

Personal Privacy through Understanding and Action: Five Pitfalls for Designers

SCOTT LEDERER (✉)

Group for User Interface Research, Computer Science Division, University of California, Berkeley, CA, USA

lederer@cs.berkeley.edu
510-643-0943
510-642-5615

JASON I. HONG

Group for User Interface Research, Computer Science Division, University of California, Berkeley, CA, USA

jasonh@cs.berkeley.edu

ANIND K. DEY

Intel Research, Berkeley, CA, USA

Group for User Interface Research, Computer Science Division, University of California, Berkeley, CA, USA

anind@intel-research.net

JAMES A. LANDAY

DUB Group, Department of Computer Science and Engineering, University of Washington, Seattle, WA, USA

Intel Research, Seattle, WA, USA

landay@cs.washington.edu

Abstract. To participate in meaningful privacy practice in the context of technical systems, people require opportunities to *understand* the extent of the systems' alignment with relevant practice and to conduct discernible social *action* through intuitive or sensible engagement with the system. It is a significant challenge to design for such understanding and action through the feedback and control mechanisms of today's devices. To help designers meet this challenge, we describe five pitfalls to beware when designing interactive systems—on or off the desktop—with personal privacy implications. These pitfalls are: obscuring potential information flow, obscuring actual information flow, emphasizing configuration over action, lacking coarse-grained control, and inhibiting existing practice. They are based on a review of the literature, on analyses of existing privacy-affecting systems, and on our own experiences designing a prototypical user interface for managing privacy in ubiquitous computing. We illustrate how some existing research and commercial systems—our prototype included—fall into these pitfalls and how some avoid

them. We suggest that privacy-affecting systems that heed these pitfalls can help users appropriate and engage them in alignment with relevant privacy practice.

Keywords: *privacy, interaction design, design guidelines, ubiquitous computing*

Introduction

One possible reason why designing privacy-sensitive systems is so difficult is that, by refusing to render its meaning plain and knowable, privacy simply lives up to its name. Rather than exposing an unambiguous public representation for all to see and comprehend, it cloaks itself behind an assortment of meanings, presenting different interpretations to different people. When sociologists look at privacy, they see social nuance that engineers overlook. When cryptologists consider privacy, they see technical mechanisms that everyday people ignore. When the European Union looks at privacy, it sees moral expectations that American policymakers do not. Amidst this fog of heterogeneous practices, technologies, and policies that characterize the current state of privacy, designers of interactive systems face increasing market pressure and a persistent moral imperative to design systems that support users' privacy needs: systems that are *privacy-sensitive*.¹

This article cannot dispel that fog, but it does attempt to shine some light through it by offering a partial set of guidelines for designers of privacy-affecting interactive systems, on and off the desktop. We say *partial* set because this article does not aspire to be a self-contained how-to guide. We do not intend that systems that follow these guidelines will decidedly support privacy. We *do* intend that systems that *ignore* any of these guidelines without careful rationale face significant risk of disrupting or inhibiting users' abilities to manage their personal privacy. For this reason, we present our guidelines as a set of *pitfalls* to avoid when designing privacy-affecting systems. Avoiding a pitfall does not ensure success, but ignoring one can potentially lead to disaster.

¹ We will use the term *privacy-affecting* as a general description for any interactive system whose use has personal privacy implications. We will use the term *privacy-sensitive* to describe any privacy-affecting system that—by whatever metrics are contextually relevant—reasonably avoids invading or disrupting personal privacy. This article is intended to help minimize the number of privacy-affecting systems that are *not* privacy-sensitive.

In addition to using our guidelines, designers of privacy-affecting *ubiquitous computing* systems should consult Bellotti and Sellen's framework for feedback and control [1], Langheinrich's transliteration of the fair information practices [2], Palen and Dourish's sociologically informed analysis of privacy as boundary negotiation [3], and Jiang *et al.*'s principle of minimum asymmetry [4]. Our work synthesizes some of the core lessons of those frameworks to inform an analysis of common privacy problems across a broad range of existing systems.

Common Design Flaws

There has been a tremendous amount of research on privacy in the context of technical systems. This includes polls showing considerable public concern about privacy on the Internet [5-7]; interviews and surveys exploring privacy design issues for context-aware systems [8-10]; studies exposing privacy perceptions and practices in groupware [11], multimedia environments [12], and location-aware systems [13]; and experiments revealing usability problems affecting privacy in email [14] and file-sharing [15] applications. Despite the consequent abundance of research and design knowledge, many systems still make it hard for people to manage their privacy.

We suggest this is largely because the designs of these systems inhibit peoples' abilities to both *understand* the privacy implications of their use and to conduct socially meaningful *action* through them. We further suggest that designs that address our five pitfalls will go a long way towards helping people achieve the understanding and action that personal privacy regulation requires. Although some of these pitfalls may appear obvious, we will demonstrate below that many systems continue to stumble into them. Some of these systems have encountered privacy controversies (*e.g.*, web browsers), while others that have avoided the pitfalls have enjoyed considerable commercial or social success (*e.g.*, instant messaging).

Our investigation into these pitfalls began when we fell into them ourselves in the design of a user interface prototype for managing personal privacy in ubiquitous computing environments [16]. Despite the input of our formative interviews, surveys, and literature review, an evaluation indicated some fundamental missteps in our design rationale. Further analysis showed that these missteps were not exclusive to our system;

we found similar problems in a number of existing commercial and research systems. Without attempting to enumerate every extant privacy design flaw, we can offer the design community descriptions of these common ones and a warning to heed them.

To help designers remember these pitfalls, we have clustered them into those that primarily affect users' *understanding* of a system's privacy implications and those that primarily affect their ability to conduct socially meaningful *action* through the system.

Understanding

1. *Obscuring potential information flow.* Designs should not obscure the nature and extent of a system's *potential* for disclosure. Users can make informed use of a system only when they understand the scope of its privacy implications.
2. *Obscuring actual information flow.* Designs should not conceal the *actual* disclosure of information through a system. Users should understand what information is being disclosed to whom.

Action

3. *Emphasizing configuration over action.* Designs should not require excessive configuration to manage privacy. They should enable users to practice privacy as a natural consequence of their normal engagement with the system.
4. *Lacking coarse-grained control.* Designs should not forgo an obvious, top-level mechanism for halting and resuming disclosure.
5. *Inhibiting established practice.* Designs should not inhibit users from transferring established social practice to emerging technologies.

Before further articulating and providing evidence supporting these suggestions, we will elaborate on the meaning of our title: Personal Privacy through Understanding and Action.

Personal Privacy

Legal and policy scholar Alan F. Westin asserted that “no definition of privacy is possible, because privacy issues are fundamentally matters of values, interests and

power” [17] (as quoted in [18]). We will not be so bold as to define privacy, but we will attempt to qualify within the scope of this article the phrase *personal privacy*.

Years prior to the assertion quoted above, Westin described information privacy as “the claim of individuals, groups or institutions to determine for themselves when, how and to what extent information about them is communicated to others” [19]. Largely intended for policymakers, the reasoning behind this formulation served as the basis for the *fair information practices*, a set of flexible policy guidelines that continue to shape privacy legislation throughout the world. Since many privacy-affecting interactive systems are developed or deployed by organizations beholden to some interpretation of the fair information practices, Westin’s formulation is a good place to start when elucidating personal privacy to designers. But we cannot end there, for there is more to privacy than this rather deterministic, libertarian formulation conveys.

Building on the work of social psychologist Irwin Altman [20], Palen and Dourish offer a more organic, sociologically-informed view that “privacy management in everyday life involves combinations of social and technical arrangements that reflect, reproduce and engender social expectations, guide the interpretability of action, and evolve as both technologies and social practices change” [3]. In this sense, privacy is less about a definitive entitlement to determine the flow of one’s personal information and more about the intuitive fulfillment and maintenance of one’s compound roles in the evolving, overlapping socio-technical contexts of everyday life.

