Managing Personal Information Disclosure in Ubiquitous Computing Environments

Scott Lederer, Jennifer Mankoff, Anind K. Dey, Christopher P. Beckmann

Report No. UCB/CSD-3-1257

July 2003

Computer Science Division (EECS)
University of California
Berkeley, California 94720

Managing Personal Information Disclosure in Ubiquitous Computing Environments

Scott Lederer[1], Jennifer Mankoff [1], Anind K. Dey [2], and Christopher P. Beckmann [1]

[1]Group for User Interface Research, University of California, Berkeley, CA 94720, USA
{lederer, jmankoff, beckmann }@cs.berkeley.edu
[2]Intel Research, Berkeley, Intel Corporation, Berkeley, CA 94720, USA
anind@intel-research.net

Abstract. Ubiquitous computing stands to redefine established notions of privacy as it introduces regular, pervasive sensing of personal information such as identity, location, and activity. To effectively and comfortably manage the disclosure of personal information, end-users will require a coherent conceptual model of privacy and convenient user interfaces to manage it. We describe a conceptual framework designed to support personal privacy management in ubiquitous computing by empowering users to adjust the precision of disclosed information. The framework relies on three key strategies: encapsulation, *a priori* configuration, and manual configuration. We describe a prototypical user interface built to instantiate this framework and we report the results of a formative evaluation of the framework. Results show our approach is superior to simple, automated disclosure paradigms but can be further refined.

1 Introduction

Ubiquitous computing implies automated disclosure of personal information. Disclosure of identity is required to personalize a service. Disclosure of location is a by-product of the situated disclosure of identity. Disclosure of activity provides input to self-reconfiguring intelligent services. Indeed, ubiquitous computing infrastructures point toward the emergence of a global heterogeneous real-time database composed of people, places, and things instead of records, tables, and fields, whereby any party with proper permissions can access one's personal information in real-time. As ubiquitous computing services start truly becoming ubiquitous, people will engage them throughout much of their lives, disclosing a stream of personal information accessible by parties near and far. This implies a crucial need for end-user control over information disclosure. Our contribution in this paper is a conceptual framework designed to support personal privacy management for end users in ubiquitous computing. We instantiated this framework in a user interface which we evaluated. We present our framework, interface, and evaluation of the framework, which we found to be an improvement over simple, automated disclosure paradigms currently in common use such as disclosing no information or all information.

We take as fundamental that people should be legally and technically empowered to decide when and what to disclose to whom. Our framework uses three techniques to empower users to actualize their privacy preferences:

- *Encapsulation.* The large set of inquiries for information, and the variety of contexts in which such inquiries occur can lead to unmanageable complexity for end users. We seek to identify the minimal set of user-level concepts that together encapsulate the multivariate nature of personal information disclosure into a usable and effective privacy management framework.
- *A priori configuration.* It seems unreasonable to expect users to be comfortable consenting (or not) to each and every inquiry in real-time, as the potential number of interruptions is prohibitively high. Logic precludes the act of consenting to disclosure from occurring after the disclosure. Hence our framework emphasizes the configuration of preferences at convenient times prior to disclosure. Preferences persist through all disclosures until the user changes them.
- *Manual configuration.* Our framework emphasizes the manual configuration of preferences. Given the sensitive nature of privacy, we believe decisions about it should be made by the people whose privacy is in question. In a 2003 Harris poll, 79% (N=1,010) of subjects said "being in control of who can get information about you" is "extremely important" [18]. Although this approach could be combined with machine learning, we believe that end user configuration is still necessary to mitigate errors and provide a feeling of control.

The central notion behind our design is that *people disclose different versions of personal information to different parties under different conditions*. This notion is rooted in the work of Goffman [10], who used a theatrical metaphor to explore the "fronts" people employ when playing different social "roles". For example, if your spouse sends you an instant message while you are at work and asks what you are doing, you might reply that you are "busy…ttyl?"[1] But if your boss were to make the same inquiry under the same conditions, you might reply that you are "Proofing the trip report. I'll have it in ten minutes." In the first case, your role is that of "employed spouse," in the second your role is "person responsible for the trip report." In each case, the front you present is tailored not only to the pertinent audience, but also to the current conditions, and it determines the amount of information you are willing to disclose to the audience. That is, if these same two people made this same inquiry while you were calculating your taxes one evening, you might reveal more information to your spouse than to your employer, rather than the converse.

Interpreting Goffman's fronts as means of controlling the immediate flow of personal information, or, managing privacy, we see that managing privacy in everyday life is an intuitive, situated social process. Managing privacy in ubiquitous computing, however, remains a complicated process involving multiple local and remote observers to whom one may want to present multiple fronts simultaneously. For example, if you have called in sick to work but are actually interviewing for a job at a competing firm, when your boss inquires about your location through a ubiquitous computing system, you may not want to disclose your precise location. But how would you prac-

[1] "ttyl" is instant messaging shorthand for "talk to you later."

tically assert this privacy preference? Would your mobile phone alert you to the inquiry and provide you an opportunity to tailor the response? Would you have preconfigured the service that morning to disclose false information in response to any inquiries from your boss? The interactional means of managing privacy in ubiquitous computing have until now remained unexamined.

