

Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems

Eran Toch · Yang Wang · Lorrie Faith Cranor

Received: 14 September 2010 / Accepted in revised form: 31 December 2010 /
Published online: 10 March 2012
© Springer Science+Business Media B.V. 2012

Abstract Personalization technologies offer powerful tools for enhancing the user experience in a wide variety of systems, but at the same time raise new privacy concerns. For example, systems that personalize advertisements according to the physical location of the user or according to the user's friends' search history, introduce new privacy risks that may discourage wide adoption of personalization technologies. This article analyzes the privacy risks associated with several current and prominent personalization trends, namely social-based personalization, behavioral profiling, and location-based personalization. We survey user attitudes towards privacy and personalization, as well as technologies that can help reduce privacy risks. We conclude with a discussion that frames risks and technical solutions in the intersection between personalization and privacy, as well as areas for further investigation. This frameworks can help designers and researchers to contextualize privacy challenges of solutions when designing personalization systems.

Keywords Privacy · Personalization · Human–computer interaction · Social networks · E-commerce · Location-based services

1 Introduction

New personalization technologies are becoming increasingly widespread, raising a multitude of privacy challenges. Three trends in personalization require special

E. Toch (✉)
Department of Industrial Engineering, Tel Aviv University, 69978 Tel Aviv, Israel
e-mail: erant@post.tau.ac.il

Y. Wang · L. F. Cranor
School of Computer Science, Carnegie Mellon University,
4720 Forbes Ave., Pittsburgh, PA 15213, USA

attention with regard to privacy: social-based personalization, behavioral profiling, and the mobile Web. The Web had become more social, a place where people use their real identities and communicate with their family, friends, and colleagues. As a result, applications have started to use information about a user's social networks to personalize advertising, search results and other content. Personalization algorithms have been steadily improving, making behavioral profiling more accurate and powerful. Finally, the Web had become mobile, frequently accessed through smartphones, providing new information and possibilities that can be used for personalization.

Personalization has the potential to amplify and complicate the Internet's inherent privacy risks and concerns. For example, personalized content in a social network system can reveal potentially embarrassing information directly to friends, family, and colleagues. Personalizing content according to the physical location of the user can reveal the location to unauthorized third-party entities. Examples of these types of personalization are readily apparent at many web services operating today in which users are facing a complicated privacy landscape. Recent surveys have described increasing concerns about privacy in the context of personalization. A 2010 survey by [Anton et al. \(2010\)](#) showed that privacy concerns regarding website personalization have grown significantly between 2002 and 2008. A survey by [Turow et al. \(2009\)](#) found that 66% of Americans do not want marketers to personalize advertisements to their interests, and that this attitude is consistent across age groups and gender.

Facebook's personalization attempts provide an alarming case study that can highlight the tension between privacy and personalization. On November 2007, Facebook introduced a feature called "Facebook Beacon," which allowed third-party sites to access Facebook's user profile and personalize advertisements and content accordingly. The user's activities could also be posted to the user's news feed, and made available to the user's social network. Companies such as eBay, Travelocity, and the New York Times, received access to the user's profile information and her social network in order to customize their services. The new feature encountered strong public opposition due to possible exposure of private information to friends and third-party sites. Several weeks after launching the feature, Facebook had hastily retracted it ([Story and Stone 2007](#)).

This study aims to reconcile the tension between privacy and personalization. Our objective is to provide designers with tools to build effective personalization while preserving the privacy of their users. Several survey articles have described privacy in personalization, including surveys by [Volokh \(2000\)](#), [Riedl \(2001\)](#), [Cranor \(2003\)](#), and [Kobsa \(2007b\)](#). In this article, we survey the novel privacy and personalization landscape which evolved over the last 3 years. In particular, we survey contemporary technologies that have a potentially transformative impact on privacy: social networks, behavioral profiling, and location-based Web services. We summarize the article by introducing a framework for classifying privacy risks and solutions according to the different stage of the personalization process. The framework can aid designers in understanding and resolving privacy challenges in personalization technology.

2 New privacy challenges

Contemporary personalization technologies pose a number of new risks to users' privacy. In this section, we discuss three domains of potential risks: social-based personalization, behavioral profiling, and location-based personalization.

2.1 Social-based personalization

The exponential growth of social network systems (SNS) in the last few years has created a huge online repository of real identities, unparalleled to anything known before. SNSs such as Facebook, Twitter, MySpace, Orkut, LinkedIn, Kaixin001, and Renren have a combined total of over 900 million users (Wang et al. 2011). Most of the social networks store rich information about their users, including real names, email addresses, list of friends, demographics, personal photos, location, inter-personal communications, and more. This information is used for various means of personalization, in domains such as application customization, social search, and online marketing. However, implementing privacy preserving personalization in SNS is particularly challenging. First, social networks include highly sensitive information, because social networks enable in-person communication, people are often willing to reveal more private information than they would otherwise (Acquisti and Gross 2006; Stutzman and Kramer-Duffield 2010). Second, personalizing content according to the user's friends may compromise not just the user's privacy, but also her friends' privacy. Third, releasing information within the social network environment has the potential to embarrass the user in the face of friends, family, and colleagues. The consequences can be severe: in 2008, 8% of U.S. companies employing 1000 workers or more had reported firing an employee because of information released on online social networks (Forrester Consulting 2008). In this section, we describe marketing and social search as prominent examples of privacy risks in social-based personalization.

Facebook's endeavors into personalization provide a telling story of the consequences of not addressing privacy concerns when deploying social-based personalization in online marketing. The "Beacon" advertising program was launched in November 2007. Partner companies were allowed to access users' profiles, and to present personalized advertising to those users. Certain activities, such as purchasing a product or adding a product to a wish list, were published to the user's friends. The service had encountered resistance in the forms of online petitions protesting the program, negative media attention, and a class action lawsuit. The fact that Beacon allowed the users friends to learn about the goods and services they buy and view online was the most cited concern against the new feature (Story and Stone 2007). As a result, the service was shut down in September 2009. A year and a half later, Facebook encountered similar resistance to another feature, called "Instant Personalization," which was introduced in April 2010 (Facebook 2010). With instant personalization, service providers were allowed access to SNS profiles automatically when users visited their site, without requiring explicit consent from users. Instant Personalization was criticized by users, media, and regulators, and was quickly retracted by Facebook (Helft and Wortham 2010).

