

Privacy-Enhanced Personalization

Alfred Kobsa¹

Multi-pronged strategies are needed to reconcile the tension between personalization and privacy.

Introduction

Consumer studies have shown that online users value personalized content. At the same time, providing personalization on websites also seems quite profitable for web vendors. This win-win situation is however marred by potential privacy threats since personalizing people's interaction entails gathering considerable amounts of data about them. Numerous consumer surveys have revealed that computer users are very concerned about their privacy online. Examples for privacy concerns in connection with valued personalized services include the following (the first three services are real and the fourth is on the horizon):

- Online shoppers who value if an online bookstore can give them personalized recommendations based on what books they bought in the past may wonder whether their purchase records will be kept truly confidential in all future.
- Online searchers who are pleased that a search engine disambiguates their queries and delivers search results geared towards their genuine interests may feel uneasy that this entails recording all their past search terms.
- Students who appreciate that a personalized tutoring system can provide individualized instruction based on a detailed model of each student's understanding of the different learning concepts may wonder whether anyone else besides the system will have access to these models of what they know and don't know.
- Office workers who value if the help component of their word processor can give them personalized advice based on a model of their individual word-processing skills may be concerned that the contents of their model become accessible to others in the company, specifically when negative consequences may arise from a disclosure of what skills they lack.

Other potential perceived privacy threats in the context of personalized systems include unsolicited marketing, computers "figuring things out" about the user, fear of price discrimination, information being revealed to other users of the same computer, unauthorized access to accounts, subpoenas by courts, and government surveillance [4].

Besides being affected by individual privacy concerns, the collection of personal data is also subject to legal regulations in many countries and states (with the scope of some laws extending

¹ Alfred Kobsa (kobsa@uci.edu) is a Professor in the Donald Bren School of Information and Computer Sciences, University of California, Irvine.

beyond the national boundaries), as well as to industry codes of conduct. Both user concerns and privacy regulations not only impact the type of data that is being collected but also the methods that are employed for processing them.

Since having fewer data about users and personalization methods available is generally regarded as detrimental to the quality of personalization, the existence of a “tradeoff” between privacy and personalization and a need to “balance” privacy and personalization were postulated around the turn of the century. This perspective would suggest that figuratively speaking, an increase in personalization would result in a decrease of privacy by about the same amount, and vice versa.

More recent research has however shown that first of all, more factors than the degree of privacy and personalization need to be taken into account when looking at the overall acceptability of a personalized system from a privacy point of view. Moreover, even when considering privacy and personalization in isolation, there seem to be a number of personalization methods around that afford a significantly higher degree of privacy than traditional methods for the same purpose, with nearly the same personalization quality. The field of Privacy-Enhanced Personalization [10, 11] aims at reconciling the goals and methods of user modeling and personalization with privacy considerations, and to strive for best possible personalization within the boundaries set by privacy. The research area is widely interdisciplinary, relying on contributions from the Information and Computer Sciences, Information Systems, Marketing Research, Public Policy, Economics and Law. This article summarizes major results that were obtained so far and their implications for the design of personalized systems.

The “Privacy Calculus”

Current privacy theory regards people’s privacy-related behavior as the result of a situation-specific cost-benefit analysis, in which the potential risks of disclosing one’s personal information are weighed against potential benefits of disclosure. However, Internet users often lack sufficient information to be able to make educated privacy-related decisions. For instance, users often underestimate the probability with which they can be identified if they disclose certain data, or are unfamiliar with a site’s privacy practices since privacy statements are difficult to understand and hardly ever read. Like all complex probabilistic decisions, privacy-related decisions are moreover affected by systematic deviations from rationality. For instance, Acquisti [1] discusses the possibility of hyperbolic temporal discounting in such decisions, which may lead to an overvaluation of small but immediate benefits and an undervaluation of future negative privacy impacts.

A number of factors have been identified that play a role in the privacy calculus of Internet users. These factors include personality and culture based privacy attitudes, the type of information to be disclosed and its deviance from the average, the recipient, the value that is being assigned to personalization benefits, the extent to which users know what information has been disclosed and can control its usage, and various trust-establishing factors. We will describe these factors in more detail and discuss their consequences for the design of privacy-enhanced personalized systems.

