Towards a Deconstruction of the Privacy Space

Scott Lederer, Jennifer Mankoff, and Anind Dey

IRB-TR-03-037

September, 2003

DISCLAIMER: THIS DOCUMENT IS PROVIDED TO YOU "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE. INTEL AND THE AUTHORS OF THIS DOCUMENT DISCLAIM ALL LIABILITY, INCLUDING LIABILITY FOR INFRINGEMENT OF ANY PROPRIETARY RIGHTS, RELATING TO USE OR IMPLEMENTATION OF INFORMATION IN THIS DOCUMENT. THE PROVISION OF THIS DOCUMENT TO YOU DOES NOT PROVIDE YOU WITH ANY LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS

Copyright 2002, Intel Corporation, All rights reserved.

Intel Research Berkeley

Towards a Deconstruction of the Privacy Space

Scott Lederer, Jennifer Mankoff
Group for User Interface Research
Computer Science Division
University of California, Berkeley
{lederer, jmankoff}@cs.berkeley.edu

Anind K. Dey
Intel Research Berkeley
2150 Shattuck Ave., Suite 1300
Berkeley, CA, 94704
anind@intel-reseach.net

ABSTRACT
When designing for or discussing privacy, it is critical to identify the conditions that create a system's or phenomenon's privacy implications. We present a set of interdependent dimensions that, when applied to the analysis of a privacy-related system or phenomenon, can expose the factors that determine the role of privacy therein. This can help focus the scope of discourse and the design of privacy in the HCI and ubiquitous computing communities.

Keywords
Privacy, HCI, Ubiquitous Computing, Privacy Space

INTRODUCTION
Privacy is an enormous topic, incorporating multiple perspectives, from the sociological (*e.g.*, [5]), to the legal (*e.g.*, [4]), to the computational (*e.g.*, [13]), to the interactive (*e.g.*, [18]). This often leads to inconclusive discourse on privacy. Parties to the discussion start with asymmetric assumptions about just what systems, actors, relations, and information are under consideration. Are we talking about individuals disclosing transactional information to businesses, or about groups keeping secrets from the state, or about individuals withholding activities from other individuals, or about encryption as a tool for anonymity, or about public surveillance cameras, or about wiretaps, or about…? These examples illustrate the range of issues often conflated when we talk about privacy. Although privacy is an important factor in each of them, they highlight different points in the privacy space.

The purpose of this paper is to offer the HCI and ubiquitous computing communities a framework for exposing key aspects of any given privacy-related phenomenon (*i.e.*, technical systems, policies, practices, and incidents) that influence the conditions and implications of privacy therein. In other words, our goal is to help researchers and designers identify the conditions that create the privacy implications specific to any given phenomenon, so they can better analyze and address them.

To do this, we disambiguate the privacy space into a set of interdependent dimensions that help define a phenomenon's implications for privacy. These dimensions are not exhaustive; we have chosen them because of their high impact on the end-user privacy experience.

We cluster these dimensions into three categories: *system properties*, *actor relations*, and *information types*. System properties determine important aspects of the mechanisms of disclosure and how participatory the disclosure is. Actor relations help determine the type of relationship between subject and observer and how much history they share, all of which affect how the observer might use the information and how much the subject might trust the observer. Information types help determine how sensitive the information is and how intentional the disclosure is.

By assessing a privacy-related phenomenon with respect to these dimensions, one can expose the elements of that phenomenon that sensitize the disclosure of personal information. Such an assessment does not explain *why* these elements sensitize disclosure; it simply highlights them amidst the cluttered privacy space. In other words, situating a privacy phenomenon into the space neither defines the phenomenon nor specifies a solution to its privacy implications, but rather narrows the scope of those implications, thereby focusing discourse and design efforts and facilitating comparative analyses.

While this work has an intrinsic emphasis on HCI and ubiquitous computing research, it may be useful to the legal, regulatory, cryptographic, and other communities as well. It applies to a range of privacy phenomena, including technical systems, social practices, and regulatory frameworks. By determining exactly how the phenomenon at hand relates to each dimension of the space, domain experts can isolate the phenomenon's pertinence to their field, *e.g.*, HCI practitioners and researchers can identify design goals and classify research efforts.