The potential enhancing personalization using information drawn from social networks has been increased dramatically by the introduction of social network application programming interfaces (API). Large SNSs allow third-party applications to access users' profiles through an API, such as Facebook's "Facebook connect." The interface allows third-party applications to access user data and to publish stories and information to the user's friends. Users often have limited control over the applications that access their information (Bonneau and Preibusch 2009), and are typically asked to authorize application access only when starting the application at the first time. This relative ease of access to social network data through the API poses a privacy risk to users. For example, Facebook applications could access almost everything on a user profile, including hometown, groups the user belong to, events attended, favorite books, and more.

Social search, personalized recommendation, and targeted advertising are the most widely used forms of social-based personalization (Paliouras 2012). In social search, search results are personalized according to the user's social network (Heymann et al. 2008; Dalal 2007). Services such as Delver, Sightix, SideStripe, and Google social search, take into account users' social networks in different ways. For example, Google +1 search feature finds relevant content from your friends and contacts and highlights it for you on the regular search results page (Google 2011). Other social search engines, such as Microsoft Bing, show user-generated content from social networks such as Twitter and Facebook, and personalize the results according to the social distance from the user posting the query. Social search exemplifies another vulnerability of social-based personalization: it can reveal information about the user's social network that might not be easily accessible through regular interaction. For example, if a friend had marked a search result about cancer treatment, this information can be revealed when an acquaintance on the user's social network will search for the same thing. This scenario exemplifies how personalization can eliminate the original context for the interaction and create unexpected results.

Users have a growing sensitivity towards privacy in SNSs. In 2005, Acquisti and Gross showed that Facebook users have concerns regarding their privacy, but only a small minority of them act upon these concerns and protect their public profile (Acquisti and Gross 2006). However, in a 2010 article by Stutzman and Kramer-Duffield, which studied users with comparable demographics to the Acquisti and Gross paper, a majority of users had friends-only profiles (Stutzman and Kramer-Duffield 2010). Surveys have shown that privacy concerns increase when the user is more active in the social network (Lewis et al. 2008) and when she has more social contexts that are included in the social network (Lampinen et al. 2009). The last factor is particularly relevant to personalization, as personalization has the potential of introducing new contexts to social networking, sharing private information in ways which the user did not foresee.

Users are not only concerned about their private information being used, but also about unintentional flow of information to their social network. Knowledge about a user's social network is in itself a private piece of information and a sensitive medium through which private information can leak. Consider, for example, the uproar concerning exposed Gmail contacts through Google Buzz (Mullins 2010). An unintended leakage of information compromises the user's ability to control her public persona and

interfere with the interaction of the user with her social network (Palen and Dourish 2003). Furthermore, surprising the user by sharing unintended information with her social network can violate social norms, thus harming the integrity of the contextual use of the social network (Nissenbaum 2004).

It is important to emphasize that not all SNS are the same. Different SNS users have different privacy approaches, and different SNS have different privacy controls, varying in expressiveness and functionality (Bonneau and Preibusch 2009). The differences can be explained by the commercial motivation behind the SNS, as well as the overall sensitivity towards sharing information between different user communities. Dwyer et al. (2007) found that Facebook users have a greater sense of trust in other Facebook users and thus reveal more information than users on MySpace. Some differences regarding privacy approaches and behaviors can be dependent on cultural approaches regarding trust, privacy, and information sharing. A recent study by Wang et al. (2011) found that American social network users have more privacy concerns than their Chinese and Indian counterparts, and that the type of information regarded as private differs between cultures.

2.2 Behavioral profiling and data aggregation

Behavioral profiling is the practice of collecting longitudinal data about a person's activities and tailoring the user experience based on those activities. Unlike systems that rely on data actively provided by the user, in most instances of behavioral profiling the system tracks a wide range of user behaviors, with little or no consent of users. In recent years, behavioral profiling has become common in many domains, including Internet advertising, Web search and electronic commerce. Profiles are created based on a variety of different activities, including sites visited, products purchased, product pages viewed, emails sent, and so forth. Most behavioral profiling systems track users over long periods of time using browser cookies, which identify the user between consecutive sessions.

Behavioral profiling poses several potential privacy risks. An FTC recently highlighted the concerns associated with behavioral advertising, including the invisibility of data collection to users and the risk that data would fall into undesired 3rd party hand (Federal Trade Commission 2009). One risk cited by many users is unsolicited marketing (Culnan and Milne 2001). Another risk associated with personalization is that personal information will be inadvertently revealed to other users of the same computer. As cookies are used to identify users, users who share the same computer and Web browser may view each other's ads. Moreover, as Internet advertisement providers such as Google are linking behavioral profiles to server-side user accounts, advertisements can be shown across different computers. Thus, there is a possibility that someone who does not share the user's computer will gain access to her personalized content. Finally, users report that the mere feeling of being watched and tracked, is disturbing (McDonald and Cranor 2010).

Sophisticated analysis can be used to derive new facts about the user model, and therefore can potentially contradict the user's initial perceptions of the system context in which the data was collected. A prominent example of these types of analysis

is collaborative filtering (Schafer et al. 2007). The underlying assumption of collaborative filtering is that a user will like things that similar users like. It computes clusters of users who share similar tastes based on their previous preferences, and then uses preferences (e.g., ratings) of these similar users to make predictions. Traditional CF systems collect and store user information (e.g., users' product ratings) in a centralized repository. This configuration can raise security risks, as well as privacy risks. Studies have shown that some users are uncomfortable being watched and analyzed (Cranor 2003). As Nissenbaum (2004) compellingly argues, the mere action of context switching holds privacy risks as it can lead to violating the integrity of the original context. In collaborative filtering, for example, the user model is enriched with information inferred from other users.

Personalization based on information aggregation is becoming widespread in industry and also as a research field in academia. The barriers towards large scale information aggregation include legal restrictions (Wang and Kobsa 2006) as well as technological hurdles. In a series of articles, Mehta and other researchers show how user models can be exchanged between systems to provide cross-system personalization using unified semantic descriptions (Mehta et al. 2005) or machine learning algorithms Mehta (2006, 2007). This functionality will enable personalization systems to use information from other systems, for example, a music personalization service can use the user's book recommendations to offer personalized content. However, the information used by the personalization system is inherently used under a different context than it was collected and has the potential of surprising or embarrassing the user. For example, the Google social graph API Google (2008) is an example for semantic data aggregation of social network information, which can result in blending different social contexts.