While neither formulation excludes the other, one might say—at the risk of oversimplification—that Westin’s formulation emphasizes privacy as *conscious process*, while Palen and Dourish’s and Altman’s emphasizes privacy as *intuitive practice*. Clearly, however, people regulate their privacy in ways both deliberate and intuitive. Drawing directly from each formulation, then, what we are trying to signify by the phrase *personal privacy* is this set of both deliberate and intuitive practices by which an individual exercises her claim to determine personal information disclosure and which constitute, in part, her participation in the co-evolving technologies and expectations of everyday life.

A useful term that can make this discussion more concrete is Palen and Dourish’s *genres of disclosure*, which are “socially-constructed patterns of privacy management,”

or “regularly reproduced arrangements of people, technology and practice that yield identifiable and socially meaningful styles of interaction, information, *etc.*” [3]. Examples might include creating and managing accounts at shopping websites; taking (or not, as the genre may oblige) photographs at social events; exchanging contact information with a new acquaintance; and the extent, nature, and accuracy of the personal history one reveals to strangers. These all involve recognizable, socially meaningful patterns of information disclosure and use. A genre of disclosure enforces social expectations of its participants. Amidst a given genre, people expect each other to disclose *this* information but not *that*, under *these* conditions but not *those*, to *this* but not *that* person, and to use information in *this* but not *that* way. A person cooperates with (or antagonizes) a genre of disclosure through her performance of her expected role in that genre, and the degree to which a system does *not* align with that genre is the degree to which it fails to support the user’s (and genre’s) privacy regulation process. In this sense, what we call personal privacy in this article is the individual’s *expected* performance within a given genre of disclosure, her *actual* performance, and the intentions and actions which determine the difference between them.

We note this difference because personal privacy can also include acting *contrary* to expectation. As technologies evolve, so do the practices that involve them. New modes and expectations of disclosure emerge as people both embrace *and* resist technologies and practices. Regardless of the case, a person can neither fully participate in nor effectively defy a genre of disclosure without *understanding* whether the system at hand aligns with that genre and without the ability to *act* in—or out of—alignment with it.

Understanding and Action

To be clear, we do not intend this dyadic formulation of *understanding* and *action* as a contribution to the theory of privacy, but simply as a conceptual framework for the arguments in this article. We frame our arguments using these terms in the hope of reaching as broad an audience as possible, for the sooner that designs improve their ability to support personal privacy regulation, the better.

With respect to genres of disclosure, we are proposing that a person cannot fulfill her role in the apposite genre of disclosure if she does not *understand* the degree to which

the system at hand aligns with that genre and if she cannot conduct socially interpretable *action* involving the system. We suggest that a system that falls into any of our pitfalls without due rationale can disrupt its users' abilities to appropriate it in accordance with the relevant genre of disclosure. In so doing, it would disrupt the genre itself or—if it is an emerging genre—make it unnecessarily complex. A privacy-sensitive interactive system will sustain the appropriate genre of disclosure—and will help its users do the same—through cues, affordances, and functions that empower users to comprehend and influence their privacy implications.

Empowering understanding and action is similar in meaning to bridging Norman's gulfs of evaluation and execution [21]. We feel the terms we have chosen convey a richer sense of the *social* implications of privacy-affecting systems than do Norman's terms, which seem to best address the perceptual/cognitive/motor problem of single-user human-system interaction. Privacy regulation does not conform to a plan-act-evaluate cycle; it is a continual, intuitive, multidimensional balancing act that requires nonlinear social dexterity. That said, at the end of this article we will examine another of Norman's canonical contributions—his elucidation of the role of mental models in the design process—and extend it to accommodate the social dimension of the privacy design process.

The rest of this article is organized as follows. First, we discuss the design and evaluation of *Faces*, our UI prototype for managing personal privacy in ubiquitous computing settings. The negative results of the evaluation motivated our investigation into the design missteps encoded in our five pitfalls. We then describe the five pitfalls, with illustrative examples from both our own and related work. We then discuss the pitfalls' implications for the design process, including an extension of Norman's elucidation of mental models. Finally, we offer negative and positive case studies of systems that, respectively, fall into and avoid the pitfalls.

Faces: (Mis)Managing Ubicomp Privacy

Our investigation into the pitfalls began after we encountered them firsthand while designing *Faces*, a software prototype for specifying privacy preferences in ubiquitous computing (ubicomp) environments. Ubicomp envisions computation

embedded throughout everyday environments to support arbitrary human activities [22], but by distributing and concealing displays and sensors, it complicates interaction [23]. This can disadvantage users by leaving them unaware of or unable to influence the disclosure of personal information—such as location and identity—as they go about their activities in augmented environments. To address this, we designed Faces to (1) support the specification of disclosure preferences, such as *who* can obtain *what* information *when* (Figure 1), and (2) provide feedback about past disclosures in an accessible log, not unlike the financial transaction logs in Quicken or Microsoft Money (Figure 2). Users would employ the feedback in the log to iteratively refine their disclosure preferences over time.²

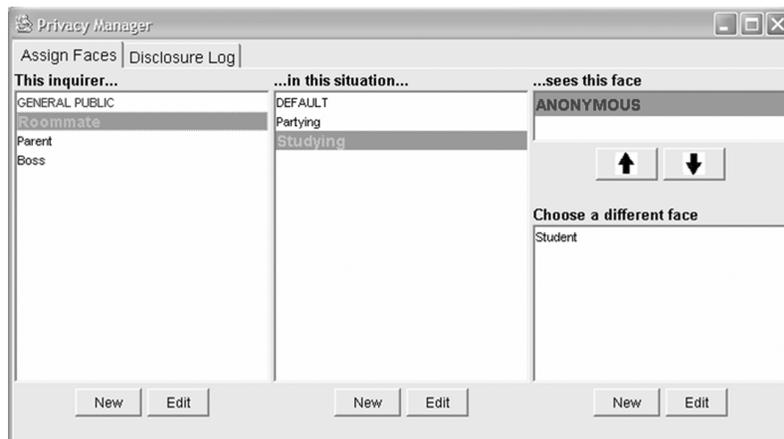


Figure 1. GUI for creating and assigning faces. Each face holds information precision preferences for disclosures to the associated

² Some might object here, noting that informing the user about a disagreeable disclosure *after the fact* is too late to be useful. While this may apply to highly sensitive disclosures, a significant component of privacy maintenance is the regulation of mundane disclosures *over time* to influence observers' historical, evolving impressions of one's self. People are remarkably capable of finessing the consequences of the occasional—and inevitable—disagreeable disclosure, and they learn to minimize repeat occurrences. The Faces disclosure log was intended to help users transfer such iterative behavior refinement to the domain of the sensed environment.

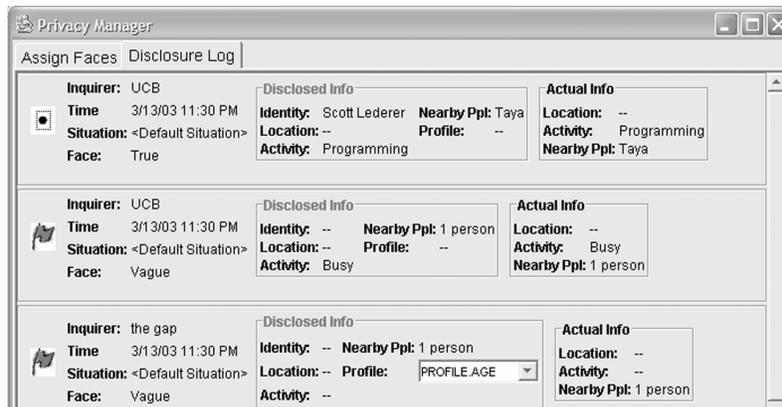


Figure 2. Users could ascertain the characteristics of disagreeable disclosures from the disclosure log, and refine their preferences to limit similar disclosures.

The next section will show that the design of Faces involved some crucial missteps that are also present in other systems. What clued us in to the fundamental nature of these missteps is that we made them despite a substantive requirements gathering effort (details in [16]). We reviewed the literature. We interviewed twelve local residents solicited from a public community website, walking them through a series of scenarios to elicit how they might think about privacy in ubicomp. We surveyed 130 people on the web to investigate factors that determine privacy preferences in ubicomp [10]. And we iterated through a series of low-fidelity designs. The functional upshot of our findings was that the identity of the inquirer is a primary determinant of users' privacy preferences, but the situation in which the information is disclosed is also important.