Individual Privacy Attitudes

Various surveys established that age, education and income are positively associated with the degree of stated Internet privacy concern. Gender effects on Internet privacy concerns could not

be clearly established so far. Several surveys since the early 1980s were able to cluster respondents into roughly three groups. *Privacy fundamentalists* generally express extreme concern about any use of their data and an unwillingness to disclose them, even when privacy protection mechanisms would be in place. The *privacy unconcerned* tend to express mild concern for privacy only, and mild anxiety about how other people and organizations use information about them. *Privacy pragmatists*, finally, are generally concerned about their privacy as well, but less than the fundamentalists. They are also far more willing to disclose personal information, e.g. when they understand the reasons for its use, see benefits for doing so, or see privacy protections in place. The size ratio between these clusters is roughly 1:1:2, but the exact numbers differ noticeably across surveys and over time, with a slight decline of fundamentalists and the unconcerned over the past two decades and a corresponding increase in the number of pragmatists.

The predictive value of these attitudinal clusters is however low. Several studies as surveyed in [9] showed that privacy fundamentalists do not act much differently in situated data-disclosure decisions than the other groups. It would seem that the mitigating factors that will be discussed below play a more important role in concrete privacy decisions than abstract attitudes that are solicited out of context. Fortunately, designers can address and bolster these mitigating factors in the design of privacy-enhanced personalized systems.

Type of Information to be Disclosed

Several surveys confirm that Internet users generally feel differently about the disclosure of different types of information [9]. They are usually quite willing to disclose basic demographic and lifestyle information as well as personal tastes and hobbies. They are slightly less willing to disclose details about their Internet behavior and purchases, followed by more extended demographic information. Financial information, contact information, and specifically credit card and social security numbers raise the highest privacy concerns. An experiment by Huberman et al. [8] suggests that not only different data categories, but also the extent to which data values deviate from the group average has an effect on people's concern about their disclosure (this was verified for age, weight, salary, spousal salary, credit rating and amount of savings). The results indicate that the more undesirable a trait is with respect to the group norm, the higher is its privacy valuation.

The lesson from these findings for the design of personalized systems seems that highly sensitive data categories should never be requested without the presence of some of the mitigating factors that we discuss below. Values that deviate considerably from socially desired norms should preferably be solicited as open intervals only whose closed boundary does not deviate too much from the expected norm (such as "weight: 250 pounds and above" for male adults).

Value of Personalization

Recent surveys indicate that about 80% of Internet users are interested in personalization. While researchers today experiment with myriads of personalization services that provide various potential benefits [2], users currently only seem to value a few of them: time savings, monetary savings and to a lesser extent pleasure received the highest approval in one survey, and customized content provision and remembering preferences in another. Chellappa and Sin found that "the consumers' value for personalization is almost two times [...] more influential than the consumers' concern for privacy in determining usage of personalization services. This suggests

that while vendors should not ignore privacy concerns, they are sure to reap benefits by improving the quality of personalized services that they offer” [3].

These findings imply that developers of personalized systems need to clearly communicate the benefits of their services to users, and ascertain that they are indeed desired. If users perceive value in personalized systems, they are considerably more likely to intend to use them and to supply the required personal information.

Awareness of and Control over the Use of Personal Information

Many privacy surveys indicate that Internet users find it important to know how their personal information is being used, and to have control over this usage. In one survey, 94% agree that they should have a legal right to know everything that a web site knows about them. In another, 63% of those who indicated having provided false information to a website or declined to provide information at all said they would have supplied the information had the site provided notice about how the information would be used prior to disclosure, and if they were comfortable with these uses. In an behavioral experiment [12], website visitors disclosed significantly more information about themselves when, for every requested piece of personal information, the website explained the user benefits and the site’s privacy practices in connection with the requested data. In another study, 69% said that “controlling what information is collected about you” is extremely important, and 24% still regarded it as somewhat important.