User attitudes towards behavioral profiling were studied in several surveys and user studies. A recent survey by the TRUSTe organization had found that only 28% of Internet users would feel comfortable with advertisers using web browsing history, and 35% believe their privacy has been invaded in the past year due to information on the Internet, even if their browsing history is not tied to their actual name (TRUSTe and TNS 2009). The concerns regarding tailored advertising are also reflected in a nationally representative phone survey conducted in 2009, which found that 66% of adults do not want tailored advertising. This increased to 86% when participants were informed of techniques such as third-party cookies (Turow et al. 2009). User understanding of the techniques behind behavioral advertising is limited. A study by McDonald and Cranor reported that a small percentage of users understand the mechanism behind behavioral tracking, and when told about these mechanisms most users found them concerning (McDonald and Cranor 2010). Furthermore, There is growing evidence that privacy concerns impact the behavior of both users and marketers. Tsai et al. found that when presented with privacy information in search results, lab study participants often paid a small premium to shop at websites with good privacy policies (Tsai et al. 2011). The timing of presenting the privacy information (e.g., while searching or while showing the product information) has a significant effect on user behavior (Egelman et al. 2009). A recent survey by the Media Post reported that some marketers are limiting their behavioral targeting due to privacy issues (Davis 2010).

2.3 Location-based personalization

Personalized location-aware services are becoming more and more widespread. Development has been triggered by the adoption of GPS-enabled phones and WiFi positioning technologies, as well as the increase in mobile data bandwidth. As the sophistication of mobile devices grows, so is the ability of service providers to continuously track the location of their users, offering them services based on their exact physical location. It has become easier to develop mobile applications that request the exact location of the user through the use of easily accessible operating system and browser application programming interfaces. For example, the iPhone operating system and Google's Android have location services programming interfaces that allow background location requests. Operating systems such as Windows 7 and Mac OS 10, as well as browsers such as Mozilla Firefox, exhibit a location programming interface that allows application developers and website builders to request the user's physical location.

Physical location is used in various ways for personalization. [Lu et al. \(2010\)](#) and [Yi et al. \(2009\)](#) use users' physical location information to improve their personalized search results. This is becoming a common practice. In addition, search results displayed on smartphones are selected according to the user's location, highlighting local venues and services. Services for mobile advertising, e.g. Apple iAd and Google, offer personalized ads based on location. Various services, including venue finders (e.g., Yelp, Urban Spoon), personal classifieds (e.g., Grindr, OkCupid, Skout), weather and news applications, and so forth use location for personalization. E-commerce services, such as Shopkick, are using location-based price discrimination, offering coupons based on location and user profile.

According to a survey by [Tsai et al. \(2010\)](#) users have several concerns regarding location-based service. These concerns are (in the order of importance): being stalked, revealing home location, being tracked by their boss, being tracked by the government, and being bothered by location-based ads. The primary privacy concerns surrounding the disclosure of this information include *context* and *use* ([Barkhuus et al. 2008](#)). Privacy concerns can depend on the situation or activity in which the user may be engaged ([Iachello et al. 2005](#)). A study by [Benisch et al. \(2011\)](#) empirically evaluated methods for privacy control in location sharing scenarios. The results show that certain contexts, identified using time or location, are considered more private by users. Expressive privacy controls that allow users to specify particular locations or times in which they will not receive advertisements were found to better express the users' preferences. A study by [Toch et al. \(2010\)](#) presented a model for privacy in location sharing applications, showing that users are less comfortable sharing places which are less frequently visited by the general population.

3 Reducing privacy risks

The privacy risks highlighted in the section above can be reduced by carefully examining technological architectures that mitigate those risks. In this section, we describe existing approaches to system design that allow personalization

methods to operate while minimizing the potential risks to users. We discuss the tradeoffs between personalization effectiveness and the privacy risks for each of the approaches.

3.1 Pseudonymous personalization

Pseudonymous personalization allows users to use pseudonyms in a personalized system. This approach enables the system to track the same pseudonym across different sessions and provide personalized services without knowing the true identity of the pseudonym. Several systems allow a user to create and maintain more than one pseudonymous persona, so that the user can separate different aspects of her online activities (e.g., work versus entertainment) and can control which service provider can have access to certain persona (Arlein et al. 2000; Hitchens et al. 2005). Kobsa proposed a system that goes one step further by hiding not only the identities of the users but also the location of the user modeling servers in the network (Kobsa and Schreck 2003). From the perspective of “identifiability of data,” pseudonymous personalization has a clear advantage because it seems to hide identity and, in most cases, privacy laws are not applicable when the interaction is anonymous. However, anonymity is still currently difficult or tedious to preserve when payments, physical goods, and non-electronic services are involved. Moreover, it has been shown repeatedly that only hiding explicit identity information (e.g., name, email address) is not sufficient to hide identity because other data sources may be used to uncover one’s identity (e.g., database entries (Sweeney 2002), web trails (Malin et al. 2003), and query terms (Nakashima 2006)).

3.2 Client-side personalization

The key idea behind client-based personalization is that users’ data are stored at the client side (e.g., users’ computers or mobile phones) and subsequent personalization processes also take place at the client side (Cassel et al. 2001; Ceri et al. 2004; Coroama and Langheinrich 2006; Mulligan and Schwartz 2000; Gerber et al. 2010). Since data collection and processing occur at the client side rather than the server side, users may perceive more control over their data and perceive less privacy risk. However, these assumptions still lack empirical evidence. Furthermore, if the server side (service provider) does not have user data and cannot identify their users with reasonable effort, privacy laws may not apply to the system. However, this approach also faces challenges. First, existing personalization algorithms such as collaborative filtering and stereotype learning (see Kobsa et al. 2001), may need to be redesigned to fit the client-side model. Second, program code that is used for personalization often incorporates confidential business rules or methods, and must be protected from disclosure through reverse engineering. Trusted computing platforms will therefore have to be developed for this purpose, similar to the one that Coroama and Langheinrich (Coroama and Langheinrich 2006) envisage to ensure the integrity of their client-side collection of personal data.