Accordingly, we designed Faces to let users assign different disclosure preferences to different *inquirers*, optionally parameterized by *situation* (a conjunction of location, activity, time, and nearby people). We employed the metaphor of *faces* to represent disclosure preferences. This is a fairly direct operationalization of Goffman, who posited that a person works to present himself to an audience in such a way as to maintain a consistent impression of his role in relation to that audience [24]—to maintain the appropriate face. Prior to any affected disclosures, users employ a desktop application to specify their preferences for subsequent disclosures by creating 3-tuples of *inquirers*, *situations*, and *faces*, with each 3-tuple meaning “if *this* inquirer wants information about me when I’m in *this* situation, show her *this* face” (Figure 3). Wildcards are allowed in

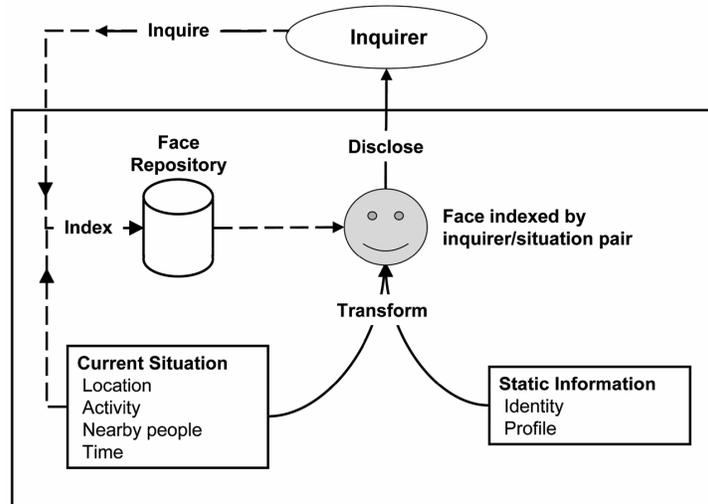


Figure 3. Disclosure precision preferences—encapsulated in *faces*—are indexed on a per inquiry basis, according to the inquirer’s identity and the user’s situation at the time of inquiry. the inquirer and situation slots to handle requests from unregistered inquirers (General Public) or when the user’s conditions do not meet the parameters of any registered situations (Default Situation). The preferences established in the desktop module are automatically synchronized with a handheld module that affords *in situ* feedback and control (Figure 4) and that we envisioned would communicate the user’s preferences to nearby ubicomp systems in the manner of Langheinrich’s Privacy Awareness System [25].

Each face alters the disclosed information by specifying the *precision* at which to disclose it. Faces supports four ordinal levels of precision—from *Undisclosed* (disclose nothing) through *Vague* through *Approximate* to *Precise* (disclose everything). Each face lets the user apply a setting from this scale to each of four information dimensions: identity, location, activity, and nearby people (Figure 5). Adjusting the precision of information can desensitize it, allowing for different versions of the same information to reach different inquirers, depending on the situation [10]. For example, a woman might permit her spouse to employ a locator system to determine that she is at her physician’s office (precise), but she might prefer that inquisitive friends learn only that she is downtown (vague).

Through its emphasis on inquirers, situations, and precision preferences, Faces operationalizes three of Adams and Sasse’s four factors that determine the perception of



Figure 4. The main screen of the handheld module allows for *in situ* feedback and control. Users could quickly override active preferences and save a snapshot of current contextual variables (e.g., location, time) for subsequent use as a situation parameter. Nested menus offer deeper configuration options.

privacy in richly sensed environments: recipient, context, and sensitivity [26]. We did not directly address the fourth factor—usage—because Faces emphasizes *a priori* preference provisioning and yet it is often impractical to predict how an observer will use observed information [1].

Formative Evaluation

A formative evaluation revealed fundamental problems with the Faces concept (details in [16]).³ After a thorough introduction and tutorial, five participants used the system to configure their privacy preferences regarding two inquirers and two situations of their choice. That is, they each created two inquirer entities in the Faces user interface to represent two parties whom they felt would regularly be interested in their location, activity, *etc.*, followed by two situation entities representing situations they often find themselves in, followed by a set of faces encoding the precision preferences they felt

³ Flaws in the visual and surface-level interaction design of the software also contributed to negative evaluation results. However, we have been careful to focus our interviews with participants and our resulting analysis on problems rooted in the conceptual model behind the interaction design—problems which even optimal interaction and visual design could not sufficiently overcome.



Figure 5. Each face contains disclosure preferences for identity, location, activity, and nearby people.

comfortable applying to disclosures to those inquirers in those situations. At a minimum, this means they would create a single face to handle both inquirers in both situations; at a maximum they would create four unique faces, one for each combination of the two inquirers and the two situations. We then described a series of hypothetical but realistic scenarios involving those same inquirers and situations and asked the participants to consider and state the precision levels at which they would prefer to disclose their information to those inquirers in those scenarios.⁴

Results showed that participants' *a priori* configured preferences often differed pointedly from their stated preferences during the scenarios. That is, when confronted with a realistic description of a specific scenario, participant's disclosure preferences

⁴ By *scenario* we mean a specific activity in a specific context (e.g., buying a pint of chocolate ice cream at the grocery store on Main Street at ten o'clock on a Saturday night). We chose our *scenarios* to be specific, somewhat sensitive events that met the constraints of the more general *situations* created in the Faces UI (e.g., shopping during the weekend).

differed from what they had previously thought they would be. Further, they had difficulty remembering the precision preferences they had specified inside their faces. This clouded their ability to predict the characteristics of any given disclosure: they might remember the *name* of the face that would be indexed by the characteristics of a given disclosure, but they would be hard pressed to recall exactly how that face would affect the disclosure.

Subsequent interviews with the participants corroborated these results and also brought the faces metaphor into significant question. Participants expressed discomfort with the indirection between faces and the situations in which they apply. In their minds, a situation and the face one “wears” in it are inseparable; they are, for practical purposes, the same thing.

Together these results illustrate the misstep of separating the privacy management process from the contexts in which it applies. While Faces *modeled* Goffman’s theory *in* the interface, it inhibited users from *practicing* identity management *through* the interface. Users had to think explicitly about privacy in the abstract—and instruct the system to model an external representation of their privacy practices—instead of managing privacy intuitively through their actions *in situ* [3].

Having identified these design flaws despite a reasonable design process, we reviewed other privacy-affecting systems in search of similar mistakes. The practicable outcome of this analysis is our set of pitfalls to beware when designing for personal privacy, presented below with evidence of designs both succumbing to and avoiding them. After articulating them, we will analyze the Faces system—and three others—with respect to the five pitfalls.

Five Pitfalls to Heed When Designing for Privacy

Our pitfalls encode common problems in interaction design across several systems, constituting a preventative guide to help designers avoid mistakes that may appear obvious in retrospect but continue to be made nonetheless. We encourage designers to carefully heed the pitfalls throughout the design cycle. Naturally, they will apply in different ways and to different degrees for each system. They should be interpreted within the context of the design task at hand.

The pitfalls fit into a history of analyses and guidelines on developing privacy-sensitive systems. They are, in part, an effort to reconcile Palen and Dourish’s theoretical insights about how people maintain privacy with Bellotti and Sellen’s practical guidelines for designing feedback and control to support it. In reaching for this middle ground, we have tried to honor the fair information practices—as developed by Westin [19] and more recently adapted to the ubicomp design space by Langheinrich [2]—and to encourage minimum information asymmetry between subjects and observers—as argued by Jiang *et al.* [4].

Concerning Understanding

Our first two pitfalls involve the user’s *understanding* of a system’s privacy implications. By illuminating (1) the system’s potential for information disclosure and (2) the actual disclosures made through it, a system can fortify users’ comprehension of its scope, its utility, and the implications of its use.

Pitfall 1: Obscuring Potential Information Flow

To whatever degree is reasonable, systems should make clear the nature and extent of their *potential* for disclosure. Users will have difficulty appropriating a system into their privacy practice if the scope of its privacy implications is unclear. This scope includes the *types* of information the system conveys, the *kinds* of observers it conveys to, the media through which it is conveyed, the length of retention, the potential for unintentional disclosure, the presence of third-party observers, and the collection of meta-information like traffic analysis. Clarifying a system’s potential for conveying personal information is vital to users’ ability to predict the social consequences of its use.