We address this issue by attempting to infuse ubiquitous computing privacy management with the intuitiveness of Goffman's fronts. While others have used similar approaches to managing personal information flow in *online* environments (*e.g.*, [4, 11]), to our knowledge, this work is the first to use this approach in *ubiquitous computing* environments.

We believe the key to enabling people to automatically present different fronts to different parties under different conditions in ubiquitous computing is the intentional, automatic adjustment of the *precision* of disclosed personal information, according to *a priori*, manual user configuration made manageable by focusing on a small number of user-level concepts. That is, under certain conditions you may want to allow certain parties to know *precisely* where you are or what you are doing, while under the same conditions you may want to allow others to know *vaguely* where you are or what you are doing. To return to our earlier example, your response to your spouse ("busy") was a less precise version than your response to your boss ("proofing trip report"). Assuming the system were able to infer your activity at some precision, it could arguably transform that inference into a less precise version depending on who was inquiring about it, thereby presenting different fronts to different inquirers.

The remainder of this paper is organized as follows. Section 2 provides our definition of privacy and describes our conceptual framework in detail, highlighting the three key abstractions of inquirer, situation, and face. Section 3 describes the user interface prototype we built, which allows a user to specify her privacy preferences using those abstractions. Section 4 focuses on a formative evaluation of the framework, which we found performed better than a variety of automated techniques despite some confusion about our abstractions. Section 5 is a discussion of the implications of the results that suggests some modifications to our framework. Section 6 differentiates our solution from other related work that did not concentrate on end-user ubicomp privacy management and Section 7 concludes the paper.

2 A Conceptual Framework for Privacy Management

This section presents our definition of privacy, and then discusses a conceptual model of privacy derived from interviews with twelve end users and a study of existing literature on end user conceptualizations of privacy.

2.1 Operational Definition of Privacy

We define privacy as control over the immediate flow of one's personal information. By "immediate" we mean initial disclosure; we do not address the control of personal information flow beyond initial disclosure in this work.

In the interests of lucidity, we have limited our scope to five dimensions of personal information: *identity*, *location*, *activity*, identities of *nearby people*, and *profile*, the latter being composed of the user's email address, phone number, street address, gender, age, and occupation.

2.2 Ordinal Precision Scale

The notion of information precision is central to privacy management in ubiquitous computing. Not only can technical conditions such as sensor noise and imperfect inferencing degrade information accuracy, but also users can intentionally adjust information precision to exert privacy control [9]. Disclosure need not be a binary decision; an individual could insist that his location be disclosed at a certain level of precision. For example, a student might allow his advisor to know which building he is in when on campus, but not his precise location within the building.

Different categories of personal information would seem to require different scales of precision. Nonetheless we are experimenting with a general ordinal precision scale that might map operationally to specific precision scales. From highest to lowest precision, our scale consists of the following values:

- Precise
- Approximate
- Vague
- Undisclosed

Table 1 shows how each precision level might affect what would be disclosed for each dimension of personal information. Alternatively, a user can specify a custom string to be disclosed in lieu of an automatic transformation. For example, a user might want to always disclose her location as "the library" whenever her advisor inquires about it, regardless of her actual location.

Table 1. Normalized precision levels of personal information dimensions

	Identity	**Location** *indoor/outdoor*	**Activity**	**Nearby People**
Precise	True name	Room/Block	Precise	Names
Approximate	Pseudonym	Building/District	Categorical	Roles
Vague	Role	Municipality	Busy / Not Busy	Number
Undisclosed	*Undisclosed*	*Undisclosed*	*Undisclosed*	*Undisclosed*

2.3 Conceptual Framework

Adams has shown that, in the context of a given audio/video-captured environment, the most important factors in determining an individual's perception of privacy are the perceived identity of the information *recipient*, the perceived *usage* of the information, the subjective *sensitivity* of the information, and the *context* in which the captures were made [2]. We uncovered further evidence of the importance of these

factors in a study we conducted involving twelve people solicited from the general public in our community, each of whom completed a questionnaire asking them to rate the importance of thirteen factors in determining the appropriate level of personal privacy in a given situation. Figure 1 illustrates the results' general alignment with Adams' findings. Inquirer ("recipient"), usage, sensitivity, and situation ("context") are among the most important factors. The next most important factor is data content, or the information being disclosed.

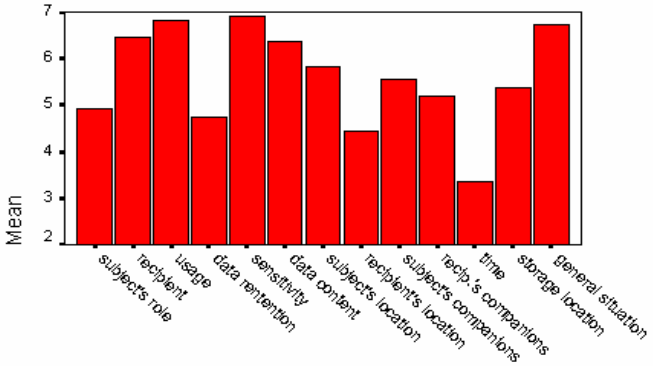

Figure 1. Mean rating of the importance of thirteen factors in determining preferred privacy in a given situation *(1=Not important, 7=Extremely Important)*.

