

Intelligence in Context

Alex S. Taylor
Microsoft Research
7 J J Thomson Avenue
Cambridge
CB3 0FB, UK
ast@microsoft.com

ABSTRACT

In this paper, I draw on some previously published ideas in HCI and ubicomp to question a number of assumptions made around computational intelligence. I pay specific attention to a central concern in the ubicomp and intelligent environments research, context-awareness, and argue that a particular idea of intelligence is implied in designing environments to be context-aware. Using three design examples taken from the ubicomp literature, I make a case against designing for this idea of intelligence, suggesting that it misjudges the capacity computers have to be intelligent. As an alternative, I propose a different perspective on intelligence and how it might be designed into our everyday environments. From this perspective, the intelligence of a technology or environment is based on how it enables people to (inter)act as they see fit—how it enables ordinary human endeavour. I develop an early basis for this position using the principle of embodiment, and attempt to orient it towards design using five concepts: augmented-awareness, off-loading, intelligibility, visibility and affordances.

Keywords

Intelligence, context-awareness, augmented-awareness, off-loading, intelligibility, visibility, affordance, embodiment.

1. INTRODUCTION

In this paper, I want to examine the idea of intelligence and what it might mean to have something like it incorporated into our everyday environments. There are, of course, a whole host of ways in which intelligence might be thought of from this perspective. We might think of intelligence as it has generally been considered in AI, where the objective has been to replicate the mind, its processes and, ultimately, aspects of human behaviour. Alternatively, we could contemplate a more partial view of intelligence in computation, one in which particular features of what is imagined to be intelligence are built into our environments. Here we have smart-cards, smart-cars, smart-homes, and so on, all with some capacity for interacting ‘intelligently’ in the world. There are also the various possibilities of form and embodiment. On the one hand intelligence could be viewed as something beyond form, something to be uncoupled from the physical constraints of the world. On the other, intelligence might be seen as inextricably tied to the material

relations between things and people, tightly interwoven into our lived experiences.

To take some initial, tentative steps into this area, I’ve chosen to centre on a particular concern that intersects with several of these positions. Specifically, I’ve chosen to focus on something I see to be implicated in the prevailing visions of intelligent environments, namely *context-awareness*. In choosing this focus, my intention here is to avoid tackling, in their entirety, the many and varied positions on intelligence. Instead, what I want to do is approach the issue from a pragmatic position, questioning some of the ideas assumed or implied in the very business of making environments intelligent, or ‘smart’. As I hope to demonstrate, this perspective has its merits in opening up the possibilities for designing intelligence into our environments

An underlying premise to my thoughts in this paper is that context-awareness is one of numerous logical extensions to computational intelligence and especially to the notion of embedding intelligence into our environments. The term implies that an environment, whether it be a home, office, shop, etc. can be designed to be aware of human behaviour and respond to it in an appropriate fashion. Thus, by building an architecture to support context-awareness, it is assumed that our environments might know when we want to eat, what music we like, whether we want to be interrupted, and so on. The ‘intelligence’ here is understood to be something associated with interpreting human behaviour, reasoning about it and acting upon it.

As it happens, this mode of thinking is precisely what dominates much of the research surrounding intelligent environments. As a case in point, the research behind smart homes appears overwhelmingly concerned with the quantity of information that can be sensed and how that information is interpreted so that embedded machines might anticipate and respond to what household members want. The corollary to this is that a home’s smartness might be improved through the use of more sensors, sensing a greater number of things and applying better methods of interpretation.

In what follows, then, I seek to unpack some of the ideas around context-awareness and reflect on some of the assumptions that appear to underlie the design of intelligent environments. This exercise begins with some background to the interpretation of context and the application of con-

text-awareness in system design. In drawing on this past material, I want to make clear that my position here is not entirely new or unique, but in one way or another brings together a set of arguments that I see as relevant to the problem of designing for intelligence. This paper should thus be seen as a vehicle to rearticulate some well-established arguments with the objective of sensitising the intelligent environments community to an alternative understanding of the intelligence-computation relationship.