Among the conveyable information types to elucidate are identifiable *personae* (*e.g.*, true names, login names, email addresses, credit card numbers, social security numbers) and monitorable *activities* (broadly, any of the user’s interpretable actions and/or the contexts in which they are performed, *e.g.*, locations, purchases, clickstreams, social relations, correspondences, audio/video records). This dichotomy of *personae* and *activities*, though imperfect and coarse, can be useful shorthand for conceptualizing a user’s identity space, with *personae* serving as indices to dynamically intersecting

subspaces and activities serving as the contents of those subspaces [27]. People work to maintain consistency of character with respect to a given audience, in effect ensuring that an audience cannot access an identity subspace to which it does not already have an index. This can require considerable effort because boundaries between subspaces are fluid and overlapping. Conveying evidence of activity out of character with the apposite persona can rupture the carefully maintained boundaries between identity subspaces, collapsing one's fragmented identities and creating opportunities for social, bodily, emotional, and financial harm [28].

Privacy-affecting systems tend to involve disclosure both between people (interpersonal) and between an individual and an organization (organizational). Designs should address the potential involvement of each, clarifying if and how primarily interpersonal disclosures (*e.g.*, chat) involve incidental organizational disclosures (*e.g.*, workplace chat monitoring) and, conversely, if and how primarily organizational disclosures (*e.g.*, workplace cameras) involve secondary interpersonal disclosures (*e.g.*, mediaspaces).

“Privacy” is a broad term whose unqualified use as a descriptor can mislead users into thinking a system protects or erodes privacy in ways it does not. Making the scope of a system's privacy implications clear will help users understand its capabilities and limits. This in turn provides grounding for comprehending the *actual* flow of information through the system, addressed in the next pitfall.

Evidence: Falling into the Pitfall

An easy way to obscure a system's privacy scope is to present its functionality ambiguously. One example is Microsoft's Windows operating systems, whose Internet control panel offers ordinal degrees of privacy protection (from Low to High). First, the functional meaning of this scale is unclear to average users. Second, despite being a component of the OS's control panel, this mechanism does not control general privacy for general Internet use through the operating system; its scope is limited only to a particular web browser's cookie management heuristics.

Similarly, Anonymizer.com's free anonymizing software can give the impression that all Internet activity is anonymous when the service is active, but in actuality it only

affects web browsing, not email, chat, or other services. A for-pay version covers those services.

Another example is found in Beckwith's report of an eldercare facility that uses worn transponder badges to monitor the locations of residents and staff [13]. Many residents perceived the badge only as a call-button (which it was) but not as a persistent location tracker (which it also was). They did not understand the disclosures it was capable of facilitating.

Similarly, some hospitals use badges to track the location of nurses for efficiency and accountability purposes but neglect to clarify what kind of information the system conveys. Erroneously thinking the device was also a microphone, one concerned nurse wrote, "They've placed it in the nurses' lounge and kitchen. Somebody can click it on and listen to the conversation. You don't need a Big Brother overlooking your shoulder" [29].

A recent example of a privacy-affecting system that has given ambiguous impressions of its privacy implications is Google's Gmail email system. Gmail's content-triggered advertisements have inspired public condemnation and legal action over claims of invading users' privacy [30]. Some critics may believe that Google discloses email content to advertisers—which Gmail's architecture prohibits—while some may simply protest the commercial exploitation—automated or not—of the content of personal communications. Despite publishing a conspicuous and concise declaration on Gmail's homepage that "no email content or other personally identifiable information is ever provided to advertisers,"⁵ the privacy implications of Gmail's use were unclear to many users when it launched.⁶

⁵ <http://gmail.google.com/gmail/help/about.html> (accessed April 16, 2004)

⁶ Equally unclear, however, is whether the confusion could have been avoided, since other factors beyond system and interaction design were at play. In particular, Google's idiosyncratic brand prominence and reputation for innovation, catalyzed by Gmail's sudden appearance, ensured an immediate—and immediately critical—market of both sophisticated and naïve users.

Evidence: Avoiding the Pitfall

Many web sites that require an email address for creating an account give clear notice on their sign-up forms that they do not share email addresses with third parties or use them for extraneous communication with the user. Clear, concise statements like these help clarify scope and are becoming more common.

Tribe.net is a social networking service that carefully makes clear that members' information will be made available only to other members within a certain number of degrees of social separation. Of course, this in no way implies that users' privacy is particularly safeguarded, but it does make explicit the basic scope of potential disclosures, helping the user understand her potential audience.

Pitfall 2: Obscuring Actual Information Flow

Having addressed the user's need to understand a system's *potential* privacy implications, we move now to instances of *actual* disclosure. To whatever degree is reasonable, designs should make clear the actual disclosure of information through the system. Users should understand *what* information is being conveyed to *whom*. The disclosure should be obvious to the user as it occurs; if this is impractical, notice should be provided within a reasonable delay. Feedback should sufficiently inform but not overwhelm the user.

By avoiding both this and the prior pitfall, designs can clarify the extent to which users' actions engage the system's range of privacy implications. This can help users understand the consequences of their use of the system thus far and predict the consequences of future use. In the Discussion section, we will elaborate on how avoiding both of these pitfalls can support the user's mental model of his personal information flow.

We will not dwell on this pitfall, for it is perhaps the most obvious of the five. We suggest Bellotti and Sellen [1] as a guide to exposing actual information disclosure.

Evidence: Falling into the Pitfall

Web browser support for cookies is a persistent example of obscuring information flow [31]. Most browsers do not, by default, indicate when a site sets a cookie or what

information is disclosed through its use. The prevalence of third-party cookies and web bugs (tiny web page images that facilitate tracking) exacerbates users' ignorance of who is observing their browsing activities.

Another example of concealed information flow is in the Kazaa P2P file-sharing application, which has been shown to facilitate the concealed disclosure of highly sensitive personal information to unknown parties [15].

Another example is worn locator badges like those described in [8, 13], which generally do not inform their wearers about who is locating them.

Evidence: Avoiding the Pitfall

Friedman *et al*'s. redesign of cookie management reveals *what* information is disclosed to *whom*. They extended the Mozilla web browser to provide prominent visual feedback about the real-time placement and characteristics of cookies, thereby showing users what information is being disclosed to what web sites [32].

Instant messaging systems tend to employ a symmetric design that informs the user when someone wants to add him to her contact list, allowing him to do the same. This way he knows who is likely to see his publicized status. Further, his status is typically reflected in the user interface, indicating exactly what others can learn about him by inspecting their buddy lists.

AT&T's mMode Find Friends service, which lets mobile phone users locate other users of the service, informs the user when someone else is locating them. They learn *who* is obtaining *what* (their location).

Concerning Action

Our last three pitfalls involve a system's ability to support the conduct of socially meaningful *action*. Rather than occurring through specific configurations of technical parameters, everyday privacy regulation often occurs through the subtle manipulation of coarse controls across devices, applications, and time. Observers discern socially meaningful actions through the accumulation of evidence across these media. Privacy-sensitive technical systems can help users intuitively shape the nature and extent of this evidence to influence the social consequences of their behavior.

Pitfall 3: Emphasizing Configuration over Action

Designs should not require excessive configuration to create and maintain privacy. They should enable users to practice privacy management as a natural consequence of their ordinary use of the system.

Palen and Dourish write, “setting explicit parameters and then requiring people to live by them simply does not work, and yet this is often what information technology requires... Instead, a fine and shifting line between privacy and publicity exists, and is dependent on social context, intention, and the fine-grained coordination between action and the disclosure of that action” [3]. But because configuration has become a universal user interface design pattern, many systems fall into the configuration pitfall.

Configured privacy breaks down for at least two reasons. First, in real settings users manage privacy semi-intuitively; they do not spell out their privacy needs in an auxiliary, focused effort [14]. Configuration imposes an awkward requirement on users, one they will often forsake in favor of default settings [11, 33]. If users are to manage their privacy at all, it needs to be done in an intuitive fashion, as a predictable outcome of their situated actions involving the system.