3.3 Distribution, aggregation and other privacy-preserving techniques

A number of distribution, aggregation, perturbation and obfuscation techniques have been proposed to help protect user privacy in recommender systems that employ collaborative filtering (CF) (Schafer et al. 2007). One possible strategy is to distribute user data, preferably across users' own machines. This is akin to the client-side personalization approach. However, this measure alone cannot solve the privacy issues in CF systems because CF needs other users' data to make recommendations for a user. (Miller et al. 2004) proposed a distributed CF algorithm that updates a user's interest model incrementally by incorporating one neighbor's ratings at a time (ratings are immediately discarded thereafter). Another strategy is to use encrypted aggregation of users' data. For instance, Canny (Canny 2002) proposed a multi-party computation scheme that allows a community of users to compute an aggregate of their data (i.e., a singular value decomposition model of the user-item matrix) based solely on vector addition so that individual data will not be disclosed. The anonymous aggregate can then be used along with one's own ratings at the client-side to make personalized recommendations. Mehta (2007) used a similar algorithm to preserve data anonymity in a distributed cross-site personalization system.

Other approaches towards privacy preserving collaborative filtering systems include perturbation and obfuscation. Perturbation means users' ratings get systematically altered using a perturbation function (e.g., adding random numbers to ratings) before submission to a central repository so as to hide users' true values from the server. For example, Polat and Du (Polat and Du 2003, 2005) showed that using perturbed data may still yield acceptable recommendations. Obfuscation means that a certain percentage of users' ratings get replaced by random values. For instance, Berkovsky et al. (2005) also showed that their obfuscation schemes (e.g., replace with the mean) worked reasonably well. Both perturbation and obfuscation affords some degree of "plausible deniability" because some user data are not true. All these alternative strategies show promising results in CF systems, but their application to other non-CF personalized systems is still an open question.

3.4 User controls and feedback

Judy Kay and her colleagues (Kay et al. 2003; Kay 2006) suggested putting scrutability into user modeling and personalized systems. The term "Scrutability" signifies the ability of users to understand and control what goes into their user model, what information from their model is available to different services, and how the model is managed and maintained. Their user modeling system, Personis, applies three privacy-enhancing mechanisms to control the protection of each unit of personal information ("evidence") in the user model (Kay et al. 2003):

- expiration dates and purging of older evidence,
- compaction, for replacing a set of evidence from a single source with an aggregate, and
- morphing, which replaces an arbitrary collection of evidence.

For controlling the usage of evidence from the user model, Personis allows users to restrict the evidences available to applications, and the methods that may generate a user model and operate on it. Despite the desirability of scrutability from a privacy point of view, its implementation and control is currently very challenging due to users' lack of understanding of these notions and of effective and efficient user interfaces to support them. Moreover, scrutability may reveal the personalization methods that a website uses, which may pose a problem in application areas in which those are considered to be competitive advantages and therefore confidential (e.g., in online retail websites).

Wang and Kobsa (2007) proposed a way for users to specify their privacy preferences and then the underlying user modeling systems will dynamically reconfigure itself by only selecting and using personalization methods that are permissible under users' current privacy preferences and applicable privacy regulations. The result of their lab experiment showed that this mechanism increases users' disclosure of data and the likelihood of making online purchases (Wang 2010).

In helping users set and control their privacy preferences, Kelley et al. (2008) proposed an approach called user-controllable policy learning that utilizes machine learning techniques to learn and predict users' privacy preferences from users' privacy decisions. More specifically, the system can suggest privacy preference decisions to a user and can learn from the user's decisions of whether to accept the system suggestions. Users can control this learning process because (1) users can decide whether to accept or reject the system suggestions, and (2) users can directly modify their privacy preferences in the system anytime. Their evaluation of applying this approach in a location-sharing application yields promising results (about 90% accuracy of preference predictions).

General solutions for controlling data collection that are relevant to personalization include opt-out cookies and the "Do-not-track" mechanism for Internet usage tracking (Mayer and Narayanan 2011). The mechanism adds an HTTP header to express the user's intent to opt out of Internet tracking.

3.5 Privacy-preserving location tracking

Research on privacy-preserving location tracking has built on existing privacy preserving techniques such as anonymity and perturbation techniques. Beresford and Stajano (2003) designed a mix network approach to enhance location privacy. Gruteser and Liu (2004) devised disclosure-control algorithms that hide users' location in sensitive areas. Hoh and Gruteser (2005) developed a scheme of perturbing path information of a user to confuse an adversary. Tang et al. (2006) designed a hitchhiking approach of achieving location privacy essentially by supporting anonymous location reporting by individual users. Ristenpart et al. (2008) built a system, Adeona, that uses cryptographic techniques and distributed hash tables to hide a mobile device's visited locations from third-party services and other parties. Gedik and Liu (2008) proposed a location privacy framework based on a k-anonymity scheme (Sweeney 2002) and a suite of location perturbation algorithms. In addition, Tsai et al. (2010) advocated support for user controls in location sharing. Benisch et al. (2011)

conducted an empirical study that asked users to rate when, where, and with whom they would feel comfortable sharing their location information. Their study demonstrated that users have fine-grained location privacy preferences, for example based on time of day, which are not currently supported by commercial location sharing tools.


4 Discussion

The changing face of personalization technologies requires new frameworks for understanding privacy risks and solutions in personalization systems. In this section, we summarize our survey, and propose a framework to help system designers to understand and address privacy risks in personalization systems.