These findings suggest that personalized systems should be able to explain to users what facts and assumptions about them are being stored, and how these are going to be used. Moreover, users should be given ample control over the storage and usage of this data. This is likely to foster users’ data disclosure, and at the same time complies with the rights of data subjects accorded by many privacy laws, industry and company privacy regulations, and Principles of Fair Information Practices that will be explained below. **Figure 1** shows a simple example from Amazon.com in which users become informed about what personal data is being used for generating recommendations, and are also given control over this use (they cannot remove data from the system though).

Trust

Trust in a website is a very important motivational factor for the disclosure of personal information. In one survey, nearly 63% of consumers who declined to provide personal information to web sites gave as the reason that they do not trust those who are collecting the data. Conversely, trust has been found to positively affect people’s stated willingness to provide personal information to websites, and their actual information disclosure to an experimental website.

Several antecedents to trust have been empirically established, and for many of them effects on disclosure have also been verified. We will discuss them in the following subsections.

Positive Past Experience

Positive experience in the past is an established factor for trust whose impact on the disclosure of personal information is well supported. Of specific importance are established, long-term

relationships. Developers of personalized systems should not regard the disclosure of personal information as a one-time matter, as is currently often the case (remember the lengthy registration form that you had to complete upon your first visit, with virtually all fields marked by an asterisk?). Users of personalized websites can be expected to become more forthcoming with personal details over time if they make positive experiences with the same or similar sites. Personalized websites should be designed in such a way that users can utilize them at least adequately with any amount of personal data they chose to disclose, and allow users to incrementally supply more information later whereupon their experience with the personalized system should improve.

Design and Operation of a Website.

Various interface design elements and operational characteristics of a website have been found to increase users' trust [6]: the absence of errors, the (professional) design and usability of a site, the presence of contact information, links from a believable website, links to outside sources and materials, updates since last visit, quick responses to customer service questions, and email confirmation for all transactions. Personalization should therefore preferably be used in professionally designed and easy-to-use websites that possess some of these trust-enhancing design elements and operational characteristics.

Reputation of the Website Operator.

Several studies found that the reputation of the organization that operates a website is a crucial factor for users' trust in the website, and for their willingness to disclose personal information [9]. In an experiment, subjects were significantly less willing to provide personally identifiable information (specifically their phone numbers, home and email addresses, and social security and credit card numbers) to lower-traffic sites which were presumably less known to them.

The lesson for the design of personalized systems seems to be that everything else being equal, users' information disclosure at sites of well-reputed companies is likely to be higher than at sites with lower reputation. Personalization is therefore likely to be more successful at more highly regarded sites, unless extra emphasis is put on other factors that foster the disclosure of personal data. Designers should not employ personalization features as a "gimmick" to increase the popularity of websites with low reputation since users may not take advantage of them if they have to disclose personal data to such sites.

Presence of a Privacy Statement

Traditional privacy statements on websites (which are often called "privacy policies") describe the privacy-related practices of these sites. The effects of privacy statements on users' trust and disclosure behavior are unfortunately somewhat unclear as yet. In several studies [9], the mere *presence* of a privacy link had a positive effect on both trust and disclosure (one experiment found a negative effect though). Inconclusive results have so far been obtained on whether the *level of privacy protection* that a privacy statement affords also has an effect on trust and disclosure. This seems unlikely for current privacy statements "in the wild" since several reading ease analyses revealed that the policies of major websites are far too difficult in their wordings to be comprehensible to the majority of web users. Not surprisingly, web server logs indicate that only a fraction of web visitors accesses privacy statements at all (less than 1% / 0.5%, according to two different sources).

The preliminary lesson for the design of personalized systems seems that traditional privacy statements should not be posted in the expectation of increasing users' trust and/or disclosure of personal information, even when statements describe good company privacy practices. There are however other good reasons for posting such statements, such as legal or self-regulatory requirements in many countries, or demonstrated good will. Evidence is mounting though that privacy-minded company practices can have a positive effect if they are communicated to web users in comprehensible forms, such as through logos that indicate the level of privacy protection (by analyzing a P3P-encoded version of the privacy policy) [7], or by explaining the implications of the privacy policy in a localized and contextualized manner [12]. Figure 1 shows examples of such strategies. More research will be needed to find such better communication forms for corporate privacy practices.