A second reason configured privacy breaks down is that the act of configuring preferences is too easily desituated from the contexts in which those preferences apply. Users are challenged to predict their needs under hypothetical circumstances, and they can forget their preferences over time. If they predict wrongly, or remember incorrectly, their configured preferences will differ from their *in situ* needs, creating the conditions for an invasion of privacy.

People generally do not set out to explicitly protect their privacy. Rather, they participate in some activity, with privacy regulation being an embedded component of that activity. Designs should take care not to extract the privacy regulation process from the activity within which it is normally conducted.

Evidence: Falling into the Pitfall

An abundance of systems emphasize explicit configuration of privacy, including experimental online identity managers [27, 34], P2P file-sharing software [15], web browsers [31], and email encryption software [14]. In the realm of ubiquitous computing,

both our Faces prototype and Bell Labs's Houdini Project [35] require significant configuration efforts prior to and after disclosures.

Evidence: Avoiding the Pitfall

Successful solutions can involve some measure of configuration, but tend to embed it into the actions necessary to use the system. Web sites like Friendster.com and Tribe.net allow users to regulate information flow by modifying representations of their social networks—a process that is embedded into the very use of these applications.

Dodgeball.com's real-time socio-spatial networking service also directly integrates privacy regulation into the primary use of the system. Dodgeball members socially advertise their location by sending a brief, syntactically constrained text message from their mobile device to Dodgeball's server, which then sends an announcement to the member's friends—and friends of friends—that are within walking distance. Identifying one's friends to the system does require specific configuration effort, but once done, regulating location privacy is integrated with the very use of the system. Each use actively publicizes one's location; concealing one's location simply involves not using the system.

Georgia Tech's In/Out Board [36] lets users reveal or conceal their presence in a workspace by badging into an entryway device. Its purpose is to convey this information, but it can be intuitively used to withhold information as well, by falsely signaling in/out status with a single gesture.

Ignoring the moral implications, another example involves camera surveillance. When someone is aware of a camera's presence, she tends to adjust her behavior to present herself in alignment with the perceived expectations of her ostensible observers [37]. She does not step outside herself to reconfigure her representation. She simply acts, albeit with "appropriate" intuition and/or intention.

Cadiz and Gupta propose a smart card that one could hand to a receptionist to grant him limited access to one's calendar to schedule an appointment; he would hand it back right afterwards. No one would have to fumble with setting permissions. They also suggest extending scheduling systems to automatically grant meeting partners access to the user's location during the minutes leading up to a meeting, so they can infer his

arrival time. The action of scheduling a meeting would imply limited approval of location disclosure [38].

Pitfall 4: Lacking Coarse-Grained Control

Designs should offer an obvious, top-level mechanism for halting and resuming disclosure. Users are accustomed to turning a thing off when they want its operation to stop. Often a power button or exit button will do the trick.

Beyond binary control, a simple ordinal control may also be appropriate in some cases (*cf.*, audio devices' volume and mute controls). Ubicomp systems that convey location or other context could incorporate both a *precision dial* (ordinal) and a *hide button* (binary), so users can either adjust the precision at which their context is disclosed or decidedly halt disclosure.

In the general case, users can become remarkably adept at wielding coarse-grained controls to yield nuanced results (*e.g.*, driving a car requires use of a wheel, a stick, and two or three pedals, but their manipulation yields tremendous results). Plus, coarse-grained controls often reflect their state, providing direct feedback and freeing the user from having to remember whether she set a preference properly. This helps users accommodate the controls and even co-opt them in ways the designer may not have intended. Examples specific to privacy include: setting a door ajar, covering up or repositioning cameras [1, 39], turning off a phone or using its invisible mode rather than navigating its privacy-related options, and removing a worn locator badge.

While some fine-grained controls may be unavoidable, the flexibility they are intended to provide is often lost to their neglect (see Pitfall 3), which is then compensated for by the nuanced manipulation of coarse-grained controls across devices, applications, and time.

Evidence: Falling into the Pitfall

E-commerce web sites typically maintain users' shopping histories. While this informs useful services like personalization and collaborative filtering, there are times when a shopper does not want the item at hand to be included in his actionable history; he effectively wants to shop anonymously during the current session (beyond the private

transaction record in the merchant’s database). For example, the shopper may not want his personalized shopping environment—which others can see over his shoulder—to reflect this private purchase. In our experiences, we have encountered no web sites that provide a simple mechanism for excluding the current purchase from our profiles.

Similarly, most web browsers still bury their privacy controls under two or three layers of configuration panels [31]. While excessive configuration may itself be a problem (see Pitfall 3), the issue here is that there is typically no top-level control for switching between one’s normal cookie policy and a “block all cookies” policy. Third-party applications that elevate cookie control widgets have begun to appear (*e.g.*, GuideScope.com).

Further, wearable locator-badges like those described in [8] and [13] do not have power buttons. One could remove the badge and leave it somewhere else, but simply turning it off would at times be more practical or preferable.

Evidence: Avoiding the Pitfall

Systems that expose simple, obvious ways of halting and resuming disclosure include easily coverable cameras [1], mobile phone power buttons, instant messaging systems with invisible modes, the In/Out Board [36], and our Faces prototype.

Pitfall 5: Inhibiting Established Practice

Designs should beware inhibiting existing social practice. People manage privacy through a range of established, often nuanced, practices. For simplicity’s sake, we might divide such practices into those that are already established and those that will evolve as new technologies of disclosure emerge. While early designs might lack elegant support for emergent practices—since, obviously, substantive practice cannot evolve around a system until after deployment—designs can at least take care to *avoid inhibiting established ones*.

This is effectively a call to employ privacy design patterns. Designers of privacy-affecting systems can identify and assess the existing genres of disclosure into which their systems will be introduced. By supporting—and possibly enhancing—the roles,

expectations, and practices already at play in those genres, designs can accommodate users' natural efforts to transfer existing skills to new media.

Beyond genre-specific practices and patterns, certain meta-practices are worth noting. In particular, we emphasize the broad applicability of plausible deniability (whereby the potential observer cannot determine whether a lack of disclosure was intentional) [40, 41] and disclosing ambiguous information (*e.g.*, pseudonyms, imprecise location). These common, broadly applicable techniques allow people to finesse disclosure through technical systems to achieve nuanced social ends. Systems that rigidly belie meta-practices like plausible deniability and ambiguous disclosure may encounter significant resistance during deployment [42].

Technical systems are notoriously awkward at supporting social nuance [43]. Interestingly, however, systems that survive long enough in the field often contribute to the emergence of new practice even if they suffer from socially awkward design in the first place (*e.g.*, see [44, 45]). In other words, emergent nuance happens. But being intrinsically difficult to predict, seed, and design for, it generally doesn't *happen* as optimally as we might like it to. Designers will continue to struggle to support these *emergent* practices, but by identifying existing genres of disclosure and successful privacy design patterns, they can at least help users transfer *established* skills to new technologies and domains.

Evidence: Falling into the Pitfall

Some researchers envision context-aware mobile phones that disclose the user's activity to the caller to help explain why their call was not answered [46]. But this prohibits users from exploiting plausible deniability. There can be value in keeping the caller ignorant of the reason for not answering.

Location-tracking systems like those described in [8] and [13] constrain users' ability to incorporate ambiguity into their location disclosures. Users can only convey their concise location or—when permitted—nothing at all.

Returning to the privacy controversy surrounding Google's email system, one possible reason for people's discomfort with Gmail's content-triggered advertising is its inconsistency with the long-established expectation that the content of one's mail is for

the eyes of the sender and the recipient only. With respect to this pitfall, the fact that Gmail discloses no private information to advertisers, third-parties, or Google employees is not the issue. The issue is the plain expectation that mail service providers (electronic *or* physical) will interpret a correspondence's meta-data (electronic headers or physical envelopes) but never its contents. Many people would express discomfort if the US Postal Service employed robots to open people's mail, scan the contents, reseal the envelopes, and send content-related junk mail to the recipient. Even if no private information ever left each robot, people would react to the violation of an established social expectation, namely, the inviolability—under normal conditions—of decidedly private communications.

Evidence: Avoiding the Pitfall

Mobile phones, push-to-talk phones [41], and instant messaging systems [40] let users exploit plausible deniability by not responding to hails and not having to explain why.