By empowering a user to specify preferences about the precision at which his personal information is to be disclosed to certain inquirers under certain conditions, our design addresses three of Adam's four factors: the identity of the inquirer, the user's situation at the time of inquiry, and the information sensitivity (made mutable by precision preferences). We do not address information usage, as it is unlikely that structured inquiries would include a description of intended usage.

Our conceptual framework of privacy management operationalizes these three factors in the form of three core abstractions:

- *Inquirer*. This is the entity requesting some subset of the user's personal information. Inquirers might be specific people familiar to the user, specific companies or organizations, or general classes of entities (*e.g.*, clothing retailers, strangers, business associates).
- *Situation*. This is an encapsulation of contextual information. A situation is a four-tuple, {*location*, *activity*, *time*, identities of *nearby people*}, where any element can be a wildcard. Hence a situation might be as simple as a specific location (*e.g.*, *at home*) with all other parameters ignored, or as complex as *walking with my son in Central Park on Sunday mornings*. Whenever all of a situation's conditions are met, that situation is considered *active*.
- *Face*. This is an encapsulation of disclosure precision preferences. In specifying a face, a user specifies (1) a level of precision on our ordinal precision scale for each of four dimensions of personal information (identity, location, activity, and nearby people), and (2) the subset of his profile that he is willing to disclose. That is, a face is a five-tuple {identity precision, location precision, activity precision, nearby

people precision, profile subset} describing the preferred levels of precision at which each dimension of information should be disclosed when that face is used. The metaphorical "face" was chosen as a more colloquial term than "front."

The essence of the framework, then, is that people present different faces to different inquirers in different situations, where a face is a set of preferences regarding the precision of personal information disclosed to the inquirer. Referring to Figure 2, which illustrates the framework, let us return to our earlier example. An inquirer (*i.e.*, your boss or your spouse) inquires for your activity. The identity of the inquirer, together with your current situation (*i.e.*, proofing your trip report at work), determines which face to use, according to your prior configuration of the system. Let us assume the face you present to your boss while at work is configured to reveal your precise activity, while the face you present to your spouse in the same situation transforms your activity to "busy" or "not busy" depending on how focused you are on your current task and whether any deadlines are imminent. The system then automatically transforms your information according to the preferences encapsulated in the appropriate face, and discloses the transformed information to the inquirer.

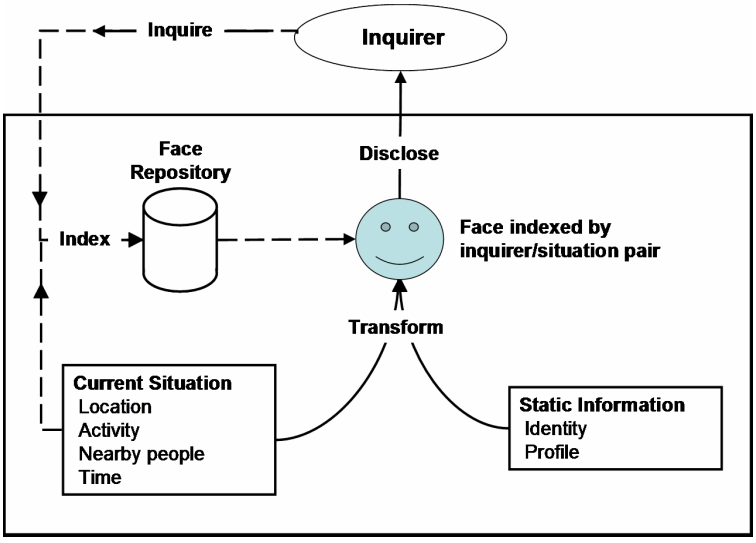

Figure 2. An illustration of the conceptual framework

3 A Prototypical User Interface for Privacy Management

In this section we describe the Java-based prototype we built to instantiate and evaluate our conceptual framework in an experimental laboratory setting. In a production setting, the system would require a ubiquitous context-aware sensing and dissemina-

tion infrastructure for determining active situations, collecting users' personal information, routing inquiries, and disclosing information to inquirers. In our experimental setting, this infrastructure, as well as the algorithm for precision adjustment of information being disclosed, is simulated using Wizard-of-Oz techniques.

Our system allows people to create inquirer, situation, and face objects and to assign faces to inquirers, optionally parameterized by situation. The specific face used to handle a given inquiry is indexed by the appropriate inquirer/situation pair (the current version allows only one face to be assigned to a given inquirer/situation pair). In earlier work, we found evidence that the identity of the inquirer is a stronger determinant of one's ubiquitous computing privacy preferences than is one's situation at the time of inquiry [**Error! Reference source not found.**]. Accordingly, our interface emphasizes inquirers as the primary objects to which faces are assigned, but these assignments can be parameterized by situation objects.

Note that since faces are situation- and inquirer-independent, they are somewhat abstract to the end user, not being situated in a specific disclosure instance. This will turn out to be a crucial issue in our evaluation and is discussed in depth in Section 5.

Situations and faces are reusable. A single face can be assigned to multiple inquirers and a single situation can be used to parameterize face assignments for multiple inquirers. Special cases arise when an inquiry is made by an unfamiliar inquirer (*i.e.*, an inquirer for whom the user has no preferences assigned) when none of the user's situations are active. We created two special objects to represent these wildcard cases for inquirers and situations. We call these the *General Public* and the *Default Situation*, respectively. These two cases intersect to create a third special case, the *Default Face*. We explain each of these below.

