Designing Mobile Awareness Cues

Antti Oulasvirta
Helsinki Institute for Information Technology HIIT
Helsinki University of Technology and University of Helsinki, Finland
aoulasvirta@acm.org

ABSTRACT
This paper considers how we may design future mobile awareness systems. Building upon research on social cognition, we suggest the need to take into account what is known about humans' interpretational capabilities. We identify design issues from the level of an individual awareness cue to the level of a product concept, systematically exposing the associated solution spaces. Using four real applications as analytical examples, we point out multiple ways in which design can affect the user's processing of awareness information and thereby yield different outcomes in the use of technology. We conclude by pointing out novel design opportunities that lie in the integration of cues with functionality and content on the mobile phone.

Categories and Subject Descriptors
H.5.3 [**Information interfaces and presentation**]: Group and Organization Interfaces – *asynchronous interaction, synchronous interaction, theory and models.*

General Terms
Design, Human Factors.

Keywords
Mobile awareness systems, awareness cues, social inference, interface design, mobile applications.

1. INTRODUCTION
Present-day mobile devices include hardware and software sensors that were not originally designed to produce data for users' consumption but rather for computers. GSM cell ID and keyboard events are prime examples. *Mobile awareness systems* is a subcategory of context-awareness applications that builds on the idea of leveraging these and other sensor data by capturing, reprocessing, distributing, and representing them for users.

Reference systems in this area include ConNexus [24], which presents location and other cues in an address book, Jabberwocky [20], which presents traces from encounters of familiar and unfamiliar passers-by in a city, iCAMS2 [12], which presents locations, proximities, and movement directors of friends in a city, and AwarePhone [1], which presents locations and statuses of a medical team. Recently, products services have emerged that employ awareness information. One of them is Jaiku (see Jaiku.com), which is a spin-off from our earlier work [15]. Interestingly, no mobile awareness system is like another, but there are salient differences in how awareness information is presented and can be interacted with. Nevertheless, there are no articles that systematically analyze recurrent design issues in this domain.

This paper's primary goal is to go beyond individual applications and identify design issues that recur from one application to another. We start by reviewing a theoretical approach to cue design that builds on work on social cognition. The premise of this approach is that the user's "awareness" is based on individual acts of inference of awareness cues. Design should build on scientific knowledge on the capabilities and limitations of human social inference. The contribution that this approach makes to design is based on two ideas: 1) breaking down awareness information to constituent *awareness cues* and 2) analyzing individual *acts of inference* as constituents of the user's cognitive state that is called awareness.

We will examine four applications in detail: two versions of ContextContacts, Jaiku, and CoMedia (Table 1). They will be used as examples in analyzing a host of design questions:

1. *Design of an individual cue* (Section 4): Which mode of presentation is the best? How abstract should the level of description be? What to do with outlier data? Should one use imperative or neutral cues? Are predictive cues better than descriptive?
2. *Cue layout* (Section 5): Are there degrees of freedom in how cues can be laid out on a screen and across multiple screens? One cue or many? How to design access to a cue? How to support mobile use by means of layout?
3. *Integration* (Section 6): Should cues be integrated into other functionality and content and, if so, how?

For most questions, no single "best" recommendation exists, but choices are a matter of striking a balance between pros and cons. We will try to sensitize the reader to the implications of one's choices and clearly state when our claims are not backed up by empirical evidence. Obviously, this is a field where empirical work is scarce. Due to limited experience with, for example, auditory cues, we have chosen to limit our discussion to the visual modality. Moreover, we will not discuss the possible types of cues, for the number of viable types is constantly growing as new sensors are added to mobile devices. The issues we address should in principle be relevant to all types of cues from emotional cues to location cues.

MobileHCI 2008, September 2–5, 2008, Amsterdam, the Netherlands.

Table 1. The four awareness applications analyzed in this paper.

Application	Main awareness cues	Integration of cues
1. ContextContacts v1, without message line	Multiple (Figure 1)	To contact information and calling functionality on the phone
2. ContextContacts v2, with message line	Same as #1, but emphasizes online status, location	Same as in #1. In addition, creates a group messaging space
3. Jaiku	Location, online status; other cues secondary	Primarily to microblogging, secondarily to calling and SMS on the phone
4. CoMedia	Location, online status, co-presence	To group media sharing

2. APPROACH

Research on *awareness* started in the beginning of 1990s. An often cited definition comes from a study of Portholes, a system that conveyed snapshots and user-controlled availability cues for workers in two distant offices [3]. Awareness was defined as the understanding of a remote person's actions that contextualizes one's own activities. The term *awareness cue* is of newer origin, introduced by Greenberg and Gutwin in an article discussing the design of groupware [7]. In parallel to CSCW and HCI that have utilized the notion of awareness, the term *social cue* was adapted to computer-mediated communication studies by Sproull and Kiesler, who examined how the lack of such cues can lead to misunderstandings in email-based communication [23]. This branch of research borrowed ideas from social cognition research and proposed principles of how technology-mediated cues are interpreted. (See [18] for a closer analysis of differences between previous positions and the one advocated here).

For the purposes of the present article, we define *awareness cues* as representations of whereabouts, actions, and situations of remote people. Awareness cues have the following properties:

1. they are computer-produced, as opposed to user-controlled cues such as presence lines
2. they are "live", i.e. automatically updated in real-time
3. they are symbolic or iconic by type, not analogous such as video or audio.

This paper advocates a particular social cognitive stance toward awareness cues. The *problem of social inference* is how people make inferences of each other, successfully and at times unsuccessfully, from the pool of information they have at their disposal. Technically, social inference refers to the premises, procedures, and outcome of making inferences of another person based on cues (any cues, not just awareness cues). The cues that a piece of awareness systems can provide are, in comparison to face-to-face situations, relatively incomplete, uncertain, and poor. Keeping in mind the fundamental limitations of human cognition, one needs an explanation of how they can form a reliable basis for social interaction.

