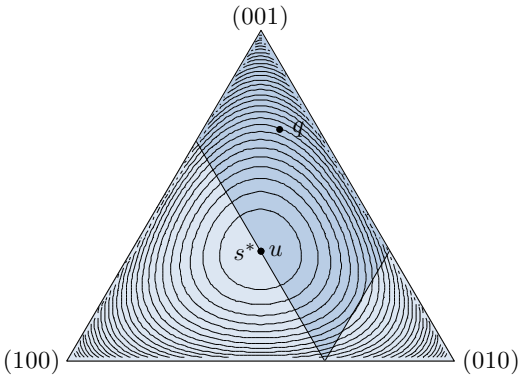

(c) $\sigma = 0.700$, $\sigma/\sigma_{\text{crit}} = 1$, $\mathcal{P}(\sigma) \simeq 1.0986$, $\mathcal{P}(\sigma)/\mathcal{P}(0) \simeq 1.3702$, $r^* = (0.000, 0.100, 0.600)$, $s^* = u$.

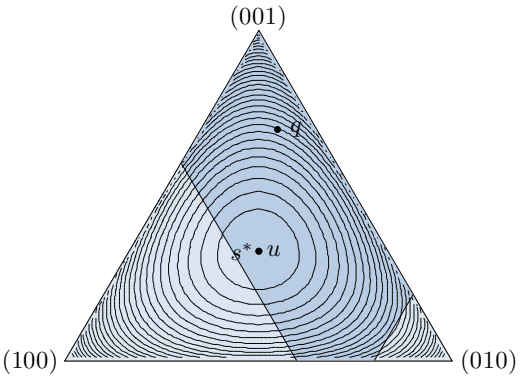

(d) $\sigma = 0.750$, $\sigma/\sigma_{\text{crit}} \simeq 1.071$, $\mathcal{P}(\sigma) \simeq 1.0986$, $\mathcal{P}(\sigma)/\mathcal{P}(0) \simeq 1.3702$, $r^* \simeq (0.017, 0.117, 0.617)$, $s^* = u$.

Figure 7: Probability simplexes showing u, q and s^* for several interesting values of σ.

suppression rates was lower than the suppression introduced. In Fig. 7(b) the suppression rate is on the interval $[\sigma_2, \sigma_{\text{crit}}]$, leading to an optimal suppression strategy r^* with $n - i + 1 = 2$ nonzero components. In this case, the solution s^* is placed on one edge of the feasible region. Additionally, note that a suppression of 55 % increments the user privacy to a 33 % of its original value. The case in which $\sigma = \sigma_{\text{crit}}$ and thus user privacy attains its maximum value is depicted in Fig. 7(c). When this happens, r^* still has $n - i + 1 = 2$ nonzero components. Precisely, note that $r_3^* = q_3 - q_1$ and $r_2^* = q_2 - q_1$, which perfectly agree with the results obtained at the end of Sec. 5.2. Finally, the case when $\sigma > \sigma_{\text{crit}}$, which certainly does not make sense, is shown in Fig. 7(d). In this particular case, r^* has $n - i + 1 = 3$ nonzero components and s^* falls into the interior of the feasible region.

7. Concluding Remarks

There exists a large number of proposals for privacy protection in the semantic Web. Within these approaches, tag suppression arises as a simple strategy in terms of infrastructure requirements, as users need not trust an external entity. Interestingly, as commented in Sec. 1.1, our technique may also be used in combination with other mechanisms such as traditional anonymous communication systems and therefore it may contribute to improve their effectiveness. Nevertheless, our approach comes at the cost of some processing overhead but more importantly at the expense of the semantic loss incurred by suppressing tags. In other words, tag suppression poses an inherent trade-off between privacy and suppression.

Our first contribution is an architecture that implements tag suppression in the semantic Web. The proposed architecture helps users refrain from proposing certain tags in order to hinder attackers in their efforts to profile users' interests. We describe the implementation details of our architecture. Specifically, the core of the system is a module responsible for obtaining an optimal tag suppression strategy. The system uses this information to warn the user when their privacy is being compromised and it is then for the user to decide whether to eliminate the tag or not.

