Strategic Pseudonym Change in Agent-Based E-Commerce

(Extended Abstract)

J. M. Such
Dep. de Sistemes Informàtics i Computació
Universitat Politècnica de València
Camí de Vera s/n, València, Spain
jsuch@dsic.upv.es

E. Serrano
Departamento de Ingeniería de la Información y las Comunicaciones
University of Murcia, Spain
emilioserra@um.es

V. Botti, A. García-Fornes
Dep. de Sistemes Informàtics i Computació
Universitat Politècnica de València
Camí de Vera s/n, València, Spain
{vbotti,agarcia}@dsic.upv.es

ABSTRACT

In agent-based e-commerce applications, vendors can construct detailed profiles about customers' preferences. These profiles can then be used to perform practices such as price discrimination, poor judgment, etc. The use of pseudonyms and, specially, changing pseudonyms from time to time are known to minimize profiling. Although there are some agent frameworks and platforms that support pseudonym change, there are few proposals that suggest or directly change the pseudonym in an automatic fashion. Instead, users are usually provided with the mechanisms to change pseudonyms but without any mechanism that aids them to decide when to change their pseudonyms. We present in this paper an approach to pseudonym change based on human privacy attitudes.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Theory, Experimentation

Keywords

Privacy

1. INTRODUCTION

The explosive growth of the Internet in the last decades has caused that as of 2011 more than 2 billion users are connected to it[1]. In this environment, privacy is of great concern. Users are constantly exposed to personal information collection and processing without even being aware of it [2].

[1] `http://www.internetworldstats.com/stats.htm`

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

In this paper, we focus on a type of information processing called buyer profiling [5], in which vendors obtain detailed profiles of their customers and tailor their offers regarding customer's tastes. These profiles can represent a serious threat to privacy. For instance, these profiles can be used to perform *price discrimination* [6]. This is when vendors charge customers different prices for the same good according to the customers' profiles, i.e., if a vendor knows that some good is of great interest to one customer, the vendor could charge this customer more money for this good than other customers for the same good.

Hansen et al. [4] encourage the use of pseudonyms to prevent buyer profiling. Specifically, they claim that pseudonyms should be changed from time to time to avoid profiling. Indeed, the most privacy-preserving option is to use transaction pseudonyms, i.e., to use a different pseudonym for each different transaction.

Pseudonym-based techniques have been integrated in agent technologies. Such et al. [8] present a pseudonym management model that has been implemented into a an agent framework [7]. Warnier and Brazier [10] also present a proposal for supporting pseudonym management in an agent framework. Both proposals include the necessary mechanisms for agents to be able to hold and change their pseudonyms but nothing is said about when a pseudonym should be changed or not. Moreover, the proposal of Warnier and Brazier [10] allows the automatic change of pseudonyms for each message sent. However, they do not consider the fact that there are many cases in which the user can be interested in reusing the same pseudonym even though this could cause a potential privacy loss, e.g., when some benefit is expected if they reuse the same pseudonym, such as price discounts, the building of a reputation, etc.

We present in this paper an approach to pseudonym change based on these general human attitudes towards privacy. In this way, agents obtain an estimation of the privacy loss and the utility of reusing a pseudonym. Thus, agents can automatically decide whether or not to change a pseudonym without the need of human intervention, but complying with its user's attitude towards privacy.

2. STRATEGIC PSEUDONYM CHANGE

Some studies have concluded that Humans have different general attitudes towards privacy [9, 11]. Privacy fundamentalists are extremely concerned about privacy and reluctant to lose privacy, they feel that they have already lost too much privacy and are reluctant to lose privacy any more. Privacy pragmatists are concerned about privacy (i.e. they are not willing to lose privacy a priory), but if they expect some utility (e.g. a monetary benefit) they may accept a privacy loss in exchange of this utility. Finally, privacy unconcerned do not consider privacy loss at all. A survey made in 2003 among 1.010 US adult citizens [9] shows that 26% of that citizens are considered privacy fundamentalists, 64% privacy pragmatists, and 10% privacy unconcerned.

To model these attitudes when it comes to pseudonym change, we consider that the decision of whether or not to change a pseudonym is based on a tradeoff between the privacy that will be lost if the pseudonym is not changed and the utility that will be earned if the pseudonym is not changed. For instance, in the case of privacy pragmatists, the agent can decide to not change its pseudonym in the next transaction if the privacy that will be lost is worth the utility that will be gained. We model this problem as a multi-objective optimization problem [1], in which an agent tries to minimize privacy loss while maximizing its utilitarian benefit.

One of the most used approaches to solve multi-objective optimization problems consists of transforming it into a single-objective problem[2] [3]. This is typically done by assigning a numerical weight to each objective (evaluation criterion) and then combining the values of the weighted criteria into a single value by adding all the weighted criteria.