Related Work
This work continues the unpacking of privacy begun by Palen and Dourish, who, employing Altman's notion of privacy as a continuous dialectic social process [5], articulated the need for the HCI community to transcend conventional role-based and access-control notions of privacy [15]. They exposed privacy as the ongoing negotiation of multiple boundaries in continuous tension. Our work unpacks privacy from a different angle, articulating key properties of the systems, actors, and information involved in a disclosure.

Recently, Brunk categorized a considerable number of online privacy tools into a feature space [7]. Our articulation of the privacy space applies not only to tools, but to privacy phenomena in general, *e.g.*, protocols, laws, norms, and even specific incidents. It aims for general applicability by teasing out common concepts and relations that determine the implications for privacy in and across phenomena.

Adams and Sasse isolated four interdependent factors of the perception of privacy in multimedia environments: information receiver, sensitivity, and usage, and the context of disclosure [1]. Our work differs by extending the privacy space beyond subjective perception, though we do emphasize the subject's perspective.

DIMENSIONS OF THE PRIVACY SPACE

We describe the privacy space by exposing an interdependent set of dimensions that help define it. When people talk about privacy, often what they are really discussing is not the nebulous whole of privacy, but some set of phenomena classifiable into subspaces of the privacy space. Such subspaces are effectively defined by revaluing the dimensions that constitute the privacy space.

While some of these dimensions are somewhat linear (*e.g.*, *familiarity*) and others are more categorical (*e.g.*, *feedback and control*), we have collapsed many of them into two categories (*e.g.*, *surveillance vs. transaction*) to keep the space manageable. The resulting loss of fidelity is worth the clarity of a simpler articulation of this complex space. Importantly, the categories of a dimension are not exclusive; a phenomenon can involve both poles of a dimension. But by articulating its relation to each pole, one disambiguates its implications with respect to that dimension.

Table 1 lists the dimensions of the privacy space we have chosen to address, classified into three categories: *system properties*, *actor relations*, and *information types*. *System properties* are the *how* of disclosure and disclosure management; they describe crucial aspects of the technical systems involved in the disclosure. *Actor relations* are the *who* of disclosure; they categorize the relations between the primary participants in the disclosure. *Information types* are the *what* of disclosure; they distinguish between critically distinct types of disclosable information.

Table 1. The dimensions of the privacy space

System Properties
Feedback and Control
Surveillance vs. Transaction
Actor Relations
Interpersonal vs. Institutional Disclosure
Familiarity
Information Types
Persona vs. Activity
Primary vs. Incidental Content

In describing each dimension below, we will illustrate some of the ways in which they relate to each other, for they are generally not orthogonal. Rather than rigidly defining the privacy space, they are a set of flexible, conceptual guides to help identify the approximate subspace pertinent to a given privacy-related phenomenon.

System Properties

Privacy phenomena differ in part by the degree to which subjects have feedback about and control over the disclosure process, and by the disclosure's status as surveillance or transaction.

Feedback and Control

Different privacy-related systems employ different ratios, degrees, and methods of *feedback* about and *control* over the disclosure process. To use a simple example to illustrate different *ratios* of feedback and control, consider two public spaces, each with a camera observing it and a sign declaring the camera's operation and purpose. In one space, there is an obvious and accessible switch to disable the camera; in the other there is no switch. Each space provides some reasonable feedback, but only one provides any direct control. The privacy implications of these systems differ fundamentally.

Degrees and *methods* of feedback and control can vary widely and obviously. For example, a red light on the camera in the public space is a different method of, and arguably provides a different degree of, feedback about the camera's operation and purpose.

We are not declaring that more feedback or control is better than less, or that one method of feedback or control is better than another, or that a perfect balance of feedback and control is optimal. Targeted guidelines for designing feedback and control mechanisms (*e.g.*, [6]) can help designers answer those questions. We are merely suggesting that discourse on any given privacy phenomenon should identify the ratio, degrees, and methods of feedback and control at play therein.

Surveillance vs. Transaction

Related to the participatory nature of feedback and control is the distinction between *surveillance* (*e.g.*, public cameras) and *transaction* (*e.g.*, credit card purchases). These means of disclosure often differ by the level of participation on the subject's part, and by the disclosure medium's amenability to machine-processing.

Our distinction between surveillance and transaction aligns approximately with Lessig's distinction between the monitored and the searched [14] and with Agre's distinction between surveillance and capture [3].

Surveillance often implies a disempowerment of the subject, whose ability to intentionally participate in disclosure is hampered and whose behavior is often socially constrained

as a result [11]. Transactions, however, often imply a sense of intention or agreement on the part of the subject.