Although privacy on the web is a common concern, a basic function of HTML allows users to practice ambiguous disclosure. Forms that let users enter false data facilitate anonymous account creation and service provision.

Tribe.net supports another established practice. It allows users to cooperatively partition their social networks into *tribes*, thereby letting both pre-existing and new groups represent themselves online, situated within the greater networks to which they are connected. In contrast, Friendster.com users each have a single set of friends that cannot be functionally partitioned.

Discussion

Having described the five pitfalls and provided evidence of systems that fall into and avoid them, we now examine some of the deeper implications they have for design. We begin by elaborating on the influence of our first two pitfalls on the user's mental model of the trajectories his information disclosures. This leads to the introduction of a new conceptual tool to help the design process. Then we present an analytical argument for why designs that avoid our five pitfalls can support the human processes of

understanding and action necessary for personal privacy maintenance. Using our Faces prototype as a case study, we then show how falling into these pitfalls can undermine an otherwise ordinary design process. Finally we discuss some successful systems that have largely avoided the pitfalls.

Mental Models of Information Flow

As we said earlier, avoiding our first two pitfalls—obscuring potential and actual information flow—can clarify the extent to which users’ actions engage the system’s range of privacy implications. Users can understand the consequences of their use of the system thus far, and they can predict the consequences of future use.

Illuminating disclosure contributes constructively to the user’s mental model of the portrayal of her identity in the context of the system. If she has a reasonable understanding of what observers can learn about her (Pitfall 1) and of what they already know about her (Pitfall 2), she can maintain and exploit this mental model to influence the portrayal of her identity and associated activities over time.

In the context of interactive systems, the personal information a user conveys is often tightly integrated with her interaction with the system. For example, by simply browsing the web, a user generates a rich clickstream that can be used by observers in ways that directly impact her life. When interaction and disclosure are integrated thusly, an informed user’s mental model of the system’s operation and her mental model of her disclosures are interdependent.

This suggests an extension to Norman’s canonical elucidation of the role of mental models in the design process. According to Norman, the designer’s goal is to design the system image (*i.e.*, those aspects of the implementation with which the user interacts) such that the user’s mental model of the system’s operation coincides with the designer’s mental model of the same [21].

When we take into account the coupling of interaction and disclosure, we see that the designer’s goal has expanded. She now strives to design the system image such that the user’s mental models of the system’s operation *and* of the portrayal of his identity and activities through it are both accurate. As with Norman’s original notion, ideally the designer’s and the user’s models of the system’s operation will coincide. But the designer

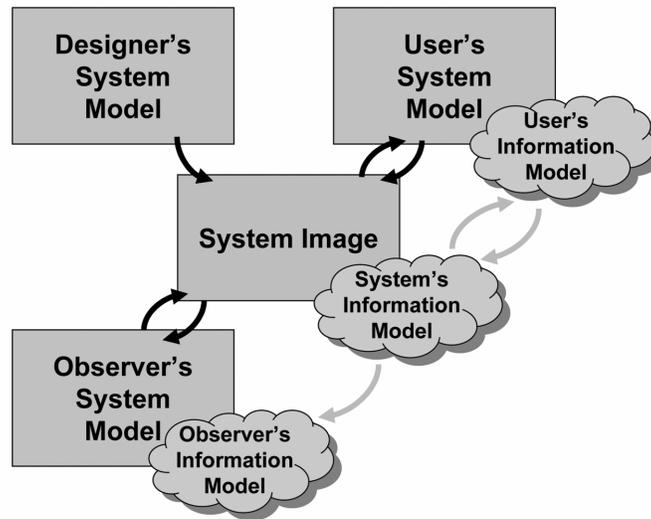


Figure 6. Building on Norman's elucidation of the role of mental models in the design process, designers can aim to harmonize the user's and the observer's understandings of the user's personal information disclosures.

generally cannot have a model of the user's personal information; that depends on the user and the context of use. Indeed, here the designer's task is not to harmonize the user's model of his information flow with her own (she likely has none), but to harmonize the user's information model with the *observer's* (Figure 6). In other words, she wants to design the system image to accurately convey a model not only of how other parties *can* observe the user's behavior through the system, but also what they *do* observe.

Generalizing this notion beyond privacy—to cooperative information flow in general—may be of further use to the computer-supported cooperative work (CSCW) community but is beyond the scope of this paper.

Opportunities for Understanding and Action

We have argued that people maintain personal privacy by *understanding* the privacy implications of their socio-technical contexts and influencing them through socially meaningful *action*. When a technical system is embedded into a social process, the primary means its designers have to engender understanding and action are its feedback and control mechanisms. We encourage designers of privacy-affecting systems to think of feedback and control mechanisms as *opportunities* for understanding and

action. They are the designer’s opportunity to empower those processes, and they are the user’s opportunity to practice them.

Thinking thusly can help designers reach across what Ackerman calls the *socio-technical gap*—the difference between systems’ technical capabilities and their social requirements [43]—just enough to empower informed social action. The challenge is to find that intermediate point where carefully designed technical feedback and control translates into social understanding and action. Reaching too far can overwhelm the user. Reaching not far enough can disempower him.

We believe that avoiding our pitfalls can help designers reach that intermediate point. Carefully designed feedback about potential (#1) and actual (#2) information flow can help users understand the representation and conveyance of their behavior through the system. Curtailing configuration (#3), providing coarse-grained control (#4), and supporting established practices (#5) can help people make productive, intuitive use of a privacy-affecting system. Designs that heed these suggestions make their consequences known and do not require great effort to use, helping people incorporate them meaningfully into the lexicons of personal privacy practices by which they engage everyday life’s genres of disclosure.

Case Studies

Negative Case Study: Faces

We return now to Faces—our prototypical ubicomp privacy UI—as a case study in how to fall into the pitfalls.

Pitfall 1: Obscuring Potential Flow. In trying to be a UI for managing privacy across any ubicomp system, Faces abstracted away the true capabilities of any underlying system. Users could not gauge its potential information flow because it aimed to address *all* information flow. Its scope was impractically broad and effectively incomprehensible.

Pitfall 2: Obscuring Actual Flow. Faces conveyed actual information flow through the disclosure log. Each record was accessible after the relevant disclosure. While this design intends to illuminate information flow, it is unclear whether postponing

notice is optimal. Embedding notice directly into the real-time experience of disclosure might foster a stronger understanding of information flow.

Pitfall 3: Configuration over Action. Faces required a considerable amount of configuration. Once configuration was done, and assuming it was done correctly, the system was designed to require little *ad hoc* configuration. The user would simply go about his business. But the sheer amount and desituated nature of configuration severely limited the system's chances of operating in alignment with the user's *in situ* preferences, positioning Faces squarely in this pitfall.

Pitfall 4: Lacking Coarse-grained Control. Faces avoided this pitfall by including an *override* function that afforded quick transitions to alternate faces.

Pitfall 5: Inhibiting Established Practice. While Faces modeled the nuance of Goffman's identity management theory, it appeared to hinder its actual practice by requiring the user to maintain virtual representations of his fragmented identities *in addition to* manifesting them naturally through intuitive, socially meaningful behavior.

Our evaluation of Faces revealed a complex, abstract configuration requirement that belies the intuitive situatedness of privacy as practiced in real settings. Faces also aimed to singularly address privacy needs across an arbitrary range of ubicomp systems and information types, a task whose futility becomes apparent upon recognizing that privacy management extends across systems, involving fluid, heterogeneous assemblies of technologies, practices, and information types. Rather than attempting to revise Faces to address our evaluation findings, we found it more appropriate to retire the Faces concept and scale our design focus down to a more isolable point in the ubicomp privacy space. In the following section, we assess an interaction concept that emerged from that process.

(Potentially) Positive Case Study: Precision Dial

One of Faces' core features—adjustable information precision—is a common privacy management technique in research (*e.g.*, [47, 48]) and could serve as the basis for a more streamlined ubicomp privacy tool. Here we briefly propose such a tool and suggest how its design might steer around the pitfalls better than Faces did. We will call this tool the *precision dial*.