In our case, agents consider two criteria: privacy loss and utility. Considering these two criteria, agents have two options: either to change or not to change its pseudonym in their next transaction. Thus, we are interested in measuring the quality in terms of the privacy loss and the utility of each of these options. An agent will choose the option with the highest quality. We formally define the option set as $\Theta = \{change, nochange\}$. Moreover, we define the quality of an option as:

Definition 1 (Option Quality). *Given a criterion function $c_p(\cdot)$ that evaluates privacy loss, a criterion function $c_u(\cdot)$ that evaluates utility, and weights $w_p, w_u \in [0,1]$ so that $w_p + w_u = 1$, the quality Q_δ of an option $\delta \in \Theta$ is:*

$$Q_\delta = w_p \cdot c_p(\delta) + w_u \cdot c_u(\delta) \quad (1)$$

The specific criterion functions $c_p(\cdot)$ and $c_u(\cdot)$ are domain-dependent. Moreover, as privacy loss units may be different from utility units, both criterion functions are expected to return a value in the interval $[0, 1]$ so that they can be comparable. Depending on the final domain, this could require a normalization process. This also implies that the quality of an option $\delta \in \Theta$ will be in that same interval, i.e., $Q_\delta \in [0, 1]$.

With the option quality formula, agents are able to obtain the quality of each of the options. Thus, they are able to choose whether or not to change their pseudonym in the next transaction. Agents will choose the option with the maximum quality. Formally, an agent will choose an option $\delta^* \in \Theta$ so that:

$$\delta^* = \arg\max_{\delta \in \Theta} Q_\delta \quad (2)$$

We model privacy attitudes by appropriately setting the values for the weights in the option quality formula (Equation 1), i.e., by setting w_p and w_u. If $w_p = 1$ (so $w_u = 0$) we are modeling privacy fundamentalists because they will only try to minimize privacy loss. Thus, they will not consider utility at all. If $w_p = 0$ (so $w_u = 1$) we are modeling privacy unconcerned because they will not consider privacy loss but the maximization of their utility. Finally, if $w_p \neq 1 \wedge w_p \neq 0$ we are modeling privacy pragmatists. Moreover, the specific value for w_p and w_u will vary according to how much a user valuates privacy in front of utility.

[2]For the sake of clarity and simplicity we only consider the transformation of multi-objective optimization problems into single-objective problems. However, there are other approaches to solve these kind of problems in the existing literature on multi-objective optimization (refer to [1] and [3]).

3. CONCLUSIONS

In this paper, agents decide whether to change a pseudonym or not based on the specific attitude towards privacy of their users. This specific attitude is what determines to what extent an agent valuates the privacy loss and the utility of changing/not changing a pseudonym.

4. ACKNOWLEDGMENTS

This work has been partially supported by CONSOLIDER-INGENIO 2010 under grant CSD2007-00022, and project TIN2008-04446.

5. REFERENCES

[1] K. Deb. Multi-objective optimization. In E. K. Burke and G. Kendall, editors, *Search Methodologies*, pages 273–316. Springer US, 2005.

[2] S. Fischer-Hübner and H. Hedbom. Benefits of privacy-enhancing identity management. *Asia-Pacific Business Review*, 10(4):36–52, 2008.

[3] A. Freitas. A critical review of multi-objective optimization in data mining: a position paper. *ACM SIGKDD Explorations Newsletter*, 6(2):77–86, 2004.

[4] M. Hansen, A. Schwartz, and A. Cooper. Privacy and identity management. *IEEE Security & Privacy*, 6(2):38–45, 2008.

[5] M. Hildebrandt and S. Gutwirth. *Profiling the European Citizen: Cross-Disciplinary Perspectives*. Springer Publishing Company, Inc., 2008.

[6] A. Odlyzko. Privacy, economics, and price discrimination on the internet. In *Proc. of ICEC*, pages 355–366, 2003.

[7] J. M. Such, A. Espinosa, and A. Garcia-Fornes. An agent infrastructure for privacy-enhancing agent-based e-commerce applications. In *Proc. of ITMAS*, pages 16–30, 2011.

[8] J. M. Such, A. Espinosa, A. Garcia-Fornes, and V. Botti. Partial identities as a foundation for trust and reputation. *Engineering Applications of Artificial Intelligence*, 24(7):1128–1136, 2011.

[9] H. Taylor. *Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits.* Harris Interactive, 2003.

[10] M. Warnier and F. Brazier. Anonymity services for multi-agent systems. *Web Intelligence and Agent Systems*, 8(2):219–232, 2010.

[11] A. Westin. *Privacy and Freedom*. New York Atheneum, 1967.