Figure 1: Communication forms for privacy practices beyond the traditional privacy statements (from www.privacyfinder.com and [12]).

Presence of a Privacy Seal

Several studies indicate that the meaning of privacy seals is not well understood by current Web users [9]. Other research found that the largest seal-awarding organization in the U.S. lacks scrutiny in the selection of its seal holders, and that sites displaying its seal are more likely to use privacy-invasive practices than sites that have no seal [5]. The presence of privacy seals at a website nevertheless has clear effects on web visitors, particularly on their perception of trust in the website, their perception of its privacy policies, and their stated willingness to disclose data. For designers of personalized systems, the pragmatic conclusion at this point is to display privacy seals as long as web users associate trust with them since doing so is likely to foster users' disclosure behavior.

Normative Approaches

To date, more than forty countries and numerous states have privacy laws enacted. Many companies and a few industry sectors additionally or alternatively adopted self-regulatory privacy guidelines. These laws and self-regulations are often based on more abstract principles of fair practices with regard to the use of personal information. In this section, we will describe some of the effects of these regulatory instruments on personalized systems.

Privacy Laws

Since personalized systems collect personal data of individual people, they are subject to privacy laws and regulations if the respective individuals are in principle identifiable. To date, more than

forty countries and numerous states have privacy laws enacted. They lay out procedural, organizational and technical requirements for the collection, storage and processing of personal data, to ensure the protection of these data and the data subjects to whom they apply. General requirements include conditions for legitimate data acquisition, transfer and processing, and the rights of data subjects. Other legal stipulations address adequate security mechanisms, disclosure duties, and the supervision and audit of personal data processing.

Some requirements imposed by privacy laws also affect the permissibility of personalization methods. Several international privacy laws prohibit the use of popular personalization methods without the user's consent. Sidebar 1 describes several such provisions.

Principles of Fair Information Practices

Over the past three decades, several collections of basic principles have been defined for ensuring privacy when dealing with personal information. So-called *Principles of Fair Information Practices* have been drafted by a number of countries as foundations for their national privacy laws (such as Australia and Canada), by supra-national organizations as guidance for their member states (such as the OECD and APEC), and by professional societies as recommendations for policy makers and as guidance for the professional conduct of their members (such as the ACM).

Developers of personalized systems should take such privacy principles into account if those are not already indirectly considered through applicable privacy laws and industry or company guidelines. Many principles have direct implications on personalized systems. Sidebar 2 shows several implications of the recommendations of the U.S. Public Policy Committee of the ACM.

1. *Value-added (e.g. personalized) services based on traffic or location data require the anonymization of such data or the user's consent.* This clause requires the user's consent for any personalization based on interaction logs if the user can be uniquely identified.
2. *Users must be able to withdraw their consent to the processing of traffic and location data at any time.* In a strict interpretation, this stipulation requires personalized systems to immediately honor requests for the termination of all traffic or location based personalization, i.e. even during the current session. A case can probably be made that users should not only be able to make all-or-none decisions, but also decisions with regard to individual aspects of traffic or location based personalization (such as agreeing to be informed about nearby sights but declining to receive commercial offers from nearby businesses).
3. *The personalized service provider must inform the user of the type of data which will be processed, of the purposes and duration of the processing, and whether the data will be transmitted to a third party, prior to obtaining her consent.* It is sometimes fairly difficult for personalized service providers to specify beforehand the particular personalized services that an individual user will receive. The common practice today is to collect as much data about the user as possible, to lay them in stock, and then to apply those personalization methods that "fire" based on the existing data (e.g. in rule-based personalization or stereotype activation). Also, internal inference mechanisms may augment the available user information by additional assumptions about the user, which in turn may trigger even more personalization processes. For meeting the disclosure requirements of privacy laws in such cases of low ex-ante predictability, it should suffice to list a number of typical personalization examples.
4. *Personal data that were obtained for different purposes may not be grouped.* This limitation affects centralized user modeling servers, which store user information from, and supply this data to, different personalized applications. Such servers must not supply user data to personalized applications for a purpose that is different from the one for which the data was originally collected.
5. *Usage data must be erased immediately after each session (except for very limited purposes).* This requirement could affect the use of machine learning methods that derive additional assumptions about users, when the learning takes place over several sessions.
6. *No fully automated individual decisions are allowed that produce legal effects concerning the data subject or significantly affect him and which are based solely on automated processing of data intended to evaluate certain personal aspects relating to him, such as his performance at work, creditworthiness, reliability, conduct, etc.* These provisions apply to, e.g., personalized tutoring applications if they assign grades to users in a fully automated manner that significantly affects them.