Examples of disclosure through surveillance include institutionally managed cameras, personal cameras, overseen and overheard behavior, and even less obvious phenomena like webblogs and camblogs, whereby information about or pictures of the subject can be posted beyond the subject's technical participatory reach.

Examples of disclosure through transaction include credit card transactions, RFID tags, and HTML forms. In each case, the subject has some technical ability to alter the disclosure by changing the content or conditions of the transaction.

Much of the distinction between surveillance and transaction is intrinsic to contemporary technology's efficiency at processing the format of the disclosed information. If a commodity PC could parse and alter a surveillance video stream as reliably and quickly as it can an XML-encoded transaction, it could be programmed to blur or remove the representations of specific individuals. Disclosure of presence and identity would become a transaction whose content could be altered by the subject. As computers improve at parsing and classifying multimedia input, the surveilled becomes transactable.

Further, conventionally surveilled information is made transactable by altering the means of information acquisition. For example, real-time location used to be conveyed to a remote party verbally, perhaps via telephone, or visually via camera. Now this information can be disclosed through more structured, computationally mutable media involving mobile phones, RFID tags, WiFi access points, and GPS receivers.

Some disclosures include elements of both surveillance and transaction. Consider the act of withdrawing cash from an automatic teller machine. A camera in the machine captures the image of the subject using the machine, and the bank maintains a record of the cash withdrawal. By not using the machine, *i.e.*, not participating in the transaction, the subject can choose to not participate in the surveillance, *i.e.*, not be captured by the camera.

When assessing privacy phenomena, it is important to distinguish the degrees to which the disclosure occurs through surveillance and through transaction. Instances of the former are often not amenable to interactive participation and may be best addressed through social processes like laws and norms, while instances of the latter may support interactive control of information flows.

With respect to ubicomp, it is important to recognize that many forms of disclosure typically considered surveillance may indeed occur through transaction (*e.g.*, location tracking), but that the scope and complexity of ubicomp ensure the continual reemergence of leaks and new forms of informal or unstructured surveillance (*e.g.*, camblogs).

Actor Relations

Privacy phenomena differ in part by the subject's relation to the observer: personal or institutional, and familiar or unfamiliar.

Interpersonal vs. Institutional

There is a crucial difference between revealing sensitive information to another person and revealing it to industry or the state. Subjects often find the need to regulate disclosure to both individual and institutional observers, but the consequences of disclosure tend to differ drastically. *Intentionally* disclosing sensitive information to another person in a social context often serves to strengthen trust. Such disclosure to an institution, however, tends to accompany the provision or maintenance of a service. *Unintentionally* disclosing sensitive information to another person often generates distrust and can even dissolve important relationships. Such disclosure to an institution, however, generally results in inconveniences like spam. Trust is an important element in both types of relations, but distrusting a software company and distrusting a spouse have fundamentally different implications.

It may be worth further dividing institutional disclosure into disclosure to institutions of which one is a member (*e.g.*, an employer) and disclosure to those of which one is not (*e.g.*, a merchant). We are considering this for subsequent articulations of the privacy space.

Subjects often disclose, implicitly or explicitly, preferential, contact, financial, and activity information to both interpersonal and institutional observers, but the format, fidelity, and use of disclosures tend to differ. Interpersonal disclosure often involves rich, dynamic, vocal or textual representations of present and past activities that influence social bonds. Institutional disclosure tends to involve dry, textual representations in databases that affect service provision or membership status.

Many phenomena involve both interpersonal and institutional disclosure. For example, personal email travels through institutionally controlled systems that often retain and use information thereabout. Spouses often share bank accounts, with each having access to the other's transaction records. An employee's relationship with her boss involves both interpersonal and institutional elements of disclosure. When classifying compound phenomena, it is important to tease apart the aspects that are interpersonal and those that are institutional.

Familiarity

The implications of a disclosure differ with the degree of *familiarity* between subject and observer. Familiarity can be bilateral[1] or unilateral, *i.e.*, both parties might be familiar with each other, or only one might be familiar with the other. The

[1] Bilateral familiarity neither implies nor excludes symmetric familiarity.

unilateral case divides further into cases where the subject is unaware of the observer (though not necessarily unaware of the observation), and cases where the potential observer is unaware of the subject, *e.g.*, a retail store that a shopper has neither visited nor purchased anything from, but whose reputation is known to the shopper.