The precision dial would be an easily accessible dial or rocker switch on a mobile phone that lets the user quickly adjust the precision of contextual information—sometimes referred to as *presence*—disclosed to his personal contacts on-the-fly. When an observer requests the user’s presence information, it would be blurred according to his current precision setting. He could quickly change precision settings as needed, similar to the practice of adjusting ringer volume when entering meetings and theaters. Pre-configuring privacy preferences would not be required, as it was in Faces.

Rather than a continuous precision scale (which seems rather implausible), we will assume the same four-point precision scale used in Faces (undisclosed < vague < approximate < precise). The dial would allow quick selection of one of these four points. In contrast to Faces’ encapsulation of separate precision settings for each dimension of information (location, nearby people, *etc.*), the dial would apply a single precision across all active dimensions. The rationale here is that, rather than being a separable dimension of context, the user’s *activity* is effectively constituted and represented by the sum of his context. Disclosing, say, where someone is and whom he is with could be tantamount to disclosing his activity, since observers can exploit personal or normative knowledge to infer activity from context. Hence if the user intends to blur the representation of his activity in a system that intentionally conveys presence, the easiest way to do so might be to apply a single transformation command to all disclosable information.

We envision the option to create groups of known observers (like friends/family/colleagues groupings in instant messaging clients) and to specify a default precision for each group. When adjusting precision *in situ*, the user could adjust for a specific group or for all observers.

To be clear, we do not intend the precision dial as a general user interface for ubicomp privacy. In fact, we hope this article makes clear the futility of such an idea. We envision the dial as a tool for managing the coarse representation of one’s activity as conveyed through real-time presence awareness systems.

Reminding the reader of the speculative nature of this assessment—since the precision dial is merely a proposed concept, not a tested tool—we suggest that, in comparison to Faces, the precision dial might heed the five pitfalls in the following ways.

Pitfall 1: Obscuring Potential Flow. Unlike Faces, this tool is deliberately scoped to a specific subspace of the privacy space: intentional interpersonal disclosure of activity—as presence—to familiar observers. In other words, it lets friends discover one’s activity, with permission. A system employing this tool should clarify its operational definition of presence. And it should clarify that information is conveyable only to people on the user’s contact list. By letting users collect observers into groups, they can know who has the potential to obtain what information about them at which precisions.

Pitfall 2: Obscuring Actual Flow. Disclosures might be exposed through real-time indicators, alerts, a disclosure log, or a combination thereof.

Pitfall 3: Configuration over Action. Adjustable disclosure precision would appear to align with the common practice of letting audiences know *just enough* information about your activity to satisfy their information needs without revealing sensitive details. A readily accessible dial could allow the timely manipulation of one’s publicized presence to achieve socially meaningful—perhaps nuanced—results. Managing groups might present a configuration a burden, but good design practices can minimize it. For instance, the user could have the option to quickly choose a group for each observer at the time he adds her to his contact list.

Pitfall 4: Lacking Coarse-grained Control. One cannot get much coarser than an ambiguous four-point ordinal precision scale. Nonetheless, we have chosen the number of points rather arbitrarily. A three-point scale might be better. Any coarser would result in a binary button, but we suspect people would prefer to leverage some gray area between the extremes of disclosing everything and disclosing nothing.

Pitfall 5: Inhibiting Established Practice. The precision dial supports both ambiguous disclosure and plausible deniability. The former is a consequence of the intrinsic ambiguity of the precision scale. The latter is supported by the observer’s ignorance of the reasons why the user employed any given precision level; it may have been due to social expectations (*i.e.*, the user may have simply been adhering to the relevant genre of disclosure), or due to technical factors (*e.g.*, signal loss), or simply the desire to be left alone.

Positive Case Study: Instant Messaging and Mobile Telephony

Interestingly, two systems that largely avoid our pitfalls—mobile phones and instant messaging (IM)—are primarily *communication* media. That is, disclosure is their central function. We will briefly assess these services against the pitfalls, focusing on their primary functions—textual and vocal communication—and on some of their secondary features that support these functions. We will not address orthogonal, controversial features like the location-tracking capabilities of some mobile phones and the capture of IM sessions, which would have to be addressed by a more robust assessment of the privacy implications of these technologies.

IM and mobile telephony each make clear the potential and actual flow of disclosed information, making for a robust, shared mental model of information flow through these cooperative interactive systems. *Potential flow* is scoped by features like Caller ID (telephony), Buddy Lists (IM), and feedback about the user's own online presence (IM). *Actual flow* is self-evident in the contents of the communications. Each technology requires minimal *configuration* for maintaining privacy (though secondary features often require excessive configuration), largely due to *coarse-grained controls* for halting and resuming information flow—*e.g.*, invisible mode (IM), application exit (IM), power button (telephony), and ringer volume (telephony). Lastly, each supports *existing practices* of plausible deniability—people can choose to ignore incoming messages and calls without having to explain why—and ambiguous disclosure—the linguistic nature of each medium allows for arbitrary customization of disclosed information [40, 41].

Indeed, communication media might serve as a model for designing other privacy-affecting systems not conventionally categorized as communication technologies. Disclosure *is* essentially communication, whether it results from the use of a symmetric linguistic medium—*e.g.*, telephony—or an asymmetric event-based medium—*e.g.*, e-commerce, context-aware systems. Systems that affect privacy but are not positioned as *communication* media do nonetheless communicate personal information to observers. Exposing and addressing these disclosure media as communication media might liberate designs to leverage users' intuitive privacy maintenance skills.

Conclusion

In this paper we described five common pitfalls to which designs of privacy-affecting systems often succumb. These pitfalls include obscuring potential information flow, obscuring actual flow, emphasizing configuration over action, lacking coarse-grained control, and inhibiting established practice. We provided several examples of systems that fall into or manage to avoid them, including Faces, our user interface prototype for managing ubicomp privacy.

We further identified some conceptual tools to help heed the pitfalls, including privacy design patterns; the metaphor of personae and activities as, respectively, indices to and contents of subspaces of a user's identity space; and an extension of Norman's elucidation of the role of mental models in the design process, in which the designer also works to align the user's mental model of his information flow with his observers'.

In closing, and in the spirit of Palen and Dourish (2003), we encourage designers of privacy-affecting systems to identify the genres of disclosure in which their systems will participate and—with the help of our guidelines and others—to design opportunities for users to (1) understand the extent of the system's alignment with those genres and (2) conduct socially meaningful action that supports them (or disrupts them, as the case may be).

Acknowledgements. This work was funded by Grant No. IIS-0205644 of the United States National Science Foundation and by a United States Department of Defense NDSEG fellowship. We are intensely grateful for the assistance and insights of Jennifer Mankoff, Chris Beckmann, danah boyd, Karen Teng, Jeff Huang, Xiaodong Jiang, the anonymous reviewers of this article and an earlier draft, and the participants of the studies mentioned herein.

References

1. Bellotti V and Sellen A (1993) Design for Privacy in Ubiquitous Computing Environments. Proceedings of the Third European Conference on Computer Supported Cooperative Work (ECSCW 93), Kluwer, Milano, Italy, pp 77-92.