According to the findings of social cognition research, people do not thoroughly evaluate all available information, as implied by normative theories of rationality, but apply shortcuts and simple rules to overcome these limitations. These rules draw heavily from pre-knowledge about the other person. They enable "jumping to conclusions" and arriving at interpretations that go beyond the literal meaning of a cue. The downside is that biases and errors are bound to occur, and effort is needed to turn or override routinely produced interpretations. Heuristics and pre-knowledge assume center stage in this paper when we analyze how design choices affect social inference.

The details of this approach have been described elsewhere [18], we do not repeat the framework in full length. Furthermore, it is not the task of this paper to evaluate this approach. Nevertheless, it is worth keeping in mind that some important design issues are not in its scope. For instance, it cannot explain how privacy is socially negotiated, how group dynamics affect the use of cues, or how cues are used to actively express oneself to others.

To sum up, the use of mobile awareness boils down to interpretation of awareness cues which has four elements: 1) processing goals triggered by the situation at hand, 2) relevant pre-knowledge, 3) used heuristics, and 4) the availability and design of awareness cues.

3. ANALYTICAL EXAMPLES

Table 1 summarizes the four applications we use as analytical examples in this paper. While many of the cues that appear in these applications are computationally analogous, the way designers have contextualized them influences how users process and act on them. The four applications differ at all three levels of design examined in this paper: an individual cue, layout of multiple cues, and their integration to functionality and content.

3.1 ContextContacts v1: Augmented Contact Book

ContextContacts v1 was a multi-cue awareness system designed for Series 60 Nokia smartphones. It replaced the native contact book with a one augmented with awareness cues [16], but retained its standard functionality. The cues were in this version were:

1. *Online status:* "the hand" icon, which turns red when the user is using the phone and grays out gradually with time
2. *Location:* district labels automatically mapped to GSM cell IDs; overridable by users
3. *Duration of stay:* shows for how long the user has stayed in that location (in parentheses next to location)
4. *Alarm profile:* two icons to show if audio and tactile alarm are on or off
5. *Number of people in proximity:* the yellow man silhouette together with an indication of count; inferred from presence of Bluetooth phones
6. *Number of friends in proximity:* the green man silhouette; the number of people present who are in the contact book.

A detail view provided textual descriptions of cue contents for any selected user.

All cue-augmented contacts were always grouped to the top of the list, making quick skimming easier in a long list of contacts.

Figure 1. A version of ContextContacts v2 with a user-controlled message line.

3.2 ContextContacts version 2: Mobile IM

Figure 1 presents ContextContacts v2. The former version (v1) included all the same cues as v2, but did not have the message line shown next to the contact's name. Parenthetically, we had two additional cues in a later version (v3): a calendar cue (current and next calendar event from the phone's calendar) and a computer environment cue (indicates with an icon if the user's own laptop, desktop PC, or PDA is within the Bluetooth range). However, we did not utilize those cues in the field studies presented in [15] so we sustain from their analysis in this section.

Although the only difference between v1 and v2 is the message line, its presence has actually so significant implications to users practices and interpretations that we decided to distinguish the two. We return to this matter in Section 6.

3.3 Jaiku: Microblogging for Mobile Clients

Jaiku.com is a spin-off company from a former project in our institute. It builds on the technology underlying ContextContacts. Jaiku is a microblogging service that enables posting brief messages to friends' phones or the Web from the mobile client. Awareness cues are in supportive role here. The mobile client splits the cues into two screens:

1. The main screen which shows the Online status cue and location cue next to the "jaiku" line (Figure 2, left)
2. A detail view that shows more cues on a selected user (Figure 2, right).

Figure 2. Jaiku is a microblogging service with awareness cues. Jaiku's main view (on left) and detail view (on right).

3.4 CoMedia: Awareness Cues in a Group Media Application

Of the four examples, CoMedia [9] is the most complex one and requires somewhat more space to describe. Comedia was inspired by ethnographic studies of spectator groups at large-scale events [8]. It provides users the means to share text, video, audio, and images within a group. Its interface is built around three *perspectives* to cues and content:

1. *Media Stories*, for creating and sharing stories together
2. *Member List*, person-based view to awareness cues
3. *Event Pamphlet*, shows event information and functionality to mark subevents for the group.

To enable a quick change of perspective, the user can move between these sections by pressing left and right on the phone's joystick. Navigating left, the user finds the Event Pamphlet that contains a schedule and next-to-real-time content. Navigating right, the user finds the Member List that is augmented with information as to the people's current locations and activities. Because the Member List is similar to ContextContacts—except for the missing communication functionality and fewer cues—we here focus on awareness cues in the Media Stories view.

A *Media Story* (Figure 3, left) organizes the group's communication into a dedicated discussion space. For each story, it shows the title, the time of the last post, the poster, including the poster's current status, the number of people who have viewed the story, and the number of messages. If there are multiple users, a number appears next to the icon, giving information about current activity inside the story.

The *Story View* (Figure 3, on right) lists the messages in the Story, with the most recent posts on the top. On the same line as the message text there are member icons representing the number of people viewing that message, analogous with the Media Story view. A scrolling ticker at the top of the view shows the names of the people viewing this story at the moment.

The *Message View* presented in Figure 4 shows the contents of a single message. The media is automatically augmented with cues about who were present (Bluetooth-based inference) when the message was written. This hints of who might have contributed to the message.

Figure 3. CoMedia's view to Media Stories (on left). Upon selecting a Story, the user can see media in that Story. Awareness cues indicate whether other users are online and if they in proximity versus distant.