4.1 Analyzing privacy risks

The new technologies presented in this article impact privacy in several distinctive ways. Some technologies enable new kinds of information to be collected, other technologies provide new ways to analyze data, and the rest enable the possibility of distributing the personalized content in new ways. To understand the effects of these new technologies, and to guide the design of privacy enhancing approaches, we construct a simple framework for classifying privacy aspects in personalization. The framework, depicted in Fig.1, is based on two dimensions: personalization phases and privacy control. The personalization phases are taken from the personalization literature Kobsa (2001, 2007a), and include data collection, user model creation, and adaptation. Each of the three phases impose different challenges on privacy and require different measures to address these challenges. The *data collection* phase represents

Data collection	User-provided information	Tracking user actions	Automatic context information
User model creation	Inference-based analysis	Inference-based analysis	Collaborative Analysis
Adaptation	Only to the user	To user’s social relation	To the World Wide Web



More user control **Less User Control**

Fig. 1 Framework for privacy management in personalization. Each of the vertical categories defines a phase in the personalization process and each of horizontal sub-categories defines a level of privacy intrusiveness

privacy risks that originate from new types of data collection and analysis. For example, personalization based on location data introduces new risks that originate from new technologies for tracking location on mobile devices. The *user model creation* phase represents risks that originate from new technologies for analyzing the data. For example, technologies such as collaborative filtering and machine learning allow new predictions to be inferred from data, posing new privacy risks. Technologies and practices for exchanging data between companies increase those risks as they allow data to be traded and sent beyond the service it was collected in. Finally, risks at the *adaptation* phase originate from distributing the adapted and personalized content in new ways. For example, in traditional personalization systems, the personalized content was shown only to the user, but as we have demonstrated in the sections above, new personalization technologies distribute content to the user’s social network or even to the World Wide Web.

The second dimension of our framework is privacy control: the amount of perceived and actual control users have over their privacy in the personalization process. Technologies that limit users’ control over their privacy are perceived as more intrusive by users [Awad and Krishnan \(2006\)](#). For example, a personalization system that relies on user-provided data for the sake of personalization has the potential of being less intrusive than a system that tracks down the user’s actions automatically. The framework highlights how technologies that are based on high levels of automation pose extended risks to users’ privacy. It can explain the inherent problems of applications such as Google Buzz, which relies on automatically collected context information (e.g., her social network).

	Privacy-by-architecture		Privacy-by-policy	
Data Collection		Client-based personalization Cassel et al, 2001, Gerber et al, 2010		Do-not-track Mayer and Narayanan, 2011
User Model Creation	Pseudonymous personalization (Arlein et al, 2000; Hitchens et al, 2005, Kobsa and Schreck, 2003)	Distributed CF Miller et al (2004) encrypted aggregation Canny, 2002, Mehta, 2007 perturbation and obfuscation Polat and Du, 2003	Scrutable personalization Kay et al, 2003; Kay, 2006	Configurable User Modeling Wang and Kobsa (2007)
Adaptation				Social network privacy controls

Fig. 2 Framework for privacy mitigating solutions for different processing phases of personalization

4.2 Designing privacy-sensitive personalization systems

The technologies presented in Sect. 3 provide designers with a rich toolbox for avoiding or mitigating privacy risks. In order to efficiently use privacy preserving technologies and to guide the development of new ones, it is important to understand which type privacy risks they can mitigate. In this section, we use the framework presented in the previous section to frame privacy technologies according to the personalization process phases they apply to. The framework, depicted in Fig. 2, includes two dimensions for categorization: personalization phase and engineering approach.

Solution approaches are classified according to their focus on the personalization phases: data collection, user model creation, and adaptation. For example, the Do-Not-Track initiative (Mayer and Narayanan 2011) limits data collection while the configurable user modeling (Wang and Kobsa 2007) provides users with control over the user model creation. Some approaches impact more than one personalization phase: pseudonymous personalization is applicable to all three personalization phases it collects user data, creates user model and applies adaptation under pseudonyms. The engineering approach categorizes approaches towards engineering personalization systems, based on a framework by Spiekermann and Cranor (2009) that maps engineering practices for privacy preserving systems. The framework distinguishes between *privacy-by-policy*, which focuses on the implementation of the notice and choice principles of fair information practices, and *privacy-by-architecture*, which aspires to create systems that minimizes the collection and processing of identifiable personal data.

Privacy technologies that are relevant to the data collection phase prevent services from compiling detailed profiles of individual users by tracking their behavior. Technologies such as client-based personalization provide privacy-by-architecture solutions by preventing services from accessing user information directly. Other technologies, such as Do-Not-Track (Mayer and Narayanan 2011) and Internet Explorer Tracking Protection Lists (Microsoft 2010), enable a privacy-by-policy solution for users who wish to opt-out of tracking.

Technologies that are relevant to the model creation phase allow individual data to be hidden from central services, using technologies such as distributed collaborative filtering, or to be configured by the user using configurable user modeling. However, it is not clear that purely architectural solutions are feasible to mitigate privacy risks in the model creation phase, as models which are based on richer data might have the potential of being more accurate. For example, Knijnenburg et al. (2012) show that in recommender systems, there is a tradeoff between the users' perceived usefulness of the system and privacy concerns.

The adaptation phase pose a new challenge to designers of privacy preserving technologies, as the applications of personalized content go far beyond the traditional application of personalizing the experience for a single user. Currently, the social network's privacy settings can be used to limit access to personalized content, but these solutions are partial at best. Specifically, the question of how to control the distribution of the product of the personalization process is still an open question.

4.3 What the future holds for privacy and personalization

It is almost impossible to predict which technologies will raise new privacy challenges in the future. However, our personalization privacy risks framework provides a way to analyze the way future development might impact users' privacy, and how research can mitigate future privacy concerns. The scope of *data collection* is increasing as new technologies capture and infer more and more information. Face recognition technologies, for example, are now becoming increasingly powerful in automatically identifying people's identities in video feeds, photo collections and so forth. Face recognition can be used to infer a wide range of knowledge about people's behavior from shopping habits to social ties. As the data is gathered and analyzed automatically, the control people have on this information is highly limited. The first examples of this technology, such as Face.com, exemplify the importance of policy regulation and proper technological architecture in handling this new technology.

New data analysis technologies provide another potential for tensions between efficient personalization and privacy. The ability of massive data centers to process huge amounts of information raise new abilities to infer information about individual users. Personalization technologies, such as those operating by Google and Amazon, exemplify how individual user profiles are compared to massive sets of profile records to identify what is relevant to the user and what is not. As the predictions become more and more accurate, and as services increase their reliance on these predictions, user privacy concerns may increase. Furthermore, the results of data leakage, authorized or unauthorized, are more serious.

The adaptation phase in personalization poses new challenges as well. Online social networks and other collaborative information systems amplify existing privacy problems by communicating possibly sensitive information to peers: friends, co-workers, and family. Nowadays, the audience of the personalization process goes beyond the individual user, and can include the user's social network and other users of a collaborative system. This aspect of collaborative systems materializes privacy risks from vague notions to concrete risks associated with personalization.