Sidebar 1: *Provisions from various European privacy laws that affect the permissibility of popular personalization methods (1, 2, 3 and 6: Europe-wide, 4: Czech Republic, 5: Germany)*

Minimization principles

- 1. Collect and use only the personal information that is strictly required for the purposes stated in the privacy policy.*
- 2. Store information for only as long as it is needed for the stated purposes.*
- 3. Implement systematic mechanisms to evaluate, reduce, and destroy unneeded and stale personal information on a regular basis, rather than retaining it indefinitely.*
- 4. Before deployment of new activities and technologies that might impact personal privacy, carefully evaluate them for their necessity, effectiveness, and proportionality: the least privacy-invasive alternatives should always be sought.*

Somewhat in contradiction to these requirements, a current tacit paradigm of personalized systems is to collect as much data as possible, to “lay it in stock”, and to let personalization be triggered by the currently available data. Research in quite a few personalization areas has meanwhile progressed so far that it should be possible to determine in hindsight which of the collected personal data hardly ever triggers personalization processes, and to hitherto disregard this less needed data even when it would be readily available.

Consent principles

- 5. Unless legally exempt, require each individual's explicit, informed consent to collect or share his or her personal information (opt-in); or clearly provide a readily-accessible mechanism for individuals to cause prompt cessation [...] including when appropriate, the deletion of that information (opt-out).*

One implication of this requirement for personalized systems is that personalization based on the users’ personal data must be an option that can be switched on and off at any time.

Openness principles

- 8. Whenever any personal information is collected, explicitly state the precise purpose for the collection and all the ways that the information might be used [...].*
- 10. Explicitly state how long this information will be stored and used, consistent with the "Minimization" principle.*
- 11. Make these privacy policy statements clear, concise, and conspicuous to those responsible for deciding whether and how to provide the data.*

The likely positive effect of such explanations on users’ willingness to disclose personal data was described in the main section of this article, and difficulties in explaining the personalization purposes completely were discussed in Sidebar 1.

Access principles

- 14. Establish and support an individual's right to inspect and make corrections to her or his stored personal information, unless legally exempted from doing so.*

This principle calls for online inspection and correction mechanisms for personal data, as discussed in the main article.

Accuracy principles

- 17. Ensure that personal information is sufficiently accurate and up-to-date for the intended purposes.*
- 18. Ensure that all corrections are propagated in a timely manner to all parties that have received or supplied the inaccurate data.*

So far, allowing users to verify their data seems to be the only solution for assuring data accuracy that has been adopted in the personalization literature. Little attention has been paid to recognizing the obsolescence of data, and to recording the provenance of data and propagating error and change notifications to the data sources.

Privacy-Enhancing Technology for Personalized Systems

This section discusses several technical approaches that may reduce privacy risks and make privacy compliance easier. They are by no means “complete technical solutions” to the privacy risks of personalized systems, and their presence is also unlikely to “charm away” users’ privacy concerns. Rather, these technologies should be employed as additional privacy protections in the context of a user-oriented system design that also takes normative aspects into account that were described above. These technologies are still in the research stage at the moment, but some of them seem deployable to practical applications.