Figure 4. CoMedia's message view shows in textual format who sent the message and who was present at the time of creation (on left). The message viewings view allows following who has seen a message (on right).

4. CUE DESIGN: REPRESENTING REMOTE STATES OF AFFAIRS

In this section, we focus on the design of an individual awareness cue.

4.1 Modality: Icons versus Text

The majority of awareness systems present awareness information either as text or icons. The choice between the two brings about four considerations:

First, icons take up less space than text.

Second, the expressiveness of small icons is more limited than that of text. For example, while one can come up with icons for a couple of different categories of locations like "at home" or "at bus stop", covering all meaningful places with distinctive icons is practically impossible.

Third, the categories underlying a cue may be less understandable with iconic representation. For example, the design of CoMedia assumes that the users learn that the blue person icon means that that person is within Bluetooth range. On the other hand, if high accuracy in interpretations is not a design goal, ambiguity can be used deliberately as a means to encourage reflection. If an icon shows a face sweating, does that mean that that person is exercising, in a difficult social situation, or taking a shower?

Fourth, visual saliency can be leveraged to support quicker search. For example, distinctive color and shape are related to attentional pop-out. Ideally, the user can check with one quick glance the group's status and note if changes have occurred, without reading through contacts exhaustively one at a time. Well-designed icons can promote the formation of visual saliency based habits.

In ContextContacts v1 and v2, we have attempted to strike a balance between these concerns by favoring icons in the main screen, except for the location cue which is textual. However, to support learning of these cues, all icons are explained in text in the detail view accessible with one key press from the main view. CoMedia, on the other hand, has icons only for the online status of a user and recent use of phone while all other cues are presented in text. This solution was preferred because there were many screens and some cues were specific to a particular screen.

4.2 Level of Description

The second cue-level issue concerns the level of description. For example, a person picking up and throwing a brick can be described at the lowest level of the movement of limbs or at any level all the way to the highest level of hooliganism or crime. While any action or event can be described at multiple levels, location cues are particularly problematic in this sense. For example, GPS enables geographic accuracy of about 1-5 meters, but it is not clear if such accuracy is needed.

We propose that the most useful level of description is the level that is the most diagnostic of situations that are of interest to the users. For example, soccer coaches monitoring players are not interested in the information that a player is running, or that a player is located on a soccer field—those are already known. Optimal awareness cues, for a coach, are related to his/her actions and decisions; the best cues can augment what he/she already knows. For example, information about a player's tendency to miss passes, or the average pulse level being over 170, may be of more interest. Such cues are diagnostic to the extent they help the coach to decide whether to keep that player on the field or call in a substitute.

Accepting this view implies that there is no universal solution; the level must be determined case by case by carefully considering the activities of the user group. This is complicated by the fact that almost all group activities involve moments when lower level information is needed. In coordination of mobility, for example, lower level feedback becomes increasingly more important as the monitored person comes closer and closer to the target.

If the application has access to both low and high level cues, this situation may be reconciled. Such application could provide a medium level of description as the default view and provide other levels as options. However, to our reading, such interfaces have not been presented.

4.3 Dealing with Outlier Data

In almost all social activities, some actions are more informative of collaborators' intentions than others. Most incoming information may actually be irrelevant, like noise, to a perceiver waiting for a particular signal. Consider for example stopping by at a kiosk on your way to work. Is this sidestep worth conveying to your co-workers or would it be more informative to simply state that you are on your way to work?

There are three somewhat orthogonal approaches to solving this problem. Again, the decision among these three must be made taking into account the pros and cons of each.

1. Omitting "outlier" information. ContextContacts, for example, omits locations of infrequently visited places, basing on an algorithm that keeps track of the most frequently visited GSM cell clusters [10]. When the user is in a new cell, a question mark is shown to others to indicate that the user is moving. However, the last known base is also presented (in text) to support understanding where the person is coming from. Omission of infrequently visited places focuses inferrers' attention to landmarks, the obvious drawback being that significant but rarely visited places are not conveyed at all.

2. Discretization. One can also discretize cues based on what is known about user needs. This hides unnecessary detail and can highlight important ones. For example, after we had used Con-

textContacts v1 for several months in our own group, we discretized the duration of stay cue so that the first hour was presented with an accuracy of a minute, but from thereon simple categories like "> 1h" or "> 1d" were utilized. This non-linear discretization seemed to work in our own activities relatively well.

A similar solution can be entertained for location information as well. For example, a group at UC Irvine has explored the possibility of sensor based recognition of bus stops and parking lots for the purpose conveying that information as a status cue in a mobile IM client [19]. Conveying only the nodal events in commuting is a significant simplification of GPS data but its utility may be limited to fewer activity domains.

3. Highlighting outliers. If outliers are infrequent but associated with abrupt changes in context information—for instance, as in the case of going abroad—the system could highlight these changes. This is a speculative option, however; to our reading, no machine inference solutions have been implemented in mobile awareness systems.

It is worth keeping in mind that adaptive solutions should not override default cues. With time, users develop a set of cue-checking habits that rely on subsets of cues [18]. An adaptive solution that displaces sought-for cues may hamper the development and execution of such habits. One possible solution is to render outlier information to a secondary status in the cue interface. For example, one can place a simple but salient indicator icon in the main cue view to denote that something extraordinary has taken place; the referred-to content, the outlier information, can be accessible in a secondary view.

4.4 Actionability

A cue can take an imperative form, a prompt to act in a certain way. Consider an availability cue implying "Do not call me now, call me later." A cue can also be designed to convey the present state of affairs as neutrally and veridically as possible, such as a calendar-based cue telling that "I'm in a meeting for the next 15 minutes." The former is called a *prescriptive* cue and the latter a *descriptive* cue. The former guides the user to a certain course of action while the latter is neutral in those terms.