To develop my thoughts and integrate the various threads from previous research, I'll use three examples taken from the ubicomp literature. Although ubicomp and intelligent environments research are not direct analogues, I see sufficient overlap, particularly around the ideas of context-awareness, to warrant their use as illustrative. Alongside my comments on context-awareness, I'll present how I think a different application of computation might be used to augment our everyday environments and how, in turn, this infers a different and possibly more manageable basis from which to think of intelligence.

1.1 The Problem of Context

Casting an inquisitive eye towards the idea of intelligent environments, the question that comes to mind is whether there is anything curious or peculiar about the problem space or, to put it more colloquially, *what's the rub?* Those of us from the socio-technical research community, in particular, would ask ourselves, above and beyond the technical challenges, what, if anything, makes the building of intelligent environments an interesting and challenging problem. The rub, as I see it, has all to do with what has emerged as a central concern for the intelligent environments 'project', that of designing technology to be sensitive to its surroundings—what has come to be referred to as *context-awareness*. It's my opinion that the issues surrounding context and the difficulties involved in building context-aware applications present a problem that will be the make or break for the project as a whole.

Arguing in a similar vein, Paul Dourish (2004) has developed what I find to be a well thought out position on context and how it has been conceived of in ubiquitous computing research. He suggests that those who have come to ubicomp from a technical perspective have, by and large, tackled context as a *representational problem* in which the significant challenge has been how to design technologies to encode and represent specific situational dimensions, e.g., location, time, co-presence, etc.

Dourish claims that this view of context neglects the more fundamental problem of *what* context is. Drawing on a philosophical tradition rooted in phenomenology, he argues that context cannot be easily reduced to predefined dimensions, but rather is something that comes about, or is *produced*, through the interactions between people and the material world in which they inhabit. Put simply, this line of argument views context as something that is continuously being made and dependent to a large degree on the

ever-changing relations between people and the resources they bring to bear in everyday settings. From this position, context ought not to be articulated in representational terms, but, instead, seen as an *interactional problem*.

These seemingly divided positions on context place much of the work around intelligent environments in an awkward position. The prevailing idea of intelligence is shaped by this notion of context as a representational problem. As I suggested earlier, the smartness or intelligence of a machine or even environment is dictated by how it is enabled to sense, encode and interpret the world and our interactions. Offering a way forward from this apparent impasse, what I'd like to consider is what it might mean for the intelligent environments project if we stopped thinking about *sensing context* and thought more about the role of computing in *context production*. Moreover, I want to consider what this would mean for the very idea of intelligence.

The way I want to develop this is to draw on a perspective of embodiment, a perspective that has multiple threads of inquiry in the social sciences and humanities. As I see it, the concept of embodiment provides a basis from which to better capture the ongoing production of context and to highlight the importance embodied interactions have on the ways we orient our behaviours in our environments (also see Clark, 2003). As a general concept, embodiment foregrounds the fundamental role material artefacts play in our everyday interactions. At one level, material artefacts and the properties they have are seen to relate to the ways we think about and find meaning in the world, and thus how we interact with it (Hutchins, 1995; Norman, 1988). The material properties of paper, for example, shape the way we read and thus make sense of printed information (O'Hara & Sellen, 1997; O'Hara, Taylor, Newman, & Sellen, 2002).

This sense of embodiment can be broadened to incorporate real-world practices, practices that are understood to evolve through people's ongoing interactions and negotiations with each other and the objects they use in the world. This position allows us to envisage computing systems embedded in our environments not as solutions for sensing and encoding context, but instead as things that can be enlisted to get on with the practical routines in everyday life.

