Towards a similarity metric for comparing machine-readable privacy policies

Inger Anne Tøndel, Åsmund Ahlmann Nyre, Karin Bernsmed

SINTEF ICT, Trondheim, Norway
{inger.a.tondel, asmund.a.nyre, karin.bernsmed}@sintef.no

Abstract. Current approaches to privacy policy comparison use strict evaluation criteria (e.g. user preferences) and are unable to state how close a given policy is to fulfil these criteria. More flexible approaches for policy comparison is a prerequisite for a number of more advanced privacy services, e.g. improved privacy-enhanced search engines and automatic learning of privacy preferences. This paper describes the challenges related to policy comparison, and outlines what solutions are needed in order to meet these challenges in the context of preference learning privacy agents.

1 Introduction

Internet users commonly encounter situations where they have to decide whether or not to share personal information with service providers. Ideally, users should make such decisions based on the content of the providers' privacy policy. In practice, however, these policies are difficult to read and understand, and are rarely used at all by users [1]. Several technological solutions have been developed to provide privacy advices to users [2–5]. A common approach is to have users specify their privacy preferences and compare these to privacy policies of sites they visit. As an example, the privacy agent AT&T Privacy Bird [2] displays icons to the user based on such a comparison, indicating whether the preferences are met or not. In general, these types of solutions provides a Yes/No answer to whether or not to accept a privacy policy. There is no information on how much the policy differs from the preferences. A policy that is able to fulfil all preferences except for a small deviation on one of the criterion, will result in the same recommendation to the user as a policy that fails to meet all the user's requirements. The user is in most cases informed about the reason for the mismatch, and can judge for himself whether the mismatch is important or not. Still, there are situations where such user involvement is inefficient or impossible, and the similarity assessment must be made automatically.

Automatic comparison of privacy policies is important to be able to give situational *privacy recommendations* to users on the web. The Privacy Finder [3] search engine ranks search results based on their associated privacy practices. Policies are classified according to a predefined set of requirements and grouped into four categories. Thus, sites that are not able to fulfil one of the basic criteria,

but offer high privacy protection on other areas will be given a low score. In order to provide more granularity and fair comparisons, a more flexible and accurate similarity metric is needed. Another application area of a similarity metric is for *preference learning in user agents* [6]. This is the application area that we focus on in this paper. To avoid having users manually specify their preferences, machine learning techniques can be utilised to *deduce* users' preferences based on previous decisions and experiences [7]. Thus, having accepted a similar policy before may suggest that the user is inclined to accept this one as well. Evidently, this approach requires a more precise mechanism to determine what constitutes a *similar* policy.

Automatic comparision of privacy policies is particularly complicated due to the subjective nature of privacy [8]. What parts of a policy are most important is dependent on the user attitude and context, and will influence how the similarity metric is to be calculated. In this paper we investigate the difficulties of defining a similarity metric for privacy policy comparison in the context of automatic preference learning. We do not consider particular privacy policy languages (such as P3P [9], PPL [10] or XACML [11]), but focus on the high-level concepts that need to be solved rather than the particular language dependent problems. The remainder of this paper is organised as follows. Section 2 gives an introduction to CBR and how it can be used to enable user agents to learn users' privacy preferences. Section 3 describes the problem of policy comparison in more detail. Section 4 takes some steps towards a solution, before Section 5 discusses the implications of our suggestions. Section 6 concludes the paper.

2 Case-Based Reasoning for privacy

Anna visits a website she has not visited before. Anna's privacy agent tries to retrieve various information on the website, including its machine-readable privacy policy. Then the agent compares its knowledge of the website with its knowledge of Anna's previous user behaviour. In this case, the agent warns Anna that the privacy policy of the website allows wider sharing than what Anna has been known to accept in the past. Anna explains to the agent that she will accept the policy since the service offered is very important to her. The agent subsequently records the decision and explanation to be used for future reference.