5 Conclusions

This article has reviewed several privacy risks related to personalization and discussed technologies and architectures that can help designers build privacy preserving personalization systems. While no silver bullet exists for designing privacy-protective personalization systems, there are technologies and principles that can be used to eliminate, reduce, and mitigate privacy risks. Furthermore, existing approaches are not mutually exclusive and should be considered as complementary in protecting users' privacy in personalized systems. Pseudonymous profiles and aggregation can be used when personalization information need not be tied to an identifiable user profile. Client-side profiles are useful when personalization services can be performed locally. User controls should always be considered on top of other technical approaches as they will likely make the personalized system more usable and trustworthy. We envision advances in all of these areas and more systems that incorporate multiple techniques in their privacy protection mechanisms.

References

- Acquisti, A., Gross, R.: Imagined communities: Awareness, information sharing, and privacy on the facebook. In: Danezis, G., Golle, P. (eds.) *Privacy enhancing technologies, Lecture notes in computer science*, vol. 4258, pp. 36–58. Springer, Berlin (2006). doi:[10.1007/119574543](https://doi.org/10.1007/119574543)
- Anton, A.I., Earp, J.B., Young, J.D.: How internet users' privacy concerns have evolved since 2002. *IEEE Secur. Priv.* **8**(1), 21–27 (2010)
- Arlein R.M., Jai B., Jakobsson M., Monrose F., Reiter M.K.: Privacy-preserving global customization. In: 2nd ACM conference on electronic commerce, pp. 176–184. ACM Press, Minneapolis (2000)
- Awad, N.F., Krishnan, M.S.: The personalization privacy paradox: an empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS Quarterly* **30**(1), 13–28 (2006)
- Barkhuus L., Brown B., Bell M., Sherwood S., Hall M., Chalmers M.: From awareness to repartee: sharing location within social groups. In: *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pp. 497–506. ACM Press, New York (2008)
- Benisch, M., Kelley, P., Sadeh, N., Cranor, L.: Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Pers. Ubiquitous Comput.* **15**(7), 679–694 (2011)
- Beresford, A.R., Stajano, F.: Location privacy in pervasive computing. *IEEE. Pervasive Comput.* **2**(1), 46–55 (2003)
- Berkovsky S., Eytani Y., Kuflik T., Ricci F.: Privacy-enhanced collaborative filtering. In: Kobza A., Cranor L. (eds.) *PEP05, UM05 workshop on privacy-enhanced personalization*, pp. 75–84. Edinburgh (2005)
- Bonneau, J., Preibusch, S.: The privacy jungle: On the market for data protection in social networks. In: Moore, T., Pym, D., Ioannidis, C. (eds.) *Economics of information security and privacy*, pp. 121–167. Springer, New York (2009)
- Canny, J.: Collaborative filtering with privacy via factor analysis. In: *Proceedings of the 25th annual international ACM SIGIR conference on research and development in information retrieval*, pp. 238–245. ACM Press, Tampere (2002)
- Cassel, L., Cassel, L., Wolz, U.: Client side personalization. In: *Proceedings of the joint DELOS-NSF workshop on personalization and recommender systems in digital libraries*, Dublin City University, Dublin (2001)
- Ceri, S., Dolog, P., Matera, M., Nejdl, W.: Model-driven design of web applications with client-side adaptation. In: *International conference on web engineering, ICWE'04*, vol. 3140, pp. 201–214. Springer, Munich (2004)
- Coroama, V., Langheinrich, M.: Personalized vehicle insurance rates—a case for client-side personalization in ubiquitous computing. *Ubiquitous Comput. Workshop Priv. Enhanced Personal. CHI* **06**(22), 56–59 (2006)
- Cranor, L.F.: I didn't buy it for myself: privacy and ecommerce personalization. In: *2003 ACM workshop on privacy in the electronic society*, pp. 111–117 ACM Press, Washington, DC (2003)
- Culnan, M.J., Milne, G.R.: The culnan milne survey on consumers and online privacy notices: summary of responses. <http://www.ftc.gov/bcp/workshops/glb/supporting/culnan-milne.pdf> (2001). Accessed Dec 2001
- Dalal, M.: Personalized social & real-time collaborative search. In: *Proceedings of the 16th international conference on World Wide Web, WWW '07*, pp. 1285–1286. ACM Press, New York (2007)
- Davis, W.: Report: Marketers limit behavioral targeting due to privacy worries. *Tech. rep., Media Post Report* (2010)
- Dwyer, C., Hiltz, S.R., Passerini, K. Trust and privacy concern within social networking sites: a comparison of facebook and myspace. In: *Proceedings of the thirteenth americas conference on information systems (AMCIS 2007)*, Keystone (2007)
- Egelman, S., Tsai, J., Cranor, L., Acquisti, A.: Timing is everything? The effects of timing and placement of online privacy indicators. In: *Proceedings of the 27th international conference on Human factors in computing systems*, pp. 319–328. ACM Press, New York (2009)
- Facebook (2010) Facebook instant personalization. <http://blog.facebook.com/blog.php?post=384733792130>. Accessed 26 April 2010
- Federal Trade Commission: self-regulatory principles for online behavioral advertising. *Tech. rep., Federal Trade Commission* (2009)
- Forrester Consulting: Outbound email and dataloss prevention in today's enterprise. <http://www.proofpoint.com/outbound> (2008). Accessed March 2008