2. CASE STUDIES

To ground some of the ideas above, in the remains of this paper I will present what I see as three reasonably representative examples of technically oriented contributions to the ubicomp project. In presenting these examples as case studies, my intention here is to offer a practical way forward. Each example is also used to discuss some concrete terms in order to do so. In the first example, describing the MediaCup, I introduce the term *augmented-awareness* in an attempt to shift the emphasis away from systems built to sense and encode context and towards solutions that support its ongoing production. Augmented-awareness will thus be presented as a basis from which to develop systems that aim to take into account how people already make sense of and

act in the settings they are in and that augment the ways in which information is made available to do so. Second, in my critique of the Nomadic Radio, I raise the concepts of *off-loading* and *intelligibility*, both of which aim to address what exactly it is a ubicomp solution should support. In the final example, the design principles of *visibility* and *affordances* are considered in response to a conceptual framework for designing context-aware applications. These last two principles are intended to offer positions from which to design solutions that are intelligible to users.

2.1 MediaCup

To begin, then, I turn to a research group that has centred on building ‘ubiquitous computing environments’ and look, specifically, at one of their projects, the [MediaCup](#) (Fig 1: Beigl, Gellersen, & Schmidt, 2001). Research into the MediaCup began in 1999 by members of the Telecooperation Office ([TecO](#)), a group with a strong emphasis in computer science located in the Telematics Department at the University of Karlsruhe, Germany.




Figure 1. MediaCup



Figure 2. With HotClock

The MediaCup has been presented as a demonstrator for exploring the augmentation of everyday objects. It was originally envisaged as a way to explore how specific features of ubiquitous computing might be introduced through existing everyday objects. Specifically, location- and context-awareness were felt to be features that could be embedded, non-intrusively, into ordinary objects—avoiding the introduction of new devices.

Able to sense temperature and acceleration and to communicate to a networked infrastructure via infrared (for technical details, see Gellersen, Beigl, & Krull, 1999), the MediaCup has been trialled at the Telecooperation Office and introduced through a number of application scenarios. As part of the ubiquitous computing environment implemented in the Telecooperation offices, the cup has been designed to communicate with objects such as the HotClock to display the temperature of the liquid it contains (Fig. 2). Using imbedded accelerometers and an ID tag system alongside infrared, ‘contextual information’ can also be detected such as whether someone is drinking from a cup or the cup’s physical location in the office. These features have been

combined with, for example, the group’s *SmartDoorplates*—LCD panels attached to the outside of doors—to display whether the cups associated with individuals are present within a room and whether the cups are in use.

The research behind the MediaCup has been driven by an overall concept of ubiquitous computing that envisages technologies pervading our lives, embedded into everyday objects, rendering computational operations ‘invisible’. The underlying premise is that these embedded technologies might work together seamlessly to augment the world in which we operate (see Beigl, 1999). The purpose of objects such as the MediaCup, then, is not merely to augment the cup, but also to contribute to the ubiquitous computing environment’s infrastructure itself and the combined features of augmented objects.

One reason for envisaging this ‘ecology’ of devices is to enable different objects to contribute to an overall representation of context. Instead of demanding each and every device to ‘sense’ context in its entirety, objects are designed to detect the different contextual features of a setting that are relevant to their use. What we see, then, is that the research behind the MediaCup and its envisaged place in a ubicomp environment presents context as something to be reduced into manageable chunks—detected, encoded and then pieced back together to augment real-world experiences. To use Dourish’s (2004) term, it’s evident that the detection of context has been conceived of as a *representational problem*. This position is made all the more clear in a paper by Beigl, Gellersen and Schmidt:

Artefact context is the information a digital artefact maintains about itself and its immediate physical environment. Digital artefacts can embody sensors and perceptual computing to establish their context. (Beigl et al., 2001, p. 403-404)

And later in the paper:

Digital artefacts communicate their context into a geometrically defined congruent communication space. Digital artefacts and/or computing devices can retrieve artefact context from any communication space they physically enter without further knowledge of the artefacts from which the contexts originate. (ibid., p. 404)

From these short quotes, we find that given sufficient computational resources, context, for Beigl et al., really is something that can be determined. Artefacts are considered to be a source of contextual information and, indeed, have their own specific ‘artefact contexts’. Thus, the movement of a cup in a room and the people assumed to be present are seen as two qualities of a setting that might add up to capturing the overall context of an activity.