- *General Public*. This special *inquirer* represents all inquirers that the system does not recognize.
- *Default Situation.* This special *situation* is active whenever none of the user's specified situations are active.
- *Default Face*. By assigning a *face* to the General Public in the Default Situation, the user indicates her preferences for handling inquiries made by unfamiliar inquirers (*e.g.*, a retail store in the mall in which she is shopping) whenever none of her situations are active.

The prototype consists of two synchronized user interface modules: a PC-based module that supports in-depth configuration, and a lightweight handheld component that affords rapid interaction on the fly. We will describe each in turn.

PC Component. The main screen of the PC component (**Figure 3**) is used to create and edit objects representing inquirers, situations, and faces, and to assign faces to inquirers, optionally parameterized by a situation. In **Figure 3**, the user has selected the *Roommate* inquirer and the *Studying* situation. The field in the upper-right corner of the screenshot displays the name of the face that will handle inquiries made by the selected inquirer when the user is in the selected situation. In this example, if the user were to click the *Student* face and then the up-arrow, she would have assigned the *Student* face to handle all inquiries made by her roommate whenever she is studying.

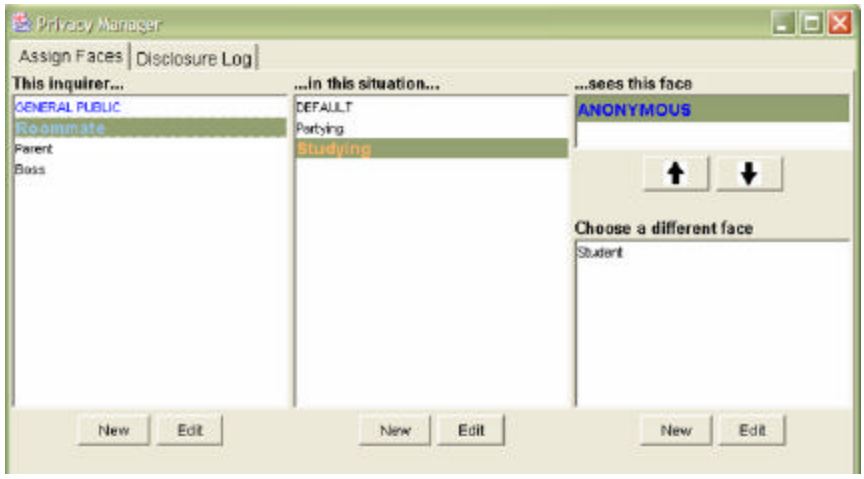

Figure 3. PC Component. Main screen for creating inquirers, situations, and faces, and binding them to each other.

A second screen displays a log of the personal information disclosed to inquirers. For each inquiry, the user can see who made the inquiry, the user's context at the time of inquiry, the situation and face that determined the disclosure precision, and the actual information disclosed. We designed the disclosure log to provide users with feedback on what they have disclosed, thereby supporting notice, a fair information practice [12]. This feedback is intended to support an iterative configuration process whereby a user can react to an unsatisfactory disclosure by altering her preferences accordingly. While the log is an important feature of the system, we should note that we did not to evaluate the log in our study.

Handheld Component. In practice, we envision the handheld component operating on users' mobile phones, however the current prototype is confined to a small window on a PC or PDA running Java. The main screen (**Figure 4**) of the handheld module is divided into two portions. The top portion displays the user's precise context. The bottom portion displays the face that will handle inquiries made by unfamiliar inquirers (*i.e.*, the face assigned to the General Public in the current situation) under the current conditions. This screen gives the user immediate feedback about what will be disclosed to inquirers permitted to obtain precise information, and it shows which face the user is showing the general public.

In addition to the main screen and a simplified disclosure log, this module has two special features worth noting:

– *Override*. The override function allows the user to choose one specific face to handle all inquiries by all inquirers under any conditions until override is disengaged. While engaged, all previously configured preferences are overridden.

– *Situation Snapshot.* The situation snapshot function creates a record of the user's current context, which he can subsequently edit and label, thereby creating a new situation that can be used to parameterize face assignments.

Figure 4. The handheld module main screen. The top portion displays the user's precise context. The bottom portion displays the face that will handle inquiries made by unfamiliar inquirers (effectively, the General Public) under the current conditions.

4 Evaluation

In this section we describe a formative evaluation of our conceptual framework as instantiated by our prototype. The evaluation focused on the usefulness of:

1. *Encapsulation.* Encapsulating multiple dimensions into a larger category minimizes the number of immediate factors users have to consider, but the importance of the inner dimensions does not necessarily diminish in the process. Does encapsulation simplify or complicate the framework? Are users able to remember the values of the inner dimensions of a given encapsulation when necessary?
2. *A priori configuration.* Our framework assumes users will not want to be alerted to each inquiry as it occurs and will prefer to configure their preferences beforehand. But will de-situated, *a priori* configuration of abstract preferences lead to *a posteriori* satisfaction with specific instances of disclosure?
3. *Manual configuration.* Manually configuring the system consumes time and cognition. Simpler strategies for privacy management, such as always disclosing all information at a fixed precision level, might require zero configuration effort. Would simpler strategies requiring less cognitive effort result in similar or better *a posteriori* satisfaction with specific instances of disclosure?
4. *Abstractions.* Are the inquirer, situation, and face abstractions useful to end-users?