- GEDIK, B., LIU, L.: Protecting location privacy with personalized k-anonymity: architecture and algorithms. *IEEE Trans. Mobile Comput.* **7**(1), 1–18 (2008)
- GERBER, S., FRY, M., KAY, J., KUMMERFELD, B., PINK, G., WASINGER, R. PersonisJ: mobile, Client-Side user modelling. In: International conference on user modeling, adaptation, and personalization, lecture notes in computer science, vol. 6075, pp. 111–122. Springer, Berlin (2010)
- Google (2008) Social graph API. <http://code.google.com/apis/socialgraph/>. Accessed 01 Feb 2008
- Google (2011) Google +1 button - social search. <http://www.google.com/+1/button/>. Accessed 01 June 2011
- GRUTESER, M., LIU, X.: Protecting privacy, in continuous location-tracking applications. *Secur. Priv. IEEE.* **2**(2), 28–34 (2004)
- HELFIT, M., WORTHAM, J.: Facebook bows to pressure over privacy. *New York Times*, New York (2010)
- HEYMANN, P., KOUTRIKA, G., GARCIA-MOLINA: Can social bookmarking improve web search? In: Proceedings of the international conference on web search and web data mining, WSDM '08, pp. 195–206. New York (2008)
- HITCHENS, M., KAY, J., KUMMERFELD, B., BRAR, A. Secure identity management for pseudo-anonymous service access. In: Hutter, D., Ullmann, M. (eds.) Security in pervasive computing: second international conference, pp. 48–55, Boppard (2005)
- HOH, B., GRUTESER, M.: Protecting location privacy through path confusion. In: Security and privacy for emerging areas in communications networks, 2005. SecureComm 2005. First international conference on security and privacy for emerging areas in communications networks, pp. 194–205. IEEE Computer Society, Washington (2005)
- IACHELLO, G., SMITH, I., CONSOLOVO, S., ABOWD, G., HUGHES, J., HOWARD, J., POTTER, F., SCOTT, J., SOHN, T., HIGHTOWER, J., LAMARCA, A.: Control, deception, and communication: Evaluating the deployment of a location-enhanced messaging service. In: Ubicomp '05, pp. 213–231. Springer-Verlag, Berlin (2005)
- KAY, J.: Scrutable adaptation: because we can and must. In: Adaptive hypermedia and adaptive web-based systems, pp. 11–19. Springer, Berlin (2006)
- KAY, J., KUMMERFELD, B., LAUDER, P. Managing private user models and shared personas. In: Workshop on user modelling for ubiquitous computing, 9th international conference on user modeling, pp. 1–11. Johnstown (2003)
- KELLEY, P.G., DRIELISMA, P.H., SADEH, N., CRANOR, L.F.: User-controllable learning of security and privacy policies. In: Proceedings of the 1st ACM workshop on AISec, pp. 11–18. ACM Press, Alexandria (2008)
- KNIJNENBURG, B.P., WILLEMSSEN, M.C., GANTNER, Z., SONCU, H., NEWELL, C.: Explaining the user experience of recommender systems. *User Model. User Adapt. Interact.* **22** (2012). doi:10.1007/s11257-011-9118-4
- KOBSA, A.: Generic user modeling systems. *User Model. User Adapt. Interact.* **11**(1–2), 49–63 (2001)
- KOBSA, A.: Generic user modeling systems. In: Brusilovsky, P., Kobsa, A., Nejdil, W. The adaptive web: methods and strategies of web personalization., pp. 136–154. Springer Verlag, Heidelberg (2007a)
- KOBSA, A.: Privacy-enhanced web personalization. In: Brusilovsky, P., Kobsa, A., Nejdil, W. The Adaptive Web., pp. 628–670. Springer-Verlag, Berlin (2007b)
- KOBSA, A., SCHRECK, J.: Privacy through pseudonymity in User-Adaptive systems. *ACM. Trans. Internet Technol.* **3**(2), 149–183 (2003)
- KOBSA, A., KOENEMANN, J., POHL, W.: Personalized hypermedia presentation techniques for improving online customer relationships. *Knowl. Eng. Rev.* **16**, 111–155 (2001)
- LAMPINEN, A., TAMMINEN, S., OULASVIRTA, A.: All my people right here, right now: management of group co-presence on a social networking site. In: GROUP '09: Proceedings of the ACM 2009 international conference on supporting group work, pp. 281–290. ACM Press, New York (2009)
- LEWIS, K., KAUFMAN, J., CHRISTAKIS, N.: The taste for privacy: an analysis of college student privacy settings in an online social network. *J. Comput. Mediat. Commun.* **14**(1), 79–100 (2008)
- LU, Y., PENG, F., WEI, X., DUMOULIN, B.: Personalize web search results with user's location. In: Proceeding of the 33rd international acm sigir conference on research and development in information retrieval, SIGIR 2010, pp. 763–764, Geneva (2010)
- MALIN, B., SWEENEY, L., NEWTON, E.: Trail re-identification: learning who you are from where you have been. Technical Report LIDAP-WP12, Carnegie Mellon University, Laboratory for international data privacy (2003)
- MAYER, J.R., NARAYANAN, A.: Do not track iab/w3c/ietf position paper. Tech. rep., W3C (2011)
- MCDONALD, A.M., CRANOR, L.F.: Beliefs and behaviors: Internet users' understanding of behavioral advertising. Tech. rep. Carnegie Mellon University, Pittsburgh (2010)