What I don’t wish to do here is simply provide yet a further example of how this conception of context as a representational problem can be set against an interactional paradigm. Rather, I want to discuss a particularly interesting tension that arises when we consider how it is the MediaCup can be understood, at one and the same time, as a means of detecting context and also as an object with qualities that are

integral to the unfolding of social interaction. My reason for this examination is to illustrate, through a practical example, why it is context-awareness is a fundamental problem for ubicomp and to introduce a possible shift in direction away from building context-aware applications to considering *augmented-awareness*.

What the designers of the MediaCup have done well is to recognise the importance of making things visible or intelligible through embodied interaction. At a simple level, we see that when a watch is brought into proximity to a coffee cup, the watch displays the temperature of the coffee. Putting to one side the usefulness of this feature, this is something someone could easily figure out because a quite 'natural' embodied interaction—bringing the hand to the cup—makes it clear that the spatial relationship between the objects triggers the display of temperature. Similarly, the change of a SmartDoorplate's display when a cup is brought into a room reveals a material relationship between the cup and the room. In this latter example, we also see that the cup's association with the room can have an impact on how that room is understood to be occupied and who might or might not enter it—if the room is seen to be occupied by people having a meeting then that necessarily includes or excludes others. The cup not only interacts with objects and people in the world, it also interacts with the social status of the room. We might say that the cup contributes to the production of a social context because it transforms the relations between people.

It's this final point that really gets to the crux of my problem with an artefact, such as a coffee cup, both sensing *and* producing context. As I see it, the trouble is that the idea of a cup's location being a representation of context, on the one hand, and it being part of the detail that contributes to the ongoing production of context, on the other, do not sit easily with one another. As we have seen, the premise behind the MediaCup is that a cup's location and whether it is being handled are adequate descriptions to determine something of the context of a situation. If someone's cup is located in their office and in use, we might hazard a guess that they are at work. Moreover, context can be further deciphered when this information is combined with additional detail, such as the relative location of other mugs, their combined use/handling, the amount of liquid they contain, etc. So a group of cups together, all containing hot liquid and all in use might indicate a meeting of some sort. The corollary to this is that the greater the computational force thrown into sensing the world, the better the prediction of a setting's context. Adding more sensing technologies to an office setting, to operate in combination with the MediaCup, for example, would increase the capacity to accurately represent context.

From an interactional standpoint, I'm not at all convinced that any amount of sensing can determine a situation's context. What a sensor does do is present information that may or may not be used to direct or shape the ongoing production of context. The location of a coffee cup, then, is

one of many things that someone might become aware of to make sense of the world and piece together a meaningful context in which to act; it might, for instance, help us to get an idea of whether a colleague is in her office and, along with other information, whether she can be interrupted with the question we've been meaning to ask her. However, at no point is it, in and of itself, a representation of a situation's context; the cup's location, or any other combination of sensed information, cannot stand in for how we have interpreted our colleague's availability and what we want to talk to her about.

What I'm suggesting here is that sensed information cannot be seen as a representation of context because it is part and parcel of that information we might choose to assemble and interpret in shaping how a situation's context unfolds. Thus, the technologies and applications designed under the rubric of ubiquitous computing cannot and probably never will be able to successfully interpret a setting's context. Put simply, this is because they are inextricably a part of producing the very context that they are meant to be removed from to accurately encode and represent. The information they sense and encode merely makes up the repertoire of resources that we might be made aware of and that can be brought to bear in the production of context.