Before describing the details of our experiment, it is worth commenting on the difficulties of evaluating ubiquitous computing systems. This well-known problem has

received much recent attention [5, 16, 19]. One major difficulty is the fact that one may have to build a ubiquitous computing system in order to evaluate it. However, this is not an appropriate approach for early stage or iterative design because of the quantity of work necessary in each iteration.

As a support tool that is supposed to function in the context of other ubiquitous computing applications, the problems we faced were particularly difficult. To truly test our system, we would have had to build a variety of ubiquitous computing applications that make use of personal information in different ways, and get a number of third party vendors and information consumers to start using the system before we could realistically test the privacy management component.

To address this problem, we have taken a multi-tiered approach. First, we conducted a scenario-based study in which we interviewed 12 end users about how they conceptualized privacy. Our results are shown in Figure 1. It was based on this initial work, summarized in Section 2.3, that we chose the abstractions of inquirer, situation, and faces.

Second, we conducted a lightweight scenario-based online questionnaire to test the importance of different factors in making privacy decisions. Our results showed that inquirer was more important than situation in determining privacy preferences [**Error! Reference source not found.**].

Third, we conducted a lab-based experiment that combined scenarios with use of our prototype. We describe this experiment in more detail below.

4.1 Hypotheses

Based on the four issues we chose to investigate in our evaluation, we began our study with the following hypotheses:

1. Participants could accurately predict the precision at which each dimension of information would be disclosed in specific instances based on their earlier configuration of the system. This hypothesis pertains to the *encapsulation* issue.
2. Participants would be satisfied *a posteriori* with the precision of information disclosed in specific instances, as determined by the preferences they set *a priori.*
3. *Manual configuration* of preferences would meet participants' disclosure needs better than simpler disclosure strategies.
4. Inquirer, situation, and face are suitable abstractions for end-user privacy management in ubiquitous computing.

We evaluated the first three hypothesis quantitatively and the fourth qualitatively. We discuss the method of evaluation next.

4.2 Method

Participants. Four women and one man, all average computer users, from the non-engineering undergraduate student body at our university, were each given a $20 gift certificate to a local bookstore for participating. They had an average of 2.2 years of experience using mobile phones and almost no experience with PDAs. Four out of five subjects rated themselves as moderately to seriously concerned about both online and everyday privacy.

Procedure. Participants were interviewed for about 10 minutes to determine their reactions to specific scenarios involving sharing information with different recipients in different situations. We asked questions such as: "Imagine you are purchasing something in a store. The cashier asks for your phone number. Do you give it? When might you and when might you not?" This served to prime participants by making them think through their feelings and opinions about privacy in ubiquitous computing.

Participants were then given a verbal description of ubiquitous computing, and a demonstration of both the handheld[2] and desktop versions of our system. They were given as much time as they liked to explore the system themselves. They then completed a series of tasks using the PC interface. Finally, participants were briefly interviewed about their experience with the system.

Participants completed the following tasks:

1. Select or create one or more faces to satisfy two example sets of requirements we provided (*e.g.*, allow your roommate to know approximately where you are when you are studying, but not whom you are with).
2. Create two situations from your own life, and select or define default faces to be used in each situation (*i.e.*, assign a face to the General Public in that situation).
3. Create two inquirers from your own life, and select or define faces to be used for each possible combination of inquirer and situation, including the Default Situation and the two situations created in the previous task.

After a 5-minute break, participants were moved away from the system and asked to imagine that one of the inquirers they had specified had asked for all of their personal information during one of the situations they had specified. Participants were then asked the following:

- What do you *think* was disclosed (precision level and actual content) about each dimension of information in this instance?
- In retrospect, what would you *want* to disclose (precision level and actual content) for each dimension of information to *this* inquirer under *these* conditions?
- For each dimension of information how important is it that the precision of the disclosed information be at the level you want in this instance?

[2] Though participants were introduced to the handheld component to fortify their understanding of our conceptual framework, this component did not play a role in the evaluation.

This process was repeated for the second inquirer and second situation that participants had specified, and again for a third inquirer of our choice (a popular clothing retailer) in a well-known nearby shopping district.

4.3 Measures and Results

Most measures were made on a per-inquiry basis, not per-subject, though we include some per-subject results. With five subjects and three inquiries per subject, there were a total of 15 inquiries. Results indicate that our approach is superior to automated disclosure strategies save for the time and cognition it consumes and warrants further investigation. We discuss these results further in Section 5. Here, we present the results for each hypothesis in turn. The first three hypotheses were measured quantitatively, while the last was measured qualitatively.

Encapsulation. To measure participants' ability to remember the values of encapsulated dimensions of faces, we compared the precision levels actually configured with the precision levels the subject *thought* she configured after a break.

Subjects accurately predicted disclosure precision levels for 73% of inquiries with respect to identity, 87% with respect to location, 71% with respect to activity, 87% with respect to nearby people, and 60% with respect to profile.