Our second and main contribution is, precisely, a systematic, mathematical approach to the problem of optimal tag suppression. We measure user privacy as the entropy of the user's apparent tag distribution, after the suppression of tags, and justify it by the rationale behind entropy maximization methods. Subsequently, we formulate and solve an optimization problem modeling the privacy-suppression trade-off.

We model user tags as random variables taking on values on a common finite alphabet of categories or topics. This allows us to describe user profiles as PMFs, which essentially leads to the representations used in *MovieLens*, *Jinni* or *Delicious*. However, the proposed model is restricted to relative frequencies, relevant against content-based attacks, but does not deal with differences in the absolute frequencies, which certainly could be exploited by traffic analysis. Besides, we assume, on the one hand, a rudimentary adversarial model where attackers are not able to estimate a particular user's tag suppression rate, and on the other, that only a small number of users adhere to this strategy.

As a result of our theoretical analysis, we present a close-form solution for the optimal tag suppression strategy and a privacy-suppression function characterizing the optimal trade-off curve. Our mathematical approach bears certain resemblance to the water-filling problem in rate-distortion theory, and is restricted to the discrete case of n tag categories. In addition, there are several interesting mathematical analogies with [19], albeit in Sec. 2 we discarded forgery as a privacy-enhancing mechanism for semantic tagging, and deemed tag suppression a more suitable strategy.

Our theoretical study first proves that the privacy-suppression function $\mathcal{P}(\sigma)$ is nondecreasing and quasi-concave. Subsequently, we show that, under the positivity assumption (2), there exists a critical suppression $\sigma_{\text{crit}} < 1$ beyond which maximum privacy is achievable. Specifically, this σ_{crit} only depends on the minimum ratio $\frac{q_j}{u_j}$ of probabilities between the user's tag distribution q and the uniform distribution u. More interestingly, for a given suppression σ the suppression of tags only affects the categories $j = i, \ldots, n$, precisely those with the highest probabilities among all categories. Not unexpectedly, the number of categories exposed to suppression, that is, $n - i + 1$, increases with σ. In the particular case when $\sigma = \sigma_{\text{crit}}$, only the category $i = 1$ remains unchanged. With regard to the optimal user's apparent distribution, the components of s^* corresponding to the categories $j = i, \ldots, n$ have the same probability, whereas the probability of the other components is obtained by normalizing the actual user distribution.

Further, we characterize $\mathcal{P}(\sigma)$ at low suppression and high privacy. Specifically, we provide a first-order approximation for $\sigma \simeq 0$ in the nontrivial case when $q \neq u$, from which we conclude that q_n determines, together with the initial privacy value, the privacy gain at low suppression. In addition, we prove that this privacy gain is lower than the suppression introduced. Besides, we provide a second-order approximation for $\sigma \simeq \sigma_{\text{crit}}$, assuming that probabilities q_j are strictly increasing. Finally, we interpret that $\mathcal{P}'(\sigma)$ vanishes at $\sigma = \sigma_{\text{crit}}$ as a consequence of a fundamental property of the Fisher information.

Acknowledgments

This work was supported in part by the Spanish Government under Projects CONSOLIDER INGENIO 2010 CSD2007-00004 "ARES" and TSI2007-65393-C02-02 "ITACA", and by the Catalan Government under Grant 2009 SGR 1362.

References

[1] T. Berners-Lee, J. Hendler, O. Lassila, The semantic web, Scient. Amer.

[2] E. Michlmayr, S. Cazer, Learning user profiles from tagging data and leveraging them for personal(ized) information access, in: Proc. Workshop Tagging and Metadata for Social Inform. Org. Workshop in Int. WWW Conf., 2007.

[3] A. John, D. Seligmann, Collaborative tagging and expertise in the enterprise, in: Proc. Col. Web Tagging Workshop WWW, 2006.