- Mehta, B.: Cross system personalization by learning manifold alignments. In: Proceedings of the 21st national conference on artificial intelligence, Vol. 2, pp. 1920–1921. AAAI Press, Menlo Park (2006)
- Mehta, B.: Learning from what others know: privacy preserving cross system personalization. In: Proceedings of the 11th international conference on user modeling, UM '07, pp. 57–66. Springer-Verlag, Berlin (2007)
- Mehta, B., Niederee, C., Stewart, A., Degenmis, M., Lops, P., Semeraro, G.: Ontologically-enriched unified user modeling for cross-system personalization. In: Ardissono L., Brna P., Mitrovic A. (eds.) User Modeling 2005, Lecture notes in computer science, vol. 3538, pp. 119–123. Springer, Berlin (2005)
- Microsoft Internet explorer tracking protection lists. <http://ie.microsoft.com/testdrive/Browser/TrackingProtectionLists> (2010). Accessed Sept 2010
- Miller, B.N., Konstan, J.A., Riedl, J.: PocketLens: toward a personal recommender system. ACM. Trans. Inf. Syst. **22**(3), 437–476 (2004)
- Mulligan, D., Schwartz, A.: Your place or mine?: privacy concerns and solutions for server and client-side storage of personal information. In: Proceedings of the tenth conference on computers, Freedom and privacy: challenging the assumptions, pp. 81–84. ACM Press, Toronto (2000)
- Mullins, R.: VentureBeat report: privacy group argues buzz breaks wiretap laws. <http://venturebeat.com/2010/02/17/privacy-group-argues-buzz-breaks-wiretap-laws/> (2010). Accessed 17 Feb 2010
- Nakashima, E.: AOL search queries open window onto users' worlds. Washington Post (2006)
- Nissenbaum, H.: Privacy as contextual integrity. Wash. Law Rev. Assoc. **79**, 119–158 (2004)
- Palen, L., Dourish, P.: Unpacking “privacy” for a networked world. In: Proceedings of the SIGCHI conference on human factors in computing systems (CHI '03), pp. 129–136. ACM Press, New York (2003)
- Paliouras, G.: Discovery of web user communities and their role in personalization. User Model. User Adapt. Interact. **22**(1–2), 151–175 (2012)
- Polat, H., Du, W.: Privacy-preserving collaborative filtering using randomized perturbation techniques. In: IEEE international conference on data mining (ICDM'03). IEEE Computer Society, Los Alamitos (2003)
- Polat, H., Du, W.: SVD-based collaborative filtering with privacy. In: 20th ACM symposium on applied computing, pp. 791–795. Santa Fe (2005)
- Riedl, J.: Personalization and privacy. Internet Comput. IEEE. **5**(6), 29–31 (2001)
- Ristenpart, T., Maganis, G., Krishnamurthy, A., Kohno, T.: Privacy-preserving location tracking of lost or stolen devices: cryptographic techniques and replacing trusted third parties with DHTs. In: Proceedings of the 17th conference on security symposium, pp. 275–290. USENIX Association, San Jose (2008)
- Schafer, J., Frankowski, D., Herlocker, J., Sen, S.: Collaborative filtering recommender systems. In: Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.) The Adaptive Web, pp. 291–324. Springer-Verlag, Berlin (2007)
- Spiekermann, S., Cranor, L.F.: Engineering privacy. IEEE. Trans. Softw. Eng. **35**(1), 67–82 (2009)
- Story, L., Stone, B.: Facebook retreats on online tracking. New York Times, New York (2007)
- Stutzman, F., Kramer-Duffield, J.: Friends only: examining a privacy-enhancing behavior in facebook. In: Mynatt ED, Schoner D, Fitzpatrick G, Hudson SE, Edwards K, Rodden T (eds.) CHI, pp. 1553–1562. ACM, New York (2010)
- Sweeney, L.: K-anonymity: a model for protecting privacy. Int. J. Uncertain. Fuzziness Knowl Based Syst. **10**(5), 557–570 (2002)
- Tang, K.P., Keyani, P., Fogarty, J., Hong, J.I.: Putting people in their place: an anonymous and privacy-sensitive approach to collecting sensed data in location-based applications. In: Proceedings of the SIGCHI conference on human factors in computing systems, pp. 93–102. ACM Press, Montréal (2006)
- Toch, E., Cranshaw, J., Drielsma, P.H., Tsai, J.Y., Kelley, P.G., Springfield, J., Cranor, L., Hong, J., Sadeh, N.: Empirical models of privacy in location sharing. In: Proceedings of the 12th ACM international conference on Ubiquitous computing, UbiComp '10, pp. 129–138. ACM Press, New York (2010)
- TRUSTe, TNS: 2009 study: consumer attitudes about behavioral targeting. Tech. rep., TRUSTe (2009)
- Tsai, J., Kelley, P.G., Cranor, L.F., Sadeh, N.: Location-sharing technologies: Privacy risks and controls. J. Law Policy Inf. Soc. **6**(2), 119–151 (2010)
- Tsai, J.Y., Egelman, S., Cranor, L., Acquisti, A.: The effect of online privacy information on purchasing behavior: an experimental study. Inf. Syst. Res. **22**, 254–268 (2011)
- Turov, J., King, J., Hoofnagle, C.J., Bleakley, A., Hennessy, M.: Americans reject tailored advertising and three activities that enable it. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1478214 (2009). Accessed 29 Sept 2009

- Volokh, E.: Personalization and privacy. *Commun ACM* **43**, 84–88 (2000)
- Wang, Y.: A framework for Privacy-Enhanced personalization. Ph.D. Dissertation, University of California, Irvine (2010)
- Wang, Y., Kobsa, A.: Impacts of privacy laws and regulations on personalized systems. In: Kobsa, A., Chellappa, R.K., Spiekermann, S. (eds.) PEP06, CHI06 workshop on privacy-enhanced personalization, pp. 44–46. Montréal (2006)
- Wang, Y., Kobsa, A.: Respecting users' individual privacy constraints in web personalization. In: Conati, C., McCoy, K., Paliouras, G. (eds.) UM07, 11th international conference on user modeling, Berlin–Heidelberg–New York, pp. 157–166. Springer-Verlag, Corfu (2007)
- Wang, Y., Norcie, G., Cranor, L.F.: Who is concerned about what? a study of american, chinese and indian users' privacy concerns on social networking sites. In: 4th international conference on trust and trust-worthy computing (TRUST2011), Springer, Pittsburgh 2011
- Yi, X., Raghavan, H., Leggetter, C.: Discovering users' specific geo intention in web search. In: WWW '09: Proceedings of the 18th international conference on World wide web, pp. 481–490. ACM Press, New York 2009

Author Biographies

Eran Toch is a faculty member at Tel Aviv University, where he is a lecturer at the Department of Industrial Engineering. His research interest lies at the intersection of several fields: privacy, human-computer interaction and artificial intelligence. Previously, he was a post-doctoral fellow at the School of Computer Science at Carnegie Mellon University. Dr. Toch received his Ph.D. from The Technion—Israel Institute of Technology.

Yang Wang is Research Scientist in CyLab at Carnegie Mellon University, working in the area of usable privacy and security, personalization, and social computing. He received his M.S. and Ph.D. in Information and Computer Science from the University of California, Irvine. His contribution to this article is based on experiences gained both from his Ph.D. work and his current research on privacy-enhancing personalization.

Lorrie Faith Cranor is an Associate Professor of Computer Science and of Engineering and Public Policy at Carnegie Mellon University where she is director of the CyLab Usable Privacy and Security Laboratory (CUPS). She is also a co-founder of Wombat Security Technologies, Inc. She was previously a researcher at AT&T-Labs Research and taught in the Stern School of Business at New York University. Dr. Cranor received her doctorate degree in Engineering and Policy from Washington University in St. Louis.