***A priori* Configuration.** To measure participants' *a posteriori* satisfaction with the precision of each specific instance of disclosure, given that faces were designed *a priori* in the absence of specific disclosure events, we compared the precision levels the subject configured *a priori* with the precision levels the subject would have preferred *a posteriori.*

Measured across all inquiries, abstract *a priori* preferences matched *a posteriori* preferences 67% of the time for identity, 47% for location, 57% for activity, 80% for nearby people, and 47% for profile. With importance ratings taken into account, preferences configured *a priori* matched *a posteriori* preferences 67% of the time for identity, 87% for location, 100% for activity, and 87% for nearby people. The derivation of these numbers is discussed in the next hypothesis, which compares manual configuration's success rate to that of automated disclosure strategies.

Manual Configuration. To measure whether manual configuration would result in higher satisfaction than simple, automated strategies, we compared the precision levels the subject would have preferred *a posteriori* with the precision levels that would have been used in each of the following strategies:

- *Precise.* Always disclose all information at the precise precision level.
- *Approximate.* Always disclose all information at the approximate precision level.
- *Vague.* Always disclose all information at the vague precision level.
- *Undisclosed.* Never disclose any information.
- *Random.* For each dimension and for each inquiry, choose one of the above four levels at random. We ran the random trials twice.

For each strategy and for each dimension[3] of personal information, we measured the difference ? between the precision that would have been used by that strategy and the subject's *a posteriori* preferred precision for *each* inquirer in *each* situation. We then looked at the importance rating (1-7) the subject gave to each dimension. If the importance was 1 (not important), then any value for ? was considered satisfactory. If the importance was 2-3, ? could be no more than two levels of precision to be considered satisfactory. If importance was 4-5, ? could be no more than one level precision. If the importance was 6-7 (extremely important), ? had to be zero.

Table 2 shows the success rate of our manual configuration strategy compared to that of the automated strategies. In all measured dimensions except identity, our approach resulted in more satisfactory disclosures than the other strategies.

When examined on a per-subject basis and averaged across all four measured dimensions of information, manual configuration met participants' needs an average of 84% (SD 21) of the time, ranging from 58% to 100%. The next best option, never disclosing any information, had an average 63% (SD 32) success rate, ranging from 25% to 80%. Despite its relative success, not disclosing any information will not be a realistic long-term strategy for most people. Table 3 shows the success rates of the remaining automated techniques.

Table 2. Percentage of disclosures, across all participants, that would have satsified participants' *a posteriori* precision preferences for a given instance of disclosure, given the importance that preferences be met. The highest number in each column is shown in bold.

	Identity	**Location**	**Activity**	**Nearby People**
Manual	67%	**87%**	**100%**	**87%**
Precise	**73%**	60%	71%	27%
Approximate	40%	53%	79%	40%
Vague	27%	60%	71%	60%
Undisclosed	40%	73%	57%	80%
Random 1	47%	53%	79%	67%
Random 2	33%	60%	79%	53%

Table 3. Average percent of disclosures that would have satsified participants' *a posteriori* precision preferences for a given instance of disclosure, given the importance that preferences be met. These results are averaged across all information dimentions and reported per-subject. The highest number in each row is shown in bold.

Subject ID	**Manual**	**Undisclosed**	**Vague**	**Approx.**	**Precise**	**Random 1**	**Random 2**
1	**89%**	80%	80%	69%	69%	88%	80%
2	58%	50%	67%	**75%**	50%	67%	67%
3	**100%**	58%	41%	25%	50%	33%	33%

[3] We did not include the *profile* dimension in these measures, because, as formulated, it did not conform to the ordinal precision scale we employed.

Subject ID	Manual	Undisclosed	Vague	Approx.	Precise	Random 1	Random 2
4	**75%**	67%	58%	58%	58%	**75%**	58%
5	**100%**	58%	25%	33%	58%	42%	42%
Mean	**84%**	63%	54%	52%	57%	61%	56%
SD	21	32	28	33	33	32	25

Suitability of Abstractions Our last hypothesis was that the abstractions represented in our interface are appropriate for managing end-user privacy. Our results are based on interview data and critical incidents noted as participants used the system. Although most of our participants were able to successfully use our interface to meet their privacy preferences, there was some confusion between the situation and face abstractions. Finally, when constructing situations, participants made heavy use of the wildcard option, often specifying details for only one inner dimension, such as "location".

5 Discussion

Our primary concern in this discussion is the gap between precision preferences buried inside faces configured *a priori* in the absence of specific inquirers or situations and actual precision preferences in situated instances of disclosure. We will call this the abstract-situated gap. Naturally an entirely different framework in which users choose a precision level for each disclosure upon being alerted to each inquiry in real-time could eliminate this gap. But as discussed earlier, ubiquitous computing promises to flood users with such alerts. This option appears intolerable. We have chosen to alleviate this flood by shifting the act of consent to an occasional (re)configuration process, thereby creating the gap between stated, abstract preferences and actual, situated preferences. The goal of our ongoing iterative design of this system is to minimize this gap to the point of acceptance for users who want control over their privacy. Below we discuss the implications of the above evaluation on this process.

With regards to our first hypothesis, the 60-87% success rate in recalling privacy preferences highlights the existence of the abstract-situated preference gap. It is important to note, however, that asking someone to remember how he or she configured an unfamiliar system five minutes ago is a far cry from asking them to do so after long-term use of a familiar system.