Somewhat predictably, what this tension confirms is that two very different conceptions of context are at work. Perhaps more usefully, what it also hints at, however, is that they are not mutually exclusive. If we see that the information being sensed and encoded by devices such as the MediaCup as possible resources in the ongoing production of context rather than accurate representations of it, we begin to catch a glimpse of how such information might be made useful. Rather than amass information with the assumption that context can be predicted, we see that as an alternative, information might simply be made available to a situation and those present so that it can act as yet another resource in the ongoing production of context.

The distinction may be subtle here, but it's my contention that the implications for design are significant. This shift in emphasis allows us to make a move away from designing systems that seek to accurately sense and interpret context and, instead, direct our attention to solutions that augment the information that is made available to us in particular settings and how we are able to act upon it. To capture this shift towards increasing people's awareness of information through the (re)design of environments, I have chosen to use the phrase *augmented-awareness*. In presenting the next two examples, I go onto develop an argument around what designing for augmented-awareness might mean and how it might be achieved, practically. I also hope to demonstrate how an environment's 'intelligence' is not necessarily achieved through the collection of more data using myriad sensors, embedded in different objects. Designing intelligence into the world can sometimes mean simply designing environments that behave as expected in ways that are understandable to us.

2.2 Nomadic Radio

The *Nomadic Radio* is a system that emerged from research by the [Speech Interface Group](#) at MIT's MedaLab between 1997-1999 (Sawhney & Schmandt, 1999, 2000). The system is designed to notify users of incoming information, using both an auditory and 'tactile' interface (Fig. 3). Email, voicemail, hourly news broadcasts and personal calendar events are managed by the system and relayed to the user over an auditory channel. Using attributes, such as message type, unread status, priority and time of arrival, the incoming information is dynamically assigned to categories. This categorisation scheme is used in combination with 'notification strategies' to determine when a user should be warned of incoming information, deciding, for example, whether the notification of an incoming email should be made instantly or delayed for later. Auditory information is presented using specifically designed cues to indicate system activity, message notification, confirmations and breakdowns. To navigate and manage their messages, users are able to use either speech or button input.




Figure 3. Nomadic Radio wearable device.

The system's notification strategies are based on the *inferred priority* of a message and *user context*. These are derived from three immediate sources of information: the message's priority, defined by the system's filtering system; how recently and the frequency to which the wearable device has been interacted with; and the level of conversation occurring in the immediate vicinity. A user's past responses to notification can also be taken into account to develop the strategies over time. Thus, a user will be immediately notified of a high priority message if it arrives when he/she has recently been using the device and where no nearby conversation is detected. The original time of incoming information is also represented spatially using stereo speakers (Fig. 4).




Figure 4. Spatial sound system on Nomadic Radio.

User preferences and the notification strategies are used to determine the type of auditory notification. Using a 'scal-

able auditory presentation', seven levels of notification are possible. These range from 'silence' to minimise interruptions, to 'foreground', where a message is read in full and presented spatially, closer to the listener. Worthy of mention is the second level of notification, 'ambient'. At this level, a constant sound of flowing water is used to indicate the system is operating. The pitch of the flow is increased during file downloads, whilst other water sounds, such as splashes and faster flowing water, are used to indicate different forms of incoming information. This spatial system is also used to simultaneously play audio streams in the foreground or background of the user's 'audio scene'. The real-time scalable feedback is dictated by the notification strategies, so that the notifications for less urgent messages or news, for example, can be played as ambient, background sounds rather than announced in the foreground.

The motivations for designing Nomadic Radio were to provide a unobtrusive system for notifying users of incoming information, and to allow them to listen to and navigate this information when convenient. As a system, the Nomadic Radio presents an interesting concept because it goes some way towards providing the vision of augmented-awareness that I raised in the discussion of the MediaCup. That is, it senses incoming information in order to augment what's immediately available in a setting; it augments a user's awareness by providing him or her with more information than he or she would ordinarily have.