Although a linear notion of familiarity is useful, two critical moments are worth noting: (1) when a subject is first known to an observer, and (2) when an observer is first known to a subject. If neither of these moments has occurred, then both subject and observer are unknown to each other; no disclosures have occurred. Interestingly, a subject can be known to an observer indirectly, *e.g.*, through reputation or rumor (interpersonal) or through the personal information market (institutional).

It is worth pointing out that familiarity is not a direct indicator of trust. One might be very familiar with a homicidal maniac, but one probably would not trust him very much. Nonetheless, the more familiar a subject is with an observer, the more informed the subject is to evaluate her level of trust in the observer, which will affect her comfort with disclosing a given set of information to him.

When assessing privacy phenomena, it is critical to evaluate the degrees of familiarity between potential and actual subjects and observers. When this factor is considered in the light of the other dimensions of the privacy space (*e.g.*, *persona vs. activity*, as will become obvious shortly), it can help explain the observer's motivations for collecting information and the subject's motivations for and comfort with disclosing it (or not).

Information Types

Privacy phenomena differ in part by the type of information being disclosed. For instance, disclosing the *existence* of a persona differs from disclosing information *about* that persona. Further, sensitive information can have different privacy implications depending on whether it is the primary or incidental content of the disclosure.

Persona vs. Activity

This dimension concerns the notion of identity. By *persona* we mean a unique identifier, to which a history of activities might be associated. Examples include true names, login names, email addresses, phone numbers, credit card numbers, frequent shopper numbers, and employee numbers. In this sense, conveying a persona primarily conveys the *existence* of an identity. Conveying a persona to an observer theretofore unfamiliar with the subject is the moment in which that observer obtains some degree of familiarity with the subject.

By *activity* we mean information *about* the subject. We connote this with "activity" because (1) arguably, any information *about* a person results from some activity, often involving the person directly, and (2) ubicomp promises to convey increasing amounts of information *about* people by conveying the contexts of their activities (two intrinsically inseparable notions [2]). Conveying activities associated with a persona increases the observer's familiarity with the subject.

The fidelity of activity conveyed is roughly proportionate to the number and/or precision of data points conveyed. Take three pieces of context, for example: location, action, and duration. If Alice's location, action, and duration are conveyed as "Using laptop in café on New Montgomery Street for two hours," a fairly constrained range of possible activities is disclosed. Lowering the number of data points by, say, disclosing only location ("café on New Montgomery Street") decreases the degree to which the subject's activity is conveyed. Maintaining the same data points but lowering their precision ("Currently using a computer in San Francisco") will also obscure the subject's activity.

Persona and activity both convey *identity*, but in different ways. Persona conveys the existence of an identity. Activity conveys its essence. Conveying a persona merely signals the existence of an identity fragment. It opens the door to a dark room. But as the activities associated with that persona are disclosed to that observer with greater precision or frequency or duration, a history of activity is increasingly revealed. The room becomes brighter, its furnishings more obvious. Notably, an observer can often accelerate the illumination of a newly encountered persona through auxiliary means (*e.g.*, reputation or the personal information market).

Conveying activity independent of an associated persona is effectively anonymity. It is worth noting, however, that anonymity dissolves as the observer develops a consolidated history of activity in association with a specific (decreasingly anonymous) persona. Eventually, a sufficiently informed observer can infer one of the subject's established personae from external sources (*e.g.*, the personal information market), or else can assign an arbitrary persona that, for all practical purposes, can become an established persona for the subject (*e.g.* as a public bus driver grows familiar with one of his regular (anonymous) passengers, his maintained conception of her identity can have practical effects for her; perhaps he holds the bus a few minutes when she is occasionally late to arrive at the bus stop).

An important thing to note is the inverse relationships that persona and activity disclosure have with the observer's familiarity with the subject. When the subject is unfamiliar, his activity is generally less sensitive, because of anonymity, yet his personae are more sensitive, because disclosing a persona immediately eliminates anonymity to some degree. On the other hand, when the subject is known, his personae are less sensitive, since the observer

already knows some subset of his personae[2], but his activities become more sensitive, since disclosing activities out of character with his known persona can reveal hidden personae and collapse the boundaries between his fragmented identities [12, 16].