Although the initial analysis of our second hypothesis suggested that participant-configured preferences did not match their actual preferences consistently, when information importance ratings were taken into account, our results were very positive. If a subject deems it not very important that a given dimension be disclosed at his preferred precision, then a difference of one, two, or even three levels of precision may be acceptable. With importance ratings taken into account, preferences configured *a priori* matched *a posteriori* preferences 67% of the time for identity, 87% for location, 100% for activity, and 87% for nearby people. In essence, this shows that when the importance of minimizing the abstract-situated gap is taken into account, our framework succeeds in meeting user preferences the majority of the time.

Our third hypothesis raised the question of whether manual configuration performed better than lightweight automatic strategies. Our framework outperformed all other strategies. Broken down across subjects, our framework performed particularly well for subjects 1, 3, and 5. Subjects 1, 3, and 5 were able to configure the system to meet their preferences (although it should be noted that subject 1 had very loose requirements, *i.e.*, she stated that it was not important that the data be disclosed precisely as she intended in most cases). This suggests that three of our subjects understood the underlying abstraction (faces) and were better able to configure it because of that. An alternative interpretation is that subjects 1, 3, and 5 maintained their preferences when moving from abstract to specific situations while subjects 2 and 4 did not. In any case, given that the disclosure of personal information in ubiquitous computing is imminent, our framework appears preferable to the other strategies we measured.

Our fourth hypothesis relates to the usability of the situation, inquirer, and face abstractions. Our data, combined with the quantitative results above, suggest the following modifications to our framework:

Preferred Precisions of Static and Dynamic Personal Information Have Inverse Relations to How Well the User Knows the Inquirer. Some of our participants pointed out a design flaw in grouping static and dynamic information precision preferences together. People need not worry about disclosing static information (*i.e.*, identity and profile) to parties with whom they have established relations because those parties already know (some of) that information. But people *are* concerned with whether and when familiar parties obtain their dynamic information (*i.e.*, location, activity, and nearby people), because that information can affect those relationships. On the other hand, people *are* concerned about disclosing static information to unfamiliar parties (*e.g.*, revealing one's identity and contact information to a retail store). But people are less concerned about revealing dynamic information to unfamiliar parties, since it will have less of an effect if one has no established relations with them.

The main implication here is that grouping identity and profile in the same encapsulation as location, activity, and nearby people complicates matters. Our next design will tease apart these concerns.

Situation and Face are Two Sides of the Same Coin. Especially if we concern ourselves only with dynamic information, then the situation is the precise version of one's dynamic information, and the face is the transformed version that one wishes to convey. Introducing a level of indirection between these concerns flies in the face of the situated nature of interaction vigorously explored by the likes of Suchman and Dourish [17,8]. This was demonstrated clearly when participants repeatedly confused these two abstractions. We are exploring ways to merge these abstractions, which should also decrease the complexity inherent in the 3-way inquirer/situation/face mapping.

Encapsulations May Introduce Unnecessary Indirection. Many participants created situations that contained a value for a single inner dimension (*e.g.*, location) and wildcards for all other dimensions. This implies it may be more intuitive to eliminate the situation abstraction and simply support parameterization of face assignments by one or more contextual dimensions *a la carte*. This would also help eliminate the confusion participants exhibited between situation and activity.

Face Is a Weak Metaphor for Dynamic Information. Face as a metaphor for identity may have merit, but it stretches one's conception as a metaphor for the precision of dynamic information. Some experts believe metaphors should be used sparingly (*e.g.*, [6]) and we agree. In our next design we are considering emphasizing precision *per se*, without masking it behind a metaphorical face.

Wildcards in a Three-Way Mapping are Confusing. The three-dimensional inquirer/situation/face space proved a bit confusing for participants. In particular, participants struggled to understand the implications of wildcards in the inquirer/situation/face mapping, *i.e.*, the General Public, Default Situation, and Default Face. Operationalizing these exceptional conditions in a complex space in an intuitive fashion proved quite challenging. By merging or eliminating the situation and/or face abstractions in our next design, we hope to mitigate this confusion.

In summary, while ongoing use of the system in realistic environments, including log-driven iterative reconfiguration, would be the real determinant of whether encapsulation, *a priori* configuration, and manual configuration are useful strategies for privacy management in ubiquitous computing, we believe our results show they hold promise and warrant further experimentation. We intend to keep these fundamental features intact while we iterate our design towards a usable framework for managing personal privacy in ubiquitous computing.

6 Related Work

A number of researchers have explored privacy and personal information management in online environments using Goffman's fronts [10]. Boyd explored the notion of fragmented identity and the social presentation of self in online communities [4]. Her SecureID system allows a user to specify "facets" of his identity, with each facet associated with certain documents and personal information. Other parties can access the data associated with a specific facet if they can answer a series of questions that guard it. Jendricke and tom Markotten experimented with identity-based security management on the Internet [11]. Using their Identity-Manager a user creates a set of identities, each encapsulating preferences about which pieces of her personal information can be disclosed to websites when that identity was engaged, and choose which identity to engage as she navigates the web. Both these systems use Goffmanesque strategies to manage users' static personal information, that is, their identities and profiles. Our work extends that approach to ubiquitous computing, where personal information changes dynamically.