Pseudonymous users and user models

It is possible for users of personalized systems to enjoy anonymity and at the same time receive full personalization [13]. In an anonymization infrastructure that supports personalization, users would be unidentifiable, unlinkable and unobservable for third parties, but linkable for the personalized system through a pseudonym. A number of authors proposed infrastructures for pseudonymous yet personalized user interaction with websites based on some or all of these properties. Several authors expect that Internet users are more likely to provide information when they are not identified, which may improve the quality of personalization and the benefits that users receive from it. To date, this claim has however not found much empirical substantiation. Designers should definitely allow for pseudonymous access and pseudonymous user models (and even allow for anonymization architectures with the above properties if one is readily available). This follows from the data minimization and security requirements of the Principles of Fair Information Practices that are discussed in Sidebar 2. Some privacy laws also mandate or recommend the provision of pseudonymous access if it is technically possible and not unreasonable. An interesting side effect of pseudonymous access is that in most cases privacy laws do not apply any more when users cannot be identified with reasonable means.

Due to a lack of relevant studies, it is unclear though whether increased anonymity will lead to more disclosure and better personalization. Anonymity is currently also difficult and/or tedious to preserve when payments, physical goods and non-electronic services are being exchanged. It harbors the risk of misuse and hinders vendors from cross-channel marketing (e.g. sending a products catalog to a web customer by postal mail). Finally, research has shown that the anonymity of database entries, web trails, query terms, ratings and textual data can be surprisingly well defeated by a resourceful attacker who has identified data available that can be partly matched with the “anonymous” data.

Client-Side Personalization

A number of researchers have worked on personalized systems in which users’ data are located at the client side rather than the server side. Likewise, all personalization processes that rely on this data are exclusively carried out at the client side. From a privacy perspective, this approach has two major advantages:

1. The privacy problem becomes smaller since very few, if any, personal data of users will be stored on the server. In fact, if a website with client-side personalization does not have control over any data that would allow for the identification of users with reasonable means, it will generally not be subject to privacy laws.

2. Users may possibly be more inclined to disclose their personal data if personalization is performed locally upon locally stored data rather than remotely on remotely stored data, since they may feel more in control of their local physical environments.

Client-side personalization also poses a number of challenges though:

1. Popular user modeling and personalization methods that rely on an analysis of data from the whole user population, such as collaborative filtering and stereotype learning, cannot be applied any more or will have to be radically redesigned (see the next section).
2. Personalization processes will also have to operate at the client side since even a mere temporary or partial transmission of personal data to the server is likely to annul the abovementioned advantages of client-side personalization. However, program code that is used for personalization often incorporates confidential business rules or methods, and must be protected from unauthorized access. Trusted computing platforms will have to be developed for this purpose.

If these drawbacks pose no problems in a specific application domain, then developers of personalized web-based systems should definitely adopt client-side personalization as soon as suitable tools become available. Doing so would constitute a great step forward in terms of the data minimization principle (see Sidebar 2) and is also likely to increase users' trust.

Privacy-enhancing techniques for collaborative filtering

Traditional collaborative filtering systems collect large amounts of information about their users in a central repository (e.g., users' product ratings, purchased products or visited web pages), to find regularities that allow for future recommendations. These central repositories may not always be trustworthy though, and may constitute attractive targets for unauthorized access. To some extent, central repositories can also be mined for individual user data by asking for recommendations using cleverly constructed profiles. A number of techniques have been proposed and partially also technically evaluated that can help protect the privacy of users of collaborative-filtering based recommender systems.

Distribution. This approach abandons central repositories containing the data of all users in favor of interacting distributed clusters that contain information about a few users only.

Aggregation of encrypted data. This approach allows users to privately maintain their own individual ratings, and a community of such users to compute an aggregate of their private data without disclosing them by using homomorphic encryption and peer-to-peer communication. The aggregate then allows personalized recommendations to be generated at the client side using the client's ratings.

Perturbation. In this approach, users' ratings are submitted to a central server, which performs all collaborative filtering. The ratings become systematically altered before submission though, to hide users' true values from the server.

Obfuscation. In this approach, a certain percentage of users' ratings become replaced by different values before the ratings are submitted to a central server for collaborative filtering. Users can then "plausibly deny" the accuracy of any of their data should they become compromised.