Primary vs. Incidental Content

Here we distinguish between whether the sensitive information is the primary content or an incidental byproduct of the disclosure. We generally consider this distinction from the subject's perspective. The consequences of disclosure may be comparable in either case, but disambiguating the origin of the sensitive information can determine how intentional the disclosure is, thereby sharpening analysis of the phenomenon and clarifying design requirements.

Perhaps the most infamous class of sensitive *incidental* content is transaction-generated information (TGI). TGI is meaningful information generated in the course of a service transaction [17]. It is typically used by service providers to personalize service provision, generate aggregate reports, and create user profiles that can be sold for profit. TGI emerges, for instance, from every credit card purchase, every phone call, and every email transmission.

Every disclosure contains both primary and incidental content, and either or both can be sensitive. For example, whispering a secret to a friend conveys the secret (primary), but also conveys, perhaps disagreeably, to a nearby observer that you are sharing secrets with the recipient (incidental).

It is particularly useful to distinguish between primary and incidental content when discussing context-aware systems. Context-aware services typically exploit incidentally generated information to personalize service provision; the disclosure of context to these systems is often incidental to the subject's primary actions. However, context can be repositioned as the primary content of activity disclosure to personal contacts (*e.g.*, the AT&T Find Friends service, which converts incidental, institutionally collected mobile phone location information into the primary content of interpersonal disclosures).

An important distinction in ubicomp is that between systems that automatically pull personal information from the user, and those that only use information pushed out by the user. This notion goes to the heart of *intentionality* of disclosure and aligns roughly with our distinction between primary (cf. push) and incidental (cf. pull) content.

[2] And accidental disclosure of a hidden persona to a familiar observer can often be smoothed over with social dexterity, *e.g.*, *"That* email address? I only use that to absorb spam."

CONNECTING THE DIMENSIONS

Having articulated all of the dimensions, it is worth examining how they relate to each other. In the interests of space, we will focus on the relation of *primary vs. incidental content* to three other dimensions.

With respect to *persona vs. activity disclosure*, personae and activity can both be disclosed as primary or incidental content. For example, if the primary content of a disclosure is "someone has been motionless in Alice's office for three hours," then both Alice's work persona and a limited range of activity have been disclosed with a high probability. Incidental disclosure of persona and activity is exemplified by the credit card purchase, which, in addition to disclosing the primary content (money) to the merchant, incidentally discloses a persona (*i.e.*, a name and credit card number) to the merchant and activity (*i.e.*, a particular purchase) to the merchant and the bank.

The implications of *surveillance vs. transaction* on primary and incidental disclosure are nuanced. One way to conceptualize them is as follows. From the subject's perspective, disclosure through surveillance is primarily incidental; being surveilled is not the subject's primary concern. But from the observer's perspective, surveillance does not distinguish content; it is all primary. Transactions, however, as exemplified by TGI, disclose both primary and incidental content, and each is reasonably distinguishable by both subject and observer.

With respect to *interpersonal vs. institutional disclosure*, each can occur through both primary and incidental disclosure, but institutions, through TGI, seem to make significant use of incidental disclosures. Notably, incidental institutional disclosure is often a byproduct of a primary interpersonal disclosure, *e.g.*, sending an email to a friend discloses part of your social network to your ISP. Conversely, incidental interpersonal disclosure can be a byproduct of primary institutional disclosure, *e.g.*, a spouse might view the subject's purchasing activity as recorded in a shared financial account.

CLASSIFYING EXISTING PRIVACY PHENOMENA

So far we have presented a conceptual argument, interspersed with some useful examples. To further concretize the discussion, Table 2 classifies some existing privacy-related phenomena into the privacy space. We hope this helps illustrate the origins of differences between the privacy implications of, say, P3P and camera surveillance. In keeping our descriptions general, both categories of a given dimension will often apply to a phenomenon (*e.g.*, camera surveillance can disclose both persona and activity), but for any *specific instance* of a phenomenon, analysis can clarify the role each category plays in the disclosure.

CONCLUSION

Privacy resists definition, and attempts at definition inevitably lead to contention and confusion. This is because privacy is neither a descriptive attribute of nor an applicable function of a phenomenon. It is a value process whose properties and operation vary with the natures and relations of the systems, actors, and information at play.

In this paper we have begun to deconstruct the privacy space. We have presented a set of dimensions that highlight the conditions that create the privacy implications of any given privacy-related phenomenon. Categorizing a phenomenon into this space can expose the qualities that shape privacy's role therein. We hope this helps to stabilize discourse on privacy and to focus the design of privacy-related systems.