Case Based Reasoning (CBR) [12, 13] resembles a form of human reasoning where previously experienced situations (cases) are used to solve new ones. The key idea is to find a stored case that closely resembles the problem at hand, and then adapt the solution of that problem. Figure 1 gives an overview of the main CBR cycle. First, the reasoner retrieves cases that are relevant for the new situation. Then the reasoner selects one or a few cases (a ballpark solution) to use as a starting point for solving the new situation. Then this ballpark solution is either adapted so that it fits the new situation better, or is used as evidence for or against some solution. The solution or conclusion reached is then criticised before it is evaluated (i.e. tried out) in the real world. It is the feedback that can be gained in the evaluation step that allows the reasoner to learn. In the end, the

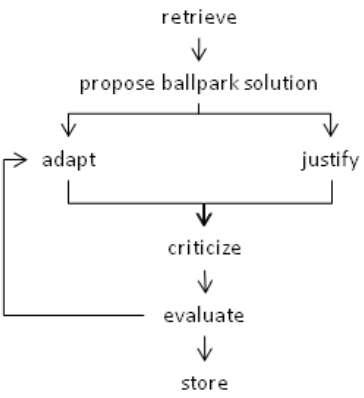

Fig. 1. CBR cycle [12]

new case is stored to be used as a basis for future decisions. Central to the CBR approach is the retrieval of *relevant* cases to use as a basis for making decisions. The prevailing retrieval algorithm is *K-Nearest Neighbour* (KNN) [14], which requires a definition of what is consider the *nearest* case. In a privacy policy setting this translates to finding the *most similar* privacy policies.

3 The similarity problem

To be able to compare privacy policies and use them in a CBR system we require the following features:

- *Similarity metric.* A method to compute a similarity metric for each of the statements of a policy. The method must be defined both for individual statements (e.g. comparing retention period) and for entire policies. Additionally it must handle missing statements so as not to simply ignore such.
- *Similarity weight.* A dynamic weight function to express the criticality or importance of a statement or part of a policy. A user or expert may for example specify that retention period is more important than conflict resolution.
- *Similarity threshold.* A threshold must be defined that determines whether two policies are *similar enough* so that one can be used as a basis to give advices on whether to accept the other.

In CBR, the similarity metric and weight function is normally what is required to compute the k-Nearest Neighbours (i.e. the k most similar policies). Thus, even if there are no policies that would be denoted similar, the algorithm will always return k policies. To cater for this, we introduce the notion of a similarity threshold such that the algorithm will return only policies that are within the threshold value.

In order to develop a similarity metric, and the corresponding similarity weight and threshold, there is a need to answer a few basic questions:

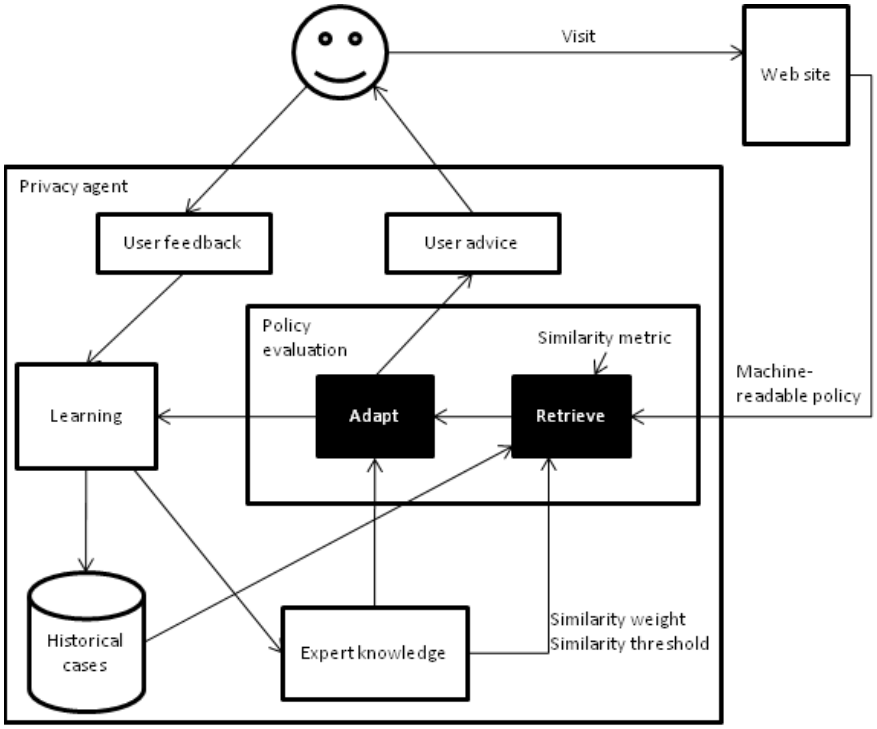

Fig. 2. Overview of solution

- *What makes policies similar?* Can policies offer roughly the same level of protection without having identical practices? And if so, how to identify this level of protection offered?
- *What type of policy content is more important for privacy decisions?* What policy changes are less likely to influence users' privacy decisions, and which changes will? And is it possible to draw conclusions in this respect without consulting the user?

There are a few surveys available that are able to give some insights into what aspects of privacy are more important to users. As an example, studies performed by Anton et al. in 2002 and 2008 [15] show that Internet users are most concerned about privacy issues related to information transfer, notice/awareness and information storage. However, in order to address the above questions in a satisfactory way in this context, more detailed knowledge is needed.

4 Towards a solution

Figure 2 gives an overview of the type of solution we envision for policy comparison in the context of a privacy agent. When the user visits a website, the

machine-readable policy of this website is used as a basis for providing the user with recommendations as to whether or not to share personal information with this site. During policy evaluation the new policy is evaluated towards the policies in historical cases in order to identify the cases that are most similar to the current situation and use those to come to a conclusion on what recommendation to give the user. Being able to identify similar cases is critical to the success of such a solution. If the cases selected as a basis for the user recommendation are not relevant, the advice the user gets will likely be irrelevant and the agent will be useless.

In order to be able to retrieve similar cases, the similarity metric is used together with the similarity weight function and the similarity threshold. To be able to compare individual statements, and assess the criticality of each statement there is a need to know, for any type of policy content, what alternative is better or worse in terms of privacy protection and how much better or worse it is. It is also necessary to have some notion of how important one statement is compared to other statements. In the figure, this is included as expert knowledge. Privacy experts are those in the best position to specify this, and let users benefit from their expertise. However, what is considered to be the most important privacy concepts will likely vary between user groups, and also individuals. It can also be dependent on the legal jurisdiction [16]. Expert knowledge can be specified in a way that takes into account some of these likely variations. However, it is also possible to make solutions that allow users to influence the expert knowledge that is used as a basis for making recommendations. We will come back to this shortly.

When the most similar cases have been retrieved, these should be adapted in order to come to a conclusion regarding the current situations. This could be done by simply taking a weighted majority-vote based on the set of cases selected. But, as the agent should be able to explain its reasoning to the user, there is also a need to build an argument that can be used to explain the agent's decision. In this process, expert knowledge also has a role to play by e.g. explaining why something is important.

The conclusion reached is presented to the user somehow. The agent may e.g. warn the user of problems with the policy, or stay silent as the policy is likely to be accepted by the user. The user, in the same way, may or may not provide feedback on this decision, e.g. by stating that he disagrees with the reasoning behind the warning. Either way, the acceptance or correction of the recommendation given is important and makes the agent able to learn and thereby improve its reasoning. A correction of the agent's reasoning may trigger a whole new re-evaluation of the policy, and result in an update to the case repository that will be in line with the preferences of the user. But the correction can also, at least in some cases, be used to improve the expert knowledge used in the policy evaluation. After all, users are experts on their own privacy preferences, and can make corrections of type "I do not care who gets my email address", or "I will never allow telemarketing, no matter the benefits"

5 Discussion

Up till now we have mainly discussed the problems related to policy comparison in a situation where historical decisions on policies are used to determine what recommendations to give to users in new situations. We have pointed at the role expert knowledge can play, and the importance that users can provide feedback on the recommendations and also influence the expert knowledge used. In the introduction we however pointed at other types of applications where automatic policy comparison can be useful, e.g. in privacy-aware search engines. So, how do our suggestions relate to such other uses?