2. Langheinrich M (2001) Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems. Proceedings of the Third International Conference on Ubiquitous Computing (UbiComp 03), Springer-Verlag LNCS, Atlanta, GA, USA, pp 273-291.
3. Palen L and Dourish P (2003) Unpacking "privacy" for a networked world. Proceedings of the Conference on Human Factors in Computing Systems (CHI 2003), ACM Press, Fort Lauderdale, FL, USA, pp 129-136.
4. Jiang X, Hong JI and Landay JA (2002) Approximate Information Flows: Socially-based Modeling of Privacy in Ubiquitous Computing. Proceedings of the Fourth International Conference on Ubiquitous Computing (UbiComp 02), Springer-Verlag LNCS, Göteborg, Sweden, pp 176-193.
5. Taylor H (2003) Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits. The Harris Poll, Rochester, NY, USA.
6. Cranor L, Reagle J and Ackerman MS (2000) Beyond Concern: Understanding Net Users' Attitudes About Online Privacy. In Vogelsang I and Compaine BM (eds). The Internet Upheaval: Raising Questions, Seeking Answers in Communications Policy, MIT Press, pp 47-70.
7. Turow J (2003) Americans and Online Privacy: The System is Broken. Annenberg Public Policy Center, University of Pennsylvania, Philadelphia, PA, USA.
8. Harper RHR, Lamming MG and Newman WH (1992) Locating Systems at Work: Implications for the Development of Active Badge Applications. *Interacting with Computers* 4(3):343-363.
9. Kaasinen E (2003) User needs for location-aware mobile services. *Personal and Ubiquitous Computing* 7(1):70-79.
10. Lederer S, Mankoff J and Dey AK (2003) Who wants to know what when? Privacy preference determinants in ubiquitous computing. Extended Abstracts of the Conference on Human Factors in Computer Systems (CHI 2003), ACM Press, Ft. Lauderdale, FL, USA, pp 724-725.
11. Palen L (1999) Social, Individual & Technological Issues for Groupware Calendar Systems. Proceedings of the Conference on Human Factors in Computing Systems (CHI 99), ACM Press, Pittsburgh, PA, USA, pp 17-24.
12. Adams A (2000) Multimedia Information Changes the Whole Privacy Ballgame. Proceedings of the Conference on Computers, Freedom, and Privacy (CFP 2000), ACM Press, Toronto, Ontario, Canada, pp 25-32.
13. Beckwith R (2003) Designing for Ubiquity: The Perception of Privacy. *IEEE Pervasive* 2(2):40-46.
14. Whitten A and Tygar JD (1999) Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. Proceedings of the 8th USENIX Security Symposium, Washington, DC, USA.
15. Good NS and Kregelberg A (2003) Usability and privacy: a study of Kazaa P2P file-sharing. Proceedings of the Conference on Human factors in Computing Systems (CHI 2003), ACM Press, Ft. Lauderdale, FL, USA, pp 137-144.
16. Lederer S, Mankoff J, Dey AK and Beckmann C (2003) Managing Personal Information Disclosure in Ubiquitous Computing Environments, Technical Report CSD-03-1257. UC Berkeley, Berkeley, CA, USA.
17. Westin A (1995) Privacy in America: An Historical and Socio-Political Analysis. Proceedings of the National Privacy and Public Policy Symposium, Hartford, CT.
18. Gellman R (1998) Does Privacy Law Work? In Agre PE and Rotenberg M (eds). *Technology and Privacy: The New Landscape*, MIT Press, Cambridge, MA, USA, pp 193-218.
19. Westin A (1967) *Privacy and Freedom*. Atheneum, New York, NY, USA.
20. Altman I (1975) *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. Brooks/Cole Publishing, Monterey, CA, USA.
21. Norman DA (1988) *The Design of Everyday Things*. Basic Books, New York, NY, USA.

22. Weiser M (1991) The Computer for the Twenty-First Century. *Scientific American* 265(3):94-104.
23. Bellotti V, Back M, Edwards WK, Grinter RE, Henderson A and Lopes C (2002) Making sense of sensing systems: five questions for designers and researchers. *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2002)*, ACM Press, Minneapolis, MN, USA, pp 415-422.
24. Goffman E (1956) *The Presentation of Self in Everyday Life*. Doubleday, New York, NY, USA.
25. Langheinrich M (2002) A Privacy Awareness System for Ubiquitous Computing Environments. *Proceedings of the Fourth International Conference on Ubiquitous Computing (UbiComp 2002)*, Springer-Verlag LNCS, Göteborg, Sweden, pp 237-245.
26. Adams A and Sasse MA (1999) Taming the Wolf in Sheep's Clothing: Privacy in Multimedia Communications. *Proceedings of the Seventh ACM International Conference on Multimedia*, Orlando, FL, USA, pp 101-107.
27. boyd d (2002) *Faceted Id/Entity: Managing representation in a digital world* (M.S. Thesis). Massachusetts Institute of Technology, Cambridge, MA, USA.
28. Phillips DJ (2002) Context, Identity, and Privacy in Ubiquitous Computing Environments. *Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing, UbiComp 2002 Conference*, Göteborg, Sweden.
29. Reang P (2002) Dozens of nurses in Castro Valley balk at wearing locators. *Mercury News*, San Jose, CA, Sept. 6, 2002.
30. Baertlein L (2004) Calif. Lawmaker Moves to Block Google's Gmail. *Reuters*, Apr 12, 2004.
31. Millett LI, Friedman B and Felten E (2001) Cookies and Web browser design: toward realizing informed consent online. *Proceedings of the Conference on Human Factors in Computing Systems (CHI 2001)*, ACM Press, Seattle, WA, USA, pp 46-52.
32. Friedman B, Howe DC and Felten EW (2002) Informed Consent in the Mozilla Browser: Implementing Value-Sensitive Design. *Proceedings of the 35th Annual Hawaii International Conference on System Sciences (HICSS 02)*, Hawaii, USA.
33. Mackay WE (1991) Triggers and barriers to customizing software. *Proceedings of the Conference on Human Factors in Computing Systems (CHI 91)*, ACM Press, New Orleans, LA, USA, pp 153-160.
34. Jendricke U and Gerd tom Markotten D (2000) Usability meets Security – The Identity-Manager as your Personal Security Assistant for the Internet. *Proceedings of the 16th Annual Computer Security Applications Conference (ACSAC 00)*, New Orleans, LA, USA, pp 344-355.
35. Hull R, Kumar B, Lieuwen D, Patel-Schneider P, Sahuguet A, Varadarajan S and Vyas A (2004) Enabling Context-Aware and Privacy-Conscious User Data Sharing. *Proceedings of the IEEE International Conference on Mobile Data Management (MDM 2004)*, Berkeley, CA, USA.
36. Dey AK, Salber D and Abowd GD (2001) A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications. *Human-Computer Interaction* 16(2-4):97-166.
37. Foucault M (1977) *Discipline and Punish*. Vintage Books, New York, NY, USA.
38. Cadiz J and Gupta A (2001) Privacy Interfaces for Collaboration, Technical Report MSR-TR-2001-82. Microsoft Corp., Redmond, WA, USA.
39. Jancke G, Venolia GD, Grudin J, Cadiz JJ and Gupta A (2001) Linking public spaces: technical and social issues. *Proceedings of the Conference on Human Factors in Computing Systems (CHI 01)*, ACM Press, Seattle, WA, USA, pp 530-537.

40. Nardi BA, Whittaker S and Bradner E (2000) Interaction and Outeraction: Instant Messaging in Action. Proceedings of the Conference on Computer Supported Cooperative Work (CSCW 00), ACM Press, New York, NY, USA, pp 79-88.
41. Woodruff A and Aoki PM (2003) How Push-to-Talk Makes Talk Less Pushy. Proceedings of the International Conference on Supporting Group Work (GROUP 03), ACM Press, Sanibel Island, FL, USA, pp 170-179.
42. Suchman L (1997) Do Categories have Politics? The language/action perspective reconsidered. In Friedman B (eds). Human Values and the Design of Computer Technology, Center for the Study of Language and Information, Stanford, CA, USA, pp 91-106.
43. Ackerman MS (2000) The Intellectual Challenge of CSCW: The Gap Between Social Requirements and Technical Feasibility. *Human-Computer Interaction* 15(2/3):181-203.
44. Green N, Lachoe H and Wakeford N (2001) Rethinking Queer Communications: Mobile Phones and beyond. Proceedings of the Sexualities, Medias and Technologies Conference, University of Surrey, Guildford, UK.
45. boyd d (2004) Friendster and Publicly Articulated Social Networks. Extended Abstracts of the Conference on Human Factors in Computing Systems, ACM Press, Vienna, Austria, in press.
46. Siewiorek D, Smailagic A, Furukawa J, Krause A, Moraveji N, Reiger K, Shaffer J and Wong F (2003) SenSay: A Context-Aware Mobile Phone. Proceedings of the IEEE International Symposium on Wearable Computers, White Plains, NY, USA.
47. Boyle M, Edwards C and Greenberg S (2000) The effects of filtered video on awareness and privacy. Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW 00), ACM Press, Philadelphia, PA, USA, pp 1-10.
48. Hudson SE and Smith I (1996) Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW 96), ACM Press, Boston, MA, USA, pp 248-257.