Acknowledgements

We thank danah boyd, Jason Hong, Bill Press, Chris Beckmann, and Natalie McGee. This material is based upon work supported by the National Science Foundation under Grant No. IIS-0205644. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

REFERENCES

1. Adams, A. and Sasse, M.A., Taming the Wolf in Sheep's Clothing: Privacy in Multimedia Communications. in *Proceedings of ACM Multimedia 1999*, (1999), 101-107.
2. Agre, P.E. Changing Places: Contexts of Awareness in Computing. *Human-Computer Interaction, 16* (2-4). 177-192.
3. Agre, P.E. Surveillance and capture: Two models of privacy. *The Information Society, 10* (2).
4. Alderman, E. and Kennedy, C. *The Right to Privacy*. Knopf, New York, 1995.
5. Altman, I. *The Environment and Social Behavior: Privacy, Personal Space, Territory, and Crowding*. Brooks/Cole Publishing Co., Monterey, CA, 1975.
6. Bellotti, V. and Sellen, A., Design for Privacy in Ubiquitous Computing Environments. in *Proceedings of the Third European Conference on Computer Supported Cooperative Work (ECSCW'93)*, (1993), 77-92.
7. Brunk, B.D. Understanding the Privacy Space. *First Monday, 7* (10).
8. CDT. Generic Principles of Fair Information Practices (http://www.cdt.org/privacy/guide/basic/generic.html), Center for Democracy and Technology, 2000.
9. Cranor, L., Langheinrich, M., Marchiori, M., Presler-Marshall, M. and Reagle, J. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, World Wide Web Consortium, 2002.
10. Cranor, L., Langheinrich, M. and Marchioro, M. A P3P Preference Exchange Language 1.0 (APPEL1.0), World Wide Web Consortium, 2002.
11. Foucault, M. *Discipline and Punish*. Vintage Books, New York, 1977.
12. Goffman, E. *The Presentation of Self in Everyday Life*. Doubleday, New York, NY, 1956.
13. Goldberg, D., Nichols, D., Oki, B.M. and Terry, D. Using collaborative filtering to weave an information tapestry *Communications of the ACM*, 1992, 61-70.
14. Lessig, L. *Code and Other Laws of Cyberspace*. Basic Books, 2000.
15. Palen, L. and Dourish, P., Unpacking "privacy" for a networked world. in *CHI 2003*, (Fort Lauderdale, FL, 2003), ACM, 129-136.
16. Phillips, D.J., Context, Identity, and Privacy in Ubiquitous Computing Environments. in *Ubicomp 2002 (Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing)*, (Göteborg, Sweden, 2002).
17. Samarajiva, R. Interactivity As Though Privacy Mattered. in Agre, P.E. and Rotenberg, M. eds. *Technology and Privacy: The New Landscape*, The MIT Press, Cambridge, MA, 1997.
18. Whitten, A. and Tygar, J.D., Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. in *8th USENIX Security Symposium*, (1999).

Table 2. Existing privacy phenomena classified into the privacy space

Phenomena / Dimensions	P3P [9]	Fair Information Practices [8]	Online Chat Rooms	Automated Location Disclosure to Boss	Window Shades Facing Domestic Neighbors	Camera Surveillance
Feedback And Control	Largely feedback, though increased deployment of the APPEL preference language [10] could lead to increased control	Both	Feedback (*e.g.*, displayed list of other personae in room, displayed history of content) and Control (*e.g.*, recourse to official moderators)	Depends on implementation; arguably both are limited.	Control	Possibly some feedback
Surveillance vs. Transaction	Transaction	Largely transaction-oriented, though inclusive of surveillance	Both	Likely surveillance, but disclosures are probably generated in a transactional manner	Surveillance	Surveillance
Interpersonal vs. Institutional	Institutional	Institutional	Interpersonal	Both	Interpersonal	Institutional
Familiarity	Variable	Variable	Variable	Bilateral	Relatively bilateral	Variable
Persona vs. Activity	Both	Both	Both	Activity	Both	Both, though typically an anonymous persona
Primary vs. Incidental	Both, with notable emphasis on TGI	Both, with notable emphasis on TGI	Both (primary to intended observers, incidental to surveillers)	Depends on implementation, though likely incidental from subject's perspective, primary from boss's	Likely incidental	Incidental from subject's perspective, primary from observer's