[4] MovieLens.
URL `http://movielens.umn.edu/`

[5] Jinni.
URL `http://www.jinni.com/`

[6] Delicious.
URL `http://delicious.com/`

[7] S. Warren, L. Brandeis, The right to privacy, Harvard Law Rev. 4 (5) (1890) 193–220.

[8] D. J. Solove, Understanding Privacy, Harvard Univ. Press, 2009.

[9] M. Deng, Privacy preserving content protection, Ph.D. thesis, Katholieke Universiteit Leuven Faculty of Engineering (2010).

[10] D. Chaum, Untraceable electronic mail, return addresses, and digital pseudonyms, Commun. ACM 24 (2) (1981) 84–88.

[11] M. G. Reed, P. F. Syverson, D. M. Goldschlag, Proxies for anonymous routing, in: Proc. Comput. Security Appl. Conf. (CSAC), San Diego, CA, 1996, pp. 9–13.

[12] D. Goldschlag, M. Reed, P. Syverson, Onion routing, Commun. ACM 42 (2).

[13] R. Dingledine, N. Mathewson, P. Syverson, TOR: The second-generation onion router, in: Proc. Conf. USENIX Security Symp., Berkeley, CA, 2004.

[14] B. N. Levine, M. K. Reiter, C. Wang, M. Wright, Timing attacks in low-latency mix systems, in: Proc. Int. Financial Cryptogr. Conf.

[15] K. Bauer, D. McCoy, D. Grunwald, T. Kohno, D. Sicker, Low-resource routing attacks against anonymous systems, Tech. rep., University of Colorado (2007).

[16] S. J. Murdoch, G. Danezis, Low-cost traffic analysis of tor, in: Proc. IEEE Symp. Security, Privacy (SP).

[17] B. Pfitzmann, A. Pfitzmann, How to break the direct rsa-implementation of mixes, in: Proc. Annual Int. Conf. Theory, Appl. of Cryptogr. Techniques (EUROCRYPT), 1990, p. 373.

[18] W. M. Grossman, alt.scientology.war (1996).
URL `http://www.wired.com/wired/archive/3.12/alt.scientology.war_pr.html`

[19] D. Rebollo-Monedero, J. Forné, Optimal query forgery for private information retrieval, IEEE Trans. Inform. Theory 56 (9) (2010) 4631–4642.

[20] T. R. Gruber, A translation approach to portable ontology specifications, Knowl. Acquisition 5 (2) (1993) 199–220.

[21] D. Brickley, R. V. Guha, RDF vocabulary description language 1.0: RDF schema, W3C recommendation, WWW Consortium (W3C) (Feb. 2004).

[22] W. O. Working Group, OWL 2 Web Ontology Language: Document Overview, W3C Recommendation, 2009.

[23] A. M. Mcdonald, R. W. Reeder, P. G. Kelley, L. F. Cranor, A comparative study of online privacy policies and formats, in: Proc. Workshop Privacy Enhanc. Technol. (PET), Springer-Verlag, Seattle, WA, 2009, pp. 37–55.

[24] C. Jensen, C. Potts, C. Jensen, Privacy practices of internet users: Self-reports versus observed behavior, Int. J. Human-Comput. Stud. 63 (1-2) (2005) 203–227.

[25] L. Kagal, T. Finin, A. Joshi, A policy based approach to security for the semantic web, in: Proc. Int. Semantic Web Conf., 2003, pp. 402–418.

[26] L. Kagal, M. Paolucci, N. Srinivasan, G. Denker, T. Finin, K. Sycara, Authorization and privacy for semantic web services, IEEE J. Intelligent Syst. 19 (4) (2004) 50–56.

[27] Y. Elovici, B. Shapira, A. Maschiach, A new privacy model for hiding group interests while accessing the web, in: Proc. ACM Workshop on Privacy in the Electron. Society, ACM, Washington, DC, 2002, pp. 63–70.

[28] B. Shapira, Y. Elovici, A. Meshiach, T. Kuflik, PRAW – The model for PRivAte Web, J. Amer. Soc. Inform. Sci., Technol. 56 (2) (2005) 159–172.