Conclusion

Research on Privacy-Enhanced Personalization aims at reconciling the goals and methods of user modeling and personalization with privacy considerations, and at achieving the best possible personalization within the boundaries set by privacy. As can be seen throughout this article, there exists no silver bullet for radically enhancing the privacy-friendliness of personalized systems, neither technical nor legal nor social/organizational. Instead, numerous small enhancements need to be introduced, which depend on the application domain as well as the types of data, users and personalization goals involved. Many of the approaches described here are ready to be deployed to practical systems, and feedback from such deployments will in turn be very informative for research purposes. Other approaches still need further technical development or evaluation in user experiments and may yield fruitful solutions in the future.

Acknowledgement

This research has been supported through NSF grant IIS 0308277 and by an Alexander von Humboldt Research Award. I would like to thank Lorrie Cranor, Vivian Olivera, Shilpa Rao, Sarah Spiekermann and the anonymous CACM reviewers for their valuable comments on an earlier draft of this paper. A much longer version that also includes numerous references appeared as [9].

References

1. Acquisti, A., Privacy in Electronic Commerce and the Economics of Immediate Gratification. In *EC'04 ACM Conference on Electronic Commerce*, (New York, NY, 2004), 21-29, DOI 10.1145/988772.988777.
2. Brusilovsky, P., Kobsa, A. and Nejdl, W. (eds.). *The Adaptive Web: Methods and Strategies of Web Personalization*. Springer Verlag, Berlin Heidelberg New York, 2007, DOI 10.1007/978-3-540-72079-9.
3. Chellappa, R.K. and Sin, R. Personalization versus Privacy: An Empirical Examination of the Online Consumer's Dilemma. *Information Technology and Management*, 6 (2-3), 2005, 181-202, DOI 10.1007/s10799-005-5879-y.
4. Cranor, L.F., 'I Didn't Buy it for Myself': Privacy and Ecommerce Personalization. In *2003 ACM Workshop on Privacy in the Electronic Society*, (Washington, DC, 2003), ACM Press, DOI 10.1145/1005140.1005158.
5. Edelman, B. Adverse Selection in Online "Trust" Certifications, Harvard University, 2006, <http://www.benedelman.org/publications/advsel-trust-draft.pdf>.
6. Fogg, B.J. *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann Publishers, San Francisco, 2003.
7. Gideon, J., Cranor, L., Egelman, S. and Acquisti, A., Power Strips, Prophylactics, and Privacy, Oh My! In *Second Symposium on Usable Privacy and Security*, (Pittsburgh, Pennsylvania, 2006), ACM Press, 133-144, DOI 10.1145/1143120.1143137.
8. Huberman, B.A., Adar, E. and Fine, L.R., Valuating Privacy. In *Fourth Workshop on the Economics of Information Security (WEIS05)*, (Cambridge, MA, 2005), <http://infosecon.net/workshop/pdf/7.pdf>.

9. Kobsa, A. Privacy-Enhanced Web Personalization. in Brusilovsky, P., Kobsa, A. and Nejdl, W. eds. *The Adaptive Web: Methods and Strategies of Web Personalization*, Springer Verlag, Berlin Heidelberg New York, 2007, 628-670, DOI 10.1007/978-3-540-72079-9_21.
10. Kobsa, A., Chellappa, R.K. and Spiekermann, S. (eds.). *Proceedings of CHI-2006 Workshop on Privacy-Enhanced Personalization*, Montréal, Canada, 2006, http://www.isr.uci.edu/pep06/papers/Proceedings_PEP06.pdf.
11. Kobsa, A. and Cranor, L. (eds.). *Proceedings of the UM05 Workshop 'Privacy-Enhanced Personalization'*, Edinburgh, Scotland, 2005, <http://www.isr.uci.edu/pep05/papers/w9-proceedings.pdf>.
12. Kobsa, A. and Teltzrow, M. Contextualized Communication of Privacy Practices and Personalization Benefits: Impacts on Users' Data Sharing Behavior. in Martin, D. and Serjantov, A. eds. *Privacy Enhancing Technologies: Fourth International Workshop, PET 2004, Toronto, Canada*, Springer Verlag, Heidelberg, Germany, 2005, 329-343, DOI 10.1007/11423409_21.
13. Schreck, J. *Security and Privacy in User Modeling*. Kluwer Academic Publishers, Dordrecht, Netherlands, 2003, <http://www.security-and-privacy-in-user-modeling.info>.