Which solution should be preferred? With pre-knowledge, the inferrer can make surprisingly deep interpretations based on descriptive cues. All other things equal, a descriptive cue is associated with a broader range of potential interpretations. This reveals a trade-off: a descriptive cue fosters interpretive flexibility [14], but does not direct the consequent actions as much as the designers may want. This may be important in systems that designed for a particular purpose.

A related issue is privacy; descriptive cues, because of interpretive flexibility, run the risk of affording intrusive interpretations. Prescritive cues, by contrast, reveal minimum or no information, assuming that the exact logic that produces these cues is not known to the user. Moreover, plausible deniability is not such a problem with prescriptive cues as it is with descriptive cues. However, if the computer's logic is unambiguous and the users know it, they can in principle retrospectively infer the underlying low level information. Designers worried about breadth of disclosure or plausible deniability should think not only about the design of cues but the transparency of the logic that produces them.

4.5 Temporal Reference

As any proposition about the world, an awareness cue can refer to the past, the present, or the future. Cues that refer to the past include, for example, previous places visited or the last time email was checked. Cues that refer to the present include, for example, current place and number of unread emails.

All cues meant to depict the present are inherently out of date when the inferrer sees them. This is due to lag caused by network transfer and sampling intervals in the sensors. However, developers cannot set arbitrarily high intervals due to battery drain and data costs. A sampling interval must be chosen.

We believe that for most cues, an interval somewhere between 15 seconds and two minutes should be acceptable, but not unequivocally. Although we have employed a universal update interval in ContextContacts, we learned that different cues require different update intervals. For example, an interval of one minute may be too coarse for the online status cue, because it creates an illusion of availability also in case of very quick checks. Many forms of turn taking require close to real-time knowledge of the other person's availability and responsiveness. Too long intervals lead to erroneous social inferences and should be avoided.

The desired interval also depends on the granularity of the cue. An interval of one minute is tolerable for a location cue that represents movement at the level of a district, but not for a cue with an accuracy of an inch.

Cues that refer to the future are, by definition, predictions and, contrary to the other types, inherently uncertain. To our knowledge, no predictive cues have been implemented although there are a few papers discussing their technical feasibility. For example, previous work has looked at the prediction of users' next location. A predictive accuracy of 60–80 % can be achieved, depending on idiosyncratic variability of routes [10]. In another paper, we showed that a mobile user's time-sharing pattern can be predicted with an accuracy of about 70–75 %. (Time-sharing is a behavioral measure of the user's deployment of attention; for example, how much the user will look at the phone versus the environment during the next 30 seconds [11]). It is an open question how large predictive inaccuracies are acceptable. If reliable, predictive cues could support such coordinative activities where users base their actions on predictions of others' actions. Reliable advance cues could help streamlining many performance-oriented and competitive situations that involve this characteristic.

5. LAYOUT AND ACCESS

In this section, we discuss issues that go beyond the scope of an individual cue: whether to present single or multiple cues, how to lay them out on a screen, how to make them understandable, and how to respond to the multitasking requirements of mobile use.

5.1 Number of Cues

There are systems that provide only a single cue (typically location) and there are systems that provide multiple cues simultaneously, such as ContextContacts. Again, we come across the issues of screen space and interpretive flexibility.

Data from our studies with ContextContacts suggest that each cue affords quite different patterns of interpretations. Intuitively, other things being equal, having more cues should improve interpretive flexibility, which in turn should imply a wider set of uses for the system. If many cues are available, users can pick the ones

deemed most relevant for the situation at hand. This is a strong argument for multi-cue systems. However, not all cues are equally important, in our studies there were cues that had more uses than others [15]. The location cue and "the hand" cue, for example, were used more frequently than others. By running user trials, designers can learn which cues are of secondary importance and can consider dropping them to save screen space.

Beyond the obvious issue of many cues taking up more space, having multiple cues introduces three difficult problems: 1) privacy, 2) disinterest, and 3) overload. Users may not want to disclose as much information as the system requires. Moreover, if some users are more interested in the cues than others, an asymmetrical interest relationship emerges. For example, in a family of one mother and three teenager boys, the boys were not generally interested in the mother's awareness information, whereas the mother was interested in boys' information [15]. This exemplifies what Grudin [6] has called "the discrepancy between those who do the work and those who reap the benefit." Having more cues is not necessarily better; it will change the social nature of the application. Moreover, having multiple cues may create a need for subtle control of disclosure. In ContextContacts, we have enable a simple reciprocal give all/get all mechanism for disclosure control, but the limits of that rough solution may become apparent in privacy-sensitive settings.

In some settings, only one or two cues may be enough. In an unpublished study of ContextContacts v2, we found that 60 % of mobile information workers' interpretations are covered by only two interpretations: availability and interruptability. Soon, with improving machine learning mechanisms, it may be possible to provide those cues directly instead of providing many lower level cues.

5.2 Order

Designers have great influence over the order in which the cues are processed by the user, which in turn will affect the outcome of interpretations. Classic studies in decision-making have provided evidence for the existence of an inference bias called *anchoring* [26]. When multiple pieces of information are presented, people will not use all available information but tend to pick one or two cues as primary cues and construct an initial hypothesis in the light of which all subsequently processed cues are interpreted. The subsequent cues typically confirm the initial hypothesis, sometimes modify it, but rarely lead to its rejection.

In a laboratory study utilizing mobile awareness cues as experimental materials, we found dramatic changes in perceived utilities when the order of presentation of a cue pair was reversed [18]. In the experimental paradigm, the participants were first given one cue and asked to guess the situation of another person based on that cue. To indicate their confidence, they placed an imaginary bet on the guess. Then another cue was given and the participant was asked to update the guess and the bet. All combinations of six cues were tested like this. For example the pair "Turnover of BT devices on range" (e.g., "20 devices recently") and the phone manipulation cue received a perceived utility of 0.34. However, when their presentation order was reversed, the perceived utility doubled (0.74).