We have considered situations in which privacy policies are compared with policies the user has accepted or rejected previously, but this is not that different from comparing a policy to those of similar types of sites to e.g. find how a web shop's privacy practices are compared to other web shops. In both cases there is no strict pre-specified matching criteria to use. Expert knowledge will be important in both cases, to assess what aspects of a policy are better or worse, and how much better or worse. But where we have been mainly concerned with identifying similar policies, other uses may be more interested in identifying policies that offer better protection, and say something about how much better this protection is.

The main research challenges that need to be addressed in order to develop a similarity metric for comparison of privacy policies are related to the expert knowledge, and in particular how expert knowledge should be presented, and how to answer the questions regarding what privacy practices are better. We do not aim at solving these research challenges in this paper, but rather take some steps towards a solution. In this respect we suggest an architecture where it is possible to start testing such metrics. With this architecture the metric and the expert knowledge can be built step-by-step by only focusing on simple policies in the beginning, and then extend the solution so that more parts of the policy can be supported.

6 Conclusion

In order to develop new and improved privacy services that can compare privacy policies in a more flexible manner than today, there is a need to develop a similarity metric that can be used to calculate how much better or worse one policy is compared to another. Expert knowledge can provide important input to such a metric. To account for individual differences in privacy attitudes, users should however have the possibility to correct and extend this expert knowledge.

References

1. C. Jensen, C. Potts, and C. Jensen, "Privacy practices of internet users: Self-reports versus observed behavior," *International Journal of Human-Computer Studies*, vol. 63, no. 1-2, pp. 203 – 227, 2005.

2. L. F. Cranor, P. Guduru, and M. Arjula, "User interfaces for privacy agents," *ACM Trans. Comput.-Hum. Interact.*, vol. 13, no. 2, pp. 135–178, 2006.
3. "Privacy Finder. http://www.privacyfinder.org."
4. J. Camenisch, A. Shelat, D. Sommer, S. Fischer-Hübner, M. Hansen, H. Krasemann, G. Lacoste, R. Leenes, and J. Tseng, "Privacy and identity management for everyone," in *Proceedings of the 2005 workshop on Digital identity management*, ser. DIM '05, 2005, pp. 20–27.
5. S. E. Levy and C. Gutwin, "Improving understanding of website privacy policies with fine-grained policy anchors," in *Proceedings of the 14th international conference on World Wide Web*, ser. WWW '05, 2005, pp. 480–488.
6. I. A. Tøndel, Å. A. Nyre, and K. Bernsmed, "Learning privacy preferences," in *Proceedings of the Sixth International Conference on Availability, Reliability and Security (AReS 2011) (to be published)*, 2011.
7. B. Berendt, O. Günther, and S. Spiekermann, "Privacy in e-commerce: stated preferences vs. actual behavior," *Commun. ACM*, vol. 48, no. 4, pp. 101–106, 2005.
8. S. A. Bagüés, L. A. R. Surutusa, M. Arias, C. Fernández-Valdivelso, and I. R. Matías, "Personal privacy management for common users," *International Journal of Smart Home*, vol. 3, no. 2, pp. 89–106, 2009.
9. "W3C. Platform for Privacy Preferences. http://www.w3.org/P3P/."
10. S. Trabelsi, "Second release of the policy engine," Prime Life, Tech. Rep. D5.3.1, 2010.
11. "OASIS eXtensible Access Control Markup Language (XACML). http://www.oasis-open.org/committees/xacml."
12. J. L. Kolodner, "An introduction to case-based reasoning," *Artificial Intelligence Review*, vol. 6, pp. 3–34, 1992.
13. A. Aamodt and E. Plaza, "Case-based reasoning: Foundational issues, methodological variations, and system approaches," *AI Communications*, vol. 7, no. 1, pp. 39–59, March 1994.
14. T. Mitchell, *Machine Learning*. McGraw Hill, 1997.
15. A. I. Anton, J. B. Earp, and J. D. Young, "How internet users' privacy concerns have evolved since 2002," *IEEE Security and Privacy*, vol. 8, pp. 21–27, 2010.
16. S. Fischer-Hübner, E. Wästlund, and H. Zwingelberg, "Ui prototypes: Policy administration and presentation version 1," Prime Life, Tech. Rep. D4.3.1, 2009.