[29] W. B. Frakes, R. A. Baeza-Yates (Eds.), Information Retrieval: Data Structures & Algorithms, Prentice-Hall, 1992.

[30] T. Kuflik, B. Shapira, Y. Elovici, A. Maschiach, Privacy preservation improvement by learning optimal profile generation rate, in: User Modeling, Vol. 2702 of Lecture Notes Comput. Sci. (LNCS), Springer-Verlag, 2003, pp. 168–177.

[31] Y. Elovici, C. Glezer, B. Shapira, Enhancing customer privacy while searching for products and services on the World Wide Web, Internet Research 15 (4) (2005) 378–399.

[32] R. Puzis, D. Yagil, Y. Elovici, D. Braha, Collaborative attack on Internet users anonymity, Internet Research 19 (1) (2009) 60–77.

[33] D. C. Howe, H. Nissenbaum, TrackMeNot (2006).
URL `http://mrl.nyu.edu/~dhowe/trackmenot`

[34] V. Toubiana, SquiggleSR (2007).
URL `http://www.squigglesr.com`

[35] D. Rebollo-Monedero, J. Forné, M. Soriano, Private location-based information retrieval via k-anonymous clustering, in: Proc. CNIT Int. Workshop Digit. Commun., Lecture Notes Comput. Sci. (LNCS), Springer-Verlag, Sardinia, Italy, 2009, invited paper.

[36] C. Chow, M. F. Mokbel, X. Liu, A peer-to-peer spatial cloaking algorithm for anonymous location-based services, in: Proc. ACM Int. Symp. Adv. Geogr. Inform. Syst. (GIS), Arlington, VA, 2006, pp. 171–178.

[37] D. Rebollo-Monedero, J. Forné, A. Solanas, T. Martnez-Ballesté, Private location-based information retrieval through user collaboration, Comput. Commun. 33 (6) (2010) 762–774.
URL `http://dx.doi.org/10.1016/j.comcom.2009.11.024`

[38] D. Chaum, Security without identification: Transaction systems to make big brother obsolete, Commun. ACM 28 (10) (1985) 1030–1044.

[39] G. Bianchi, M. Bonola, V. Falletta, F. S. Proto, S. Teofili, The SPARTA pseudonym and authorization system, Sci. Comput. Program. 74 (1–2) (2008) 23–33.

[40] V. Benjumea, J. López, J. M. T. Linero, Specification of a framework for the anonymous use of privileges, Telemat., Informat. 23 (3) (2006) 179–195.

[41] C. Gülcü, G. Tsudik, Mixing email with Babel, in: Proc. IEEE Symp. Netw. Distrib. Syst. Security (SNDSS), Washington, DC, 1996, pp. 2–16.

[42] G. Danezis, R. Dingledine, N. Mathewson, Mixminion: Design of a type III anonymous remailer protocol, in: Proc. IEEE Symp. Security, Privacy (SP), Berkeley, CA, 2003, pp. 2–15.
[43] G. Danezis, Statistical disclosure attacks: Traffic confirmation in open environments, in: Proc. Security, Privacy, Age Uncertainty, (SEC), Athens, Greece, 2003, pp. 421–426.
[44] M. J. Freedman, R. Morris, Tarzan: A peer-to-peer anonymizing network layer, in: Proc. ACM Conf. Comput., Commun. Security (CCS), Washington, DC, 2002.
[45] M. J. Freedman, E. Sit, J. Cates, R. Morris, Introducing Tarzan, a peer-to-peer anonymizing network layer, in: Proc. ACM Conf. Comput., Commun. Security (CCS), Washingtion, DC, 2003, pp. 193–206.
[46] G. Salton, A. Wong, C. S. Yang, A vector space model for automatic indexing, Commun. ACM 18 (11) (1975) 613–620.
[47] T. M. Cover, J. A. Thomas, Elements of Information Theory, 2nd Edition, Wiley, New York, 2006.
[48] E. T. Jaynes, On the rationale of maximum-entropy methods, Proc. IEEE 70 (9) (1982) 939–952.
[49] E. T. Jaynes, Information theory and statistical mechanics II, Phys. Review Ser. II 108 (2) (1957) 171–190.
[50] C. E. Shannon, Communication theory of secrecy systems, Tech. j., Bell Syst. (1949).
[51] A. Wyner, The wiretap channel, Tech. J. 54, Bell Syst. (1975).
[52] I. Csiszár, J. Körner, Broadcast channels with confidential messages, IEEE Trans. Inform. Theory 24 (1978) 339–348.
[53] C. Díaz, S. Seys, J. Claessens, B. Preneel, Towards measuring anonymity, in: Proc. Workshop Privacy Enhanc. Technol. (PET), Vol. 2482 of Lecture Notes Comput. Sci. (LNCS), Springer-Verlag, 2002.
[54] C. Díaz, Anonymity and privacy in electronic services, Ph.D. thesis, Katholieke Univ. Leuven (Dec. 2005).
[55] S. Boyd, L. Vandenberghe, Convex Optimization, Cambridge University Press, Cambridge, UK, 2004.