This phenomenon can be exploited by means of design. In ContextContacts initial design, we (admittedly, accidentally) emphasized the location cue and the online status cue by placing them to the *left hand side* in the interface. Users typically read the contact list top-to-bottom and left-to-right, and are thus more likely to fixate on cues on the left hand side. Indeed, these two cues were the ones that saw the broadest repertoires of interpretations in our field studies [15]. Another ordering on the screen could have yielded other kinds of interpretations and, consequently, other uses of the system.

Other ways to guide the order of processing include using visual saliency. Of course, this can work counter to the designer's original intentions. Careless coloring of such cues that are of secondary importance may draw users' attention to those cues first, with the unfortunate consequence that they anchor on those cues and neglect more critical information.

5.3 Division across Screens

There are two logical alternatives in distributing awareness cues: making all cues visible on the same page and *pagination*—splitting the set of cues in several views.

ContextContacts does not paginate cues but presents all of them (except the "previous location" cue, which is available in the detail screen) in the main display. This solution is preferable because all information is directly available; it also supports the formation of cue-checking habits.

The Jaiku mobile client, on the other hand, uses pagination in the following way: in the main view, all contacts are presented with a location cue and an online status cue next to the person's name and her latest message. All other cues, which are available if the user has decided to disclose them, must be accessed from a detail view. The splitting of cues across screens means that, most likely, users often make first interpretation based on the first screen and then, if they look at the detailed view, mainly update or modify their interpretations. Or, if users are overwhelmed with these cognitive demands, they limit their intake of cues solely to the main view. Consequently, interpretive flexibility is compromised, because the secondary cues are cognitively more distant from the initial screen and because anchoring on only two cues narrows the space of interpretations.

ContextContacts' solution is not perfect, either, in this respect. ContextContacts can only present three contacts on a single display. This means that when a group is chatting, not all messages can be seen at the same time; for comparisons the user has to keep in mind non-visible information.

The clustering of group members in ContextContacts and Jaiku is arbitrary in the sense that its listing of contacts is based on alphabetical order. A more meaningful way to cluster group members, particularly in a larger user group, can be based on social network analysis (SNA) metrics. For example, the PASION (an on-going EU FP6-IST project, www.ist-pasion.com) knowledge work prototype uses an atom metaphor where the user is in the middle, the nucleus, and co-workers are on the orbits. Contacts are placed in close or distant orbits according to SNA metrics such as the frequency of interpersonal communication. Thus, the only stable anchor in the display is the user herself, all others are displayed adaptively as response to communication patters. The downside is that users may have to pay more attention to find the sought-for contact when the layout suddenly changes.

To sum up, the issue of division is a matter of two trade-offs: 1) screen estate versus interpretive flexibility and 2) guidance (to particular kinds of interpretation) versus formation of habits.

5.4 Inspectability

The inspectability (or "scrutability") of a computer's interpretation is important in all ubicomp applications that leverage complex machine inference [2]. For example, CMU's "interruptability" signs for office doors not only show Bayesian interpretations of the office dwellers' interruptability; they also show the confidence levels of those estimations with the three most important sensors on which they were based [25].

Inspectability is beneficial in two pursuits: 1) finding out more information about another person and 2) learning about the computer's inference logic. The roughest form of inspectability is represented by a cue that shows when awareness cues are *not* up to date. For example, if a user does not have a wireless connection, that can be shown to others with a cue. In ContextContacts, a contact was grayed if for some reason he/she could not be updated. When the user herself could not contact the network, all contacts on the display were grayed. To our reading, more sophisticated solutions for scrutability have not been presented for mobile awareness systems. Generally speaking, it may be natural to prefer those forms of machine inference that are understandable by humans, such as case-based reasoning and Bayesian networks.

5.5 Design for Mobile Use

As a final topic that is somewhat an outlier here, but nevertheless of relevance, we raise the question of design for mobile use.

Particularly in larger user groups, awareness cues may change frequently; the frequency of updates is a function of the number of cues times the number of contacts. Receiving an audible notification for all status changes would result in a beepy nightmare. On the other hand, some status changes are of relevance to the users. There are three alternative solutions to this problem: 1) make the system "pull only"—no notifications and 2) let users mark certain contacts as interesting and deliver notifications selectively. 3) A third option, which is as speculative as the second one, is to construct an algorithm to notify when somebody is doing something out of the ordinary.

Besides supporting users' awareness over changes, the system needs to support rapid access to the cues. This is a clear implication from experiments on mobile users' cognitive resources. Research has indicated that, at worst, mobile users can attend their devices for couple of seconds at a time [17]. This implies that accessing the primary cues should not take more than a few seconds from anywhere in the device's application space. ContextContacts succeeded to support access in less than 3 seconds from the standby screen of the phone (calculated for a list of 6 cue-augmented contacts). All middle-steps like startup screens and logins were eliminated to streamline access.

Data from our field studies [15] indicate that about 5-15% of ContextContacts' use took place when mobile. With even faster access, the proportion of mobile use could increase. For this end, one can consider "preview cues" in the phone's standby screen, cue-augmented screensavers, and vibrotactile notifications.

6. INTEGRATION

Integration of cues is an important design topic and one of the less exploited opportunities in mobile awareness research. The topic of integration expands our analysis to two new levels: the level of application and the level of personal content. In practice, integration means two things:

1. *Positioning* of cues close (visually or interactionally) to content in an application. Content can be something as simple as the name of a contact (as in ContextContacts) or a more complex entity like a posting (as in Jaiku) or media (as in CoMedia).
2. *Associating cues to communication functionality* such as SMS, calling, logs, or replies.