Adams conducted empirical investigations into individuals' perceptions of privacy in an audio/video-captured environment. As mentioned, her findings played a key role in the design of our conceptual framework [2].

Langheinrich issued a call-to-arms to ubiquitous computing researchers to adhere to a set of privacy-sensitive design principles honoring fair information practices [12]. He subsequently outlined a privacy awareness system [13], extending P3P [7] to ubiquitous computing and routing all inquiries through the user's privacy proxy,

which handles them according to the user's preferences. Our work could be a user interface to such a system.

Palen and Dourish describe privacy management as "a dynamic response to circumstance rather than a static enforcement of rules" and emphasize the nuanced effects disruptive technologies have on privacy [15]. We hope to imbue our future work with the sort of careful sociological awareness for which they call.

7 Conclusions and Future Work

Ubiquitous computing is driving a fundamental shift in the paradigm of electronic privacy. As human-computer interaction moves beyond the networked desktop, and as automated data collection becomes a greater part of everyday life, privacy management efforts must address the greater scope of situated data collection and the increased complexity of personal information disclosure. To this end, we have introduced and evaluated a conceptual framework designed to support personal privacy management for end-users in ubiquitous computing.

Our framework emphasizes encapsulation, *a priori* and manual configuration of preferences, and an ordinal precision scale for managing personal information disclosure in ubiquitous computing. A formative evaluation showed that our approach is superior to simple, automated disclosure strategies, but can be simplified further still. We are refining our framework and will incorporate its changes into the next revision of our user interface prototype.

To be clear, privacy is a far more nuanced process than a set of preferences can reasonably represent [15,3]. In attempting to operationalize the management of this sensitive process, we have tried to infuse our framework with its intrinsic subtleties. If our solution is sub-optimal in this respect, we maintain that, in the words of Ackerman, we offer a "first-order approximation" to close part of the "socio-technical gap" between what systems are *supposed* to do socially and what they *can* do technically [1]. That is, our solution is neither final nor complete, but is a first step towards fulfilling the complex promise of manageable privacy in ubiquitous computing.

Acknowledgements

Thanks to Karen Teng, Jeff Huang, and danah boyd.

References

1. Ackerman, M.S. The Intellectual Challenge of CSCW: The Gap between Social Requirements and Technical Feasibility. In John M. Carroll, editor, Human-Computer Interaction in the New Millennium. Addison-Wesley, Reading, MA, 2002.

2. Adams, A. Multimedia information changes the whole privacy ballgame. In Proceedings of Computers, Freedom, and Privacy, 2000.
3. Agre, P., Changing Places: Contexts of Awareness in Computing. Human-Computer Interaction, 2001. 16(2-4): p. 177-192.
4. boyd, d. Faceted Id/entity: Managing representation in a digital world. M.S. Thesis, Mass. Inst. of Tech., (2002).
5. Consolvo, S. *et al.* User study techniques in the design and evaluation of a ubicomp environment. In Proceedings of *Ubicomp 2002*. pp. 73-90. 2002.
6. Cooper, A. The Myth of Metaphor. Visual Basic Programmer's Journal, June 1995.
7. Cranor, L., Langheinrich, M., Marchiori, M., and Reagle, J. The platform for privacy preferences 1.0 (P3P1.0) specification. W3C Recommendation, HTML Version at www.w3.org/TR/P3P/, April 2002.
8. Dourish, P. Where the Action Is. MIT Press, Cambridge, MA, 2001.
9. Geographic Location/Privacy IETF Working Group. Geopriv requirements. IETF Internet-Draft. March 2003.
10. Goffman, E. The Presentation of Self in Everyday Life. Anchor, Doubleday, New York, 1959.
11. Jendricke, U. and tom Markotten, D.G. Usability meets security- the identity-manager as your personal security assistant for the Internet. In Proceedings of *ACSAC 2000*. pp. 344-351, 2000.
12. Langheinrich, M. Privacy by design – Principles of privacy-aware ubiquitous systems. In Proceedings of *Ubicomp 2001*. pp. 273-291, 2001.
13. Langheinrich, M. A privacy awareness system for ubiquitous computing environments. In Proceedings of *Ubicomp 2002*. pp. 237-245, 2002.
14. Lederer, S. Mankoff, J., Dey, A.K.. Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. Short Talk in the Extended Abstracts of CHI 2003, ACM Conference on Human Factors in Computing Systems, pp. 724-725, April 5-10, 2003.
15. Palen, L. and Dourish, P. Unpacking "privacy" for a networked world. In Proceedings of *CHI 2003*. pp. 129-136, 2003.
16. Scholtz, J. *et al*. User-centered evaluations of ubicomp applications. Intel Research Technical Report IRS-TR-02-014. May 2002.
17. Suchman, L. Plans and Situated Actions. Cambridge University Press, New York, NY, 1987.
18. Taylor, H. Most people are "privacy pragmatists" who, while concerned about privacy, will sometimes trade it off for other benefits. The Harris Poll #17, March 19, 2003.
19. Trevor, J. *et al.* Issues in personalizing shared ubiquitous devices. In Proceedings of *Ubicomp 2002*. pp. 56-72. 2002
20. Weiser, M. The computer for the 21st century. *Scientific American* 265(3). pp. 94-104, 1991.