Javier Parra-Arnau was awarded the M.S. degree in electrical engineering by the Universitat Politècnica de Catalunya (UPC) in 2004. After finishing his degree, he gained a position as a project engineer in the communications department of an important Spanish engineering company. Four years later he joined the Information Security Group in the Department of Telematics Engineering at the UPC and continued to further develop his training. He was awarded the M.S. degree in Telematics Engineering in 2009 and decided to engage in research. He is currently a Ph.D. candidate at UPC, where he investigates mathematical models dealing with the inherent trade-off between privacy and data utility in information systems.

David Rebollo-Monedero received the M.S. and Ph.D. degrees in electrical engineering from Stanford University, in California, USA, in 2003 and 2007, respectively. His doctoral research at Stanford focused on data compression, more specifically, quantization and transforms for distributed source coding. Previously, he was an information technology consultant for PricewaterhouseCoopers, in Barcelona, Spain, from 1997 to 2000, and was involved in the Retevisión startup venture. During the summer of 2003, still as a Ph.D. student at Stanford, he worked for Apple Computer with the QuickTime video codec team in California, USA. He is currently a postdoctoral researcher with the Information Security Group, in the Department of Telematics of the Universitat Politècnica de Catalunya (UPC), also in Barcelona, where he investigates the application of data compression formalisms to privacy in information systems.

Jordi Forné received the M.S. degree in telecommunications engineering from the Universitat Politècnica de Catalunya (UPC) in 1992, and the Ph.D. degree in 1997. In 1991, he joined the Cryptography and Network Security Group, in the Department of Applied Mathematics and Telematics. Currently, he is an associate professor of the Telecommunications Engineering School of Barcelona (ETSETB), and works with the Information Security Group, both affiliated to the Department of Telematics Engineering of UPC in Barcelona. He is coordinator of the Ph.D. program in Telematics Engineering (holding a Spanish Quality Mention) and director of the research M.S. program in Telematics Engineering. His research interests span a number of subfields within information security and privacy, including network security, electronic commerce and public-key infrastructures. He has been part of the program committee of a number of security conferences, and he is editor of the Computer Standards & Interfaces Journal (Elsevier).

Jose L. Muñoz received the M.S. degree in telecommunication engineering from the Universitat Politècnica de Catalunya (UPC) in 1999 and the Ph.D. degree in 2003. In 2000, he joined the Information Security Group at the Department of Telematics Engineering of the UPC. His research interests include security and privacy in computer

networks. Currently, he is an associate professor at the Department of Telematics Engineering of the UPC.

Óscar Esparza received the M.S. degree in telecommunication engineering from the Universitat Politècnica de Catalunya (UPC) in 1999 and the Ph.D. degree in 2004. In 2001, he joined the Information Security Group at the Department of Telematics Engineering of the UPC. His research interests include security and privacy in computer networks and mobile agent platforms. Currently, he is an associate professor at the Department of Telematics Engineering of the UPC.