From the perspective of social cognition, the role of integration is in the guidance of users' processing. Integration guides the user's attention to process particular pieces of content alongside with the cues, but it also primes the processing goals with which users approach the cues. By means of integration, designers can interweave the awareness cues with action and sense-making.

Before discussing three examples, it is worth mentioning that "stand-alone" cues can be powerful as well; cues *need* not be embedded anywhere. Stand-alone cues do not direct user's attention too much to a given activity or meaning, but may leave more freedom for interpretation. At the same time, stand-alone cues are detached from actions wherein designers may want them to participate. In the worst case, this may render the system as an opaque "panopticon-like" tool for monitoring others. In our studies, we find little support for stand-alone solutions.

In the following, we discuss the effects of integration. ContextContacts was studied in a set of three longer-term trials with teenagers; high-granularity logging and interviews were used for data collection [15]. An A-B intervention methodology was utilized where we first collected a baseline level of observations during period A, to which we compare results gathered in period B. CoMedia was studied at two large-scale events [9].

6.1 Integration to Contact Book

In the 2005 paper that reported the first version of ContextContacts [16], we wrote that our aim is to solve the problem of "acontextuality" of phone calls: a caller looking at the standard contact book does not have any information on the callee. According to some sources, the success rate of mobile phone calls is notoriously low—more than one third of call attempts fail to reach the intended recipient.

In a set of trials of ContextContacts v1 [15], which integrated cues to the contact book, the collected logs revealed that users tend to look at the cues for a time period of 1 to 4 seconds before placing a call. By comparison, when they did *not* have the cues, 60% of calls were placed almost immediately—within one second after moving the selector to the contact. During this brief period of time, users looked at the cues and made their decisions on whether to call or not. In one group that used the system very frequently, the proportion of successful call attempts increased by 12 percentage points due to the introduction of cues. Thus, placing cues close to communication functionality had a tangible, albeit small, effect on success rate. These inferences mostly relied on the "hand cue," the location cue, and the alarm profile cues; the key inferences being availability for communication, interruptability, and responsiveness. Part of this coordination was of the compensatory type; for example, looking at cues to understand *why* a recent call attempt was rejected or not answered.

It is intriguing that the availability of a location cue in ContextContacts v1 did not seem to release conversants from telling each other their location in the beginning of phone calls [15]. Conversation-analytic studies of mobile phone calls have pointed

out that the disclosure of location serves two functions in interpersonal interactions: First, it provides a basis for mutual communication and the planning of activity, and second, it facilitates the building of interpersonal understanding and trust. In studies where we have measured location disclosure, we have found that mobile awareness do not make these functions obsolete. However, in one user group, the cues frequently directed the conversants to disclosure location *with higher granularity* (e.g., "I'm in the computer room" instead of "I'm in school") than when they did not have those cues.

To understand the contribution of integration, it is useful to look at the kinds of uses that emerged that were *not* communication-related. We observed two main categories of that kind: 1) coordination of mobility and 2) companionship. Coordinations of mobility involved inferences of place, proximity, movement, and activity. Here, automatic cues were of decisive importance because they could be relied on as being updated and timely, and thus provided a reliable resource for moment-to-moment coordinative efforts. Through automatic cues, the participants initiated face-to-face meetings when others were seen to be close by—and knew when to avoid certain actions. Cues were not only used in a compensatory way (e.g., to explain why a person is late after an event has unfolded); they were also used to generate *anticipations* of future events—of what the other will do in a given situation. Users also monitored the progress of others in agreed-upon group rendezvous.

Awareness cues also supported a feeling of mediated companionship, mainly through their use as weak "signals" of another's situation and presence. Awareness cues became a sort of proxy for another person, particularly in the way they can act and be used in the place of a distant person, like having that someone somehow "with you." However, this use was much more pronounced when the message line was introduced (Section 6.2). The role of integration with messaging lies in the fact that messages provide a more controlled means of expression and are thus a better resource for reciprocally deepening companionship than automatic cues. Parenthetically, we believe, although there is not much direct evidence in our data to back this up, that in this kind of use the cues are approached in a more exploratory way than when using them for coordination or expression. The goal is not so much to decide between a few alternative inferences (e.g., available or not) but to construct a more holistic representation of the other so that he or she can be felt as "present." A secondary role of such states may be that they inspire informal encounters and interactions.

These two categories of use do not rely on integration with communication functionality although they benefit from it indirectly. If coordination or companionship related inferences lead to action, they are likely to lead to communication attempts.

6.2 Integration to "Mobile IM"

ContextContacts v2 introduced a message line to the main view, allowing users to chat in the space of the group's awareness information. Over the period of three weeks, over 5000 message were sent in a teenager group [15].

We believe that two features of ContextContacts v2 contributed to the phenomenon that could be characterized as its "emergence as a "locale"—"a digital place that offers a group the site and means for maintaining awareness of another and for rapidly moving into interaction" [4]. First, the grouping of the cue-augmented contacts at the beginning of the contact book allowed for quick access to the situations and messages of the group. Second, the automatic cues, particularly the hand cue, allowed for presence, availability, and responsiveness inferences, which have been observed to be important in all asynchronous communications (see e.g., [13]).

Thus, this medium puts user-created content on center stage, while relegating automatic cues to a secondary, supportive role. The hand cue was in effect appropriated for understanding and negotiating when to send a message so that others will see it, to infer conversational availability, to check to see who is "online," to estimate the rapidity of response to one's own turn, to infer if others have received a message after it was sent, and to signal one's own availability to others.

Seeing cues as *expressions* produced for others, rather than consumed impressions of others, might be fruitful also for designers, because efforts can be directed to features that support this use. As extreme examples, we saw that the cues were used to reproduce poetry and to play a game ("The ship is loaded with the letter ..."). The expressiveness of the medium itself can be extended by allowing sharing of richer multimedia content or by creating cues that can be appropriated for expression. For example, calendar markings shared as cues can become expressive mediums. If expression is supported, feedback mechanisms must be designed as well. Notification and other turn taking mechanisms can be provided to support the coordination of discussion. We believe that the strategy of providing multiple automatic cues, instead of a few, increases the probability that for a given expression there is cue-based information that can help in sense-making.

6.3 Integration to Media-Sharing

CoMedia utilized two strategies of integrating cues to media-sharing: 1) redundancy and 2) consistency. Redundancy was achieved by providing multiple views of the same cue contents: on the one hand, cues were integrated with media content and, on the other hand, cues were available separately in the Member list view. Consistency was pursued by using the same status icons throughout all views:

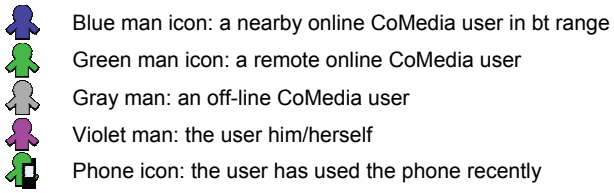

We arranged two field trials of CoMedia, one in Finland during a world championship rally event, and the other at an electronic music festival in Cologne [9]. Logs, media content, video-taped observations, and interviews were analyzed to understand the role of awareness cues. The following roles were found:

- "On-site reporting," reporting activities and events to a remote friend. Here, location and online status cues were important; they enable the remote spectator to engage in interaction and the reporters to know who is online and follows their reports.
- "Keeping up to date with others," following what others are doing at the moment. Here, the location cue was of importance, especially when making sense of media sent by others.

- "Remote spectating," following events, almost as they appear, through video sent by on-site friends. Here, online status cues were important when establishing a connection, but after that they were not as relevant as media itself.
- "Coordinating," planning and monitoring the mobility of other group members. Here, location and proximity cues were important.

To sum up, four uses of media appeared that may not have emerged without awareness cues.

7. DISCUSSION

Attempts at finding a conceptual basis for awareness systems have been disappointing. According to Schmidt [22] and Gross et al. [5], the term awareness has not been used consistently; its very definition is problematic. Multiple flavors of awareness, such as general awareness, collaboration awareness, peripheral awareness, background awareness, passive awareness, reciprocal awareness, mutual awareness, and workspace awareness, have been presented in the literature. From the social cognitive perspective this dispersion has a probable explanation: attributes of action—as in collaboration awareness—have been confounded with attributes of cognitive processing—as in peripheral awareness.

We have started developing a social cognitive approach to mobile awareness first in an empirical paper [15] and then in a more theoretical paper [18]. The present paper has taken this approach one step further and looked at implications for design. Table 2 summarizes the design considerations discussed in the paper. The main contribution of the proposed theoretical approach lies in its ability to expose ways in which interface design is linked to users' belief states (awareness) and actions. We believe that an applied theory of social inference is needed to demystify awareness.

In his review of awareness research, Schmidt [22] concludes: "From a cognitivist point of view, the very notion that an actor is able to pick up and relate to occurrences beyond the scope of his or her line of action and without interrupting that line of action, is difficult if not outright impossible to fathom." In light of the present approach, skilled interpretation of awareness cues becomes understandable; it can be explained by interpretational skills that develop with practice. Such skills allow the user to anticipate, select, filter, elaborate, and generally go beyond the literal meaning of the awareness cues. We have discussed how various design choices promote these skills. Well-designed cues can support the use of pre-knowledge and efficient heuristics, they can promote the formation of habits; well-designed cues are usable in mobile situations where attentional resources are scarce.

These goals may seem almost antithetical to concerns raised in privacy research. If one follows the directions given here to enhance the breadth and accuracy of inferences, privacy issues seem unavoidable. Nonetheless, we believe that to certain extent the positive (inferences) and negative (privacy) sides of awareness can be researched independently—just like an engine manufacturer can try to enhance the acceleration of a car independent of what is known about friction and tires. An intersection exists where both aspects are relevant, however. As soon as users start using awareness cues expressively, the system needs to offer more sophisticated means to control self-disclosure. People need means to manage impressions and influence the inferences others make of them. Analyzing ContextContacts as the application, we have started to explore the space of relevant solutions [21].

The rapid development of the mobile phone, particularly its sensors and networking capabilities, represents perhaps the most promising opportunity for mobile awareness. The phone will be more knowledgeable of its environment; it will be connected to a wider pool of data and computational resources; and it will host more personal content. By bringing awareness cues closer to media and content, a designer can provide resources for users to make sense of that content. By interweaving cues with communication functionality, a designer ensures that whatever is noticed in the cues and content, can be acted on. This paper has discussed mechanisms through which integration affects social inference. Although awareness cues will most likely remain in a secondary role in such applications, they will enable appropriations that would not be possible without them.

Table 2. Summary of the design considerations analyzed in Sections 4 and 5.

Topic	Design choices	Trade-offs
Modality	Textual, iconic	Screen space, expressiveness, learnability, visual search (pop-out)
Level of description	Low, basic, high	Redundancy, informativeness
Realism	Omitting outliers, discretization, highlighting outliers	Screen space, information overload, predictability
Actionability	Neutral description, prompt to act in a certain way	Interpretive flexibility, guidance, privacy, plausible deniability
Temporal reference	Past, present, future; update interval	Battery consumption, timeliness of cues, support for turntaking activities
Number of cues	One, many	Screen space, interpretive flexibility, privacy, disinterest
Order	Making some cues "anchors" by means of prominent positioning, coloring, or shape	Range of interpretations, inadvertent neglect of important information
Division across screens	Making all visible at once, pagination	Interpretive flexibility, "cognitive distance" between cues, visual search
Inspectability of complex cues	Hiding machine logic, exposing it	Deepening interpretations, learning about functioning of cues
Design for mobile use	Notification mechanism; placement of cues in the application space of the phone; rapidity of access	Interruptiveness, suitability for mobile use

ACKNOWLEDGEMENTS

This work was jointly supported by the Academy of Finland project ContextCues and the EU FP-6 project PASION. I express gratitude to several colleagues who have influenced my thinking; particularly Giulio Jacucci, Antti Lindqvist, John Evans, Sakari Tamminen, and Mika Raento.

REFERENCES

[1] Bardram, J., and Hansen, T. The AWARE architecture: Supporting context-mediated social awareness in mobile cooperation. In *Proc. CSCW'04*, ACM Press (2004), 192-201.

[2] Cheverst, K., Byun, H., Fitton, D., Sas, C., Kray, C., and Villar, N. Exploring issues of user model transparency and proactive behaviour in an office environment control system. *User Modeling and User-Adapted Interaction 15*, 3 (2005), 235-273.

[3] Dourish, P., and Bellotti, V. Awareness and coordination in shared workspaces. In *Proc. CSCW'92*, ACM Press (1992), pp. 107-114.

[4] Fitzpatrick, G. *The Locales Framework: Understanding and Designing for Wicked Problems*, Kluwer Academic Publishers (2003).

[5] Gross, T., Stary, C., and Totter, A. User-centered awareness in computer-supported cooperative work systems: Structured embedding of findings from social sciences. *International Journal of Human-Computer Interaction 18* (2005), 323-360.

[6] Grudin, J. Groupware and social dynamics: Eight challenges for developers. *Communications of the ACM 37*, 1 (1994), 92-105.

[7] Gutwin, C., and Greenberg, S. A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work (CSCW) 11*, 3 (2002), 411-446.

[8] Jacucci, G., Oulasvirta, A., and Salovaara, A. Active construction of experience through multimedia: A field study with implications for recording and sharing. *Personal and Ubiquitous Computing* (2007).

[9] Jacucci, G., Oulasvirta, A., Ilmonen, T., Evans, J., and Salovaara, A. CoMedia: Mobile group media for active spectatorship. In *Proc. CHI'07*, ACM Press (2007), pp. 1273-1282.

[10] Laasonen, K., Raento, M., and Toivonen, H. Adaptive on-device location recognition. In *Proc. Pervasive'04*, Springer (2004).

[11] Miettinen, M., and Oulasvirta, A. Predicting time-sharing in mobile interaction. *User Modeling and User-Adapted Interaction (UMUAI) 17*, 5 (2007), 475-510.

[12] Nakanishi, Y., Kumazawa, S., Tsuji, T., and Hakozaki, K. ICAMS2: Developing a mobile communication tool using location information and schedule information with J2ME. In *Proc. MobileHCI'03*, ACM Press (2003).

[13] Nardi, B., Whittaker, S., and Bradner, E. Interaction and outeraction: instant messaging in action. In *Proc. CSCW'00*, ACM Press (2000), pp. 79-88.

[14] Orlikowski, W. J. The duality of technology: rethinking the concept of technology in organizations. *Organization Science 3,* 3 (1992), 398-427.

[15] Oulasvirta, A., Petit, R., Raento, M., and Tiitta, S. Interpreting and acting on mobile awareness cues. *Human-Computer Interaction 22*, 1&2 (2007), 97-135.

[16] Oulasvirta, A., Raento, M., and Tiitta, S. ContextContacts: Re-designing smartphone's contact book to support mobile awareness and collaboration. In *Proc. MobileHCI'05*, ACM Press (2005), pp. 167-174.

[17] Oulasvirta, A., Tamminen, S., Roto, V., and Kuorelahti, J. Interaction in 4-second bursts: The fragmented nature of attentional resources in mobile HCI. In *Proc. CHI'05*, ACM Press (2005), pp. 919-928.

[18] Oulasvirta, A. Social inference through technology. To appear In P. Markopoulos, W. MacKay, & B. de Ruyter (Eds.), *Awareness Systems: Advances in Theory, Methodology and Design*, Springer, Series on HCI.

[19] Patterson, D.J., Ding, X., and Noack, N. Nomatic: Location By, For, and Of Crowds. In. *Proc. LoCA 2006* (2006), pp. 186-203.

[20] Paulos, E., and Goodman, E. The familiar stranger: anxiety, comfort, and play in public places. In *Proc. CHI'04*, ACM Press (2004), pp. 223-230.

[21] Raento, M., and Oulasvirta, A. Designing for privacy and self-presentation in social awareness. *Personal and Ubiquitous Computing*, to appear.

[22] Schmidt, K. The problem with awareness: Introductory remarks on awareness in CSCW. *Computer Supported Cooperative Work (CSCW) 11*, 3 (2002), 285-298.

[23] Sproull, L., and Kiesler, S. Reducing social context cues: Electronic mail in organizational communications. *Management Science 32*, 11 (1986), 1492-1512.

[24] Tang, J., Yankelovich, N., Begole, J., Li, F., and Bhalodia, J. Connexus to awarenex: Extending awareness to mobile users. In *Proc. CHI'01*, ACM Press (2001), pp. 221-228.

[25] Tullio, J., Dey, A., Chalecki, J., and Fogarty, J. How it works: a field study of non-technical users interacting with an intelligent system. In *Proc. CHI'07*, ACM Press (2007), pp. 31-40.

[26] Tversky, A., and Kahneman, D. Judgment under uncertainty: Heuristics and biases. *Preference, Belief, and Similarity: Selected writings.* Bradford Book (2004).