My problem with the Nomadic Radio, however, is that it continues along the same lines as the MediaCup insofar as it aims to predict the significance or meaningfulness of information to a user by attempting to interpret context. Again, some notion of intelligence is assumed on the part of Nomadic radio because it is able to interpret the world, in the absence of human intervention. Specifically, as well as sensing information in the world, it seeks to aggregate and then act on it by predicting how people might respond to incoming news, messages and calendar events. Not surprisingly, it's all too easy to see where the system would fail because of its insensitivity to the nuance of social interaction. How well, for example, would the full foreground notifications be received in non-verbal interactions, perhaps in the case of an intimate embrace? As we have seen before, the proposed solution to such contingencies is to throw more weight into the sensing and the so-called intelligence of working out context. In a paper published in the *ACM Transactions on Computer-Human Interaction*, the system's architects explain:

In Nomadic Radio, context dynamically scales the notifications for incoming messages. The primary contextual cues used include message priority from email filtering, usage level based on time since last user action, and the likelihood of conversation estimated from real-time analysis of the auditory scene. In our experience these parameters provide sufficient context to scale notifications; however, data from motion or location sensors can also be integrated in such a model (Sawhney & Schmandt, 2000).

In the same paper, Sawhney and Schmandt refer to a study by O’Conaill and Frohlich (1995) to explain the reasoning behind Nomadic Radio:

The study suggests development of improved filtering techniques that are especially *lightweight*, i.e., do not require more attention from the user and are less disruptive than the interruption itself. By moving interruptions to asynchronous media, messages can be stored for retrieval and delivery at more appropriate times. Personal messaging and communication, demonstrated in Nomadic Radio, provides a simple and constrained problem domain in which to develop and evaluate a contextual notification model (emphasis added, Sawhney & Schmandt, 2000, pp. 368-369).

What is assumed by the authors is that by alleviating the user of having to make judgements about incoming messages they reduce the overall demands on attention, thus making the handling of incoming information less disruptive. Moreover, they assume that integrating motion and location sensing equipment into their architecture will improve the system’s judgements and further reduce the demands on the user.

Commenting on context-aware applications, Jonathan Grudin (2001) raises a problem with these somewhat crude assumptions. He reveals that matters can be made more complex by having systems interpret context by pointing out that contextual information’s very nature is altered when aggregated and interpreted in software. He argues that such profound changes can cause disruptions to well-understood and established practices, and the more context work that is done the more likely and fundamental these disruptions are likely to be. Under these terms, the system described by Sawhney and Schmandt would appear to be far from lightweight because it relies so heavily on forcing the user to reinterpret and make new (or at least different) judgements about the information they receive. A good deal of work is being done on the incoming information by the system that is somehow opaque to users, requiring them to drastically re-learn the way they interact with emails, voicemails, etc. The proposal to integrate additional sensors threatens to further complicate the system and make its workings more obscure. The end result of Nomadic Radio’s so-called context-aware features is to significantly alter routine practice and hence the production of context. It is not, as it is assumed, to sense context in a detached fashion.

To my mind, the architects of Nomadic Radio have misjudged what demands to place on their system. Not only have they fallen foul of the same problem of trying to interpret context, but they’ve also attempted to replicate and in some cases replace the human’s capacity to make judgements about information. In essence, they’ve attempted to imbue the device with some idea of intelligence. By making it ‘smart’ in the way they have, their argument is that the Nomadic Radio can reduce the burden, or workload, of continually checking for messages and judging whether particular messages warrant an interruption.

The trouble is the Nomadic Radio seems to be pretty unintelligent because it’s doing something it’s likely to get wrong and that we’re reasonably good at as humans anyway. Indeed, we only need to look around us to discover that people are very resourceful at managing interruptions without the need for context-sensitive systems. The use of the mobile phone stands as a case in point. Our methods of answering the mobile phone (Laurier, 2001; Weilenmann, 2003) or simply switching the phone off stand as lightweight solutions to handling incoming calls. The provision of caller ID offers information that helps us manage this, but importantly, in sensing the caller, the phone in no way tries to interpret context and make decisions on our behalf. What’s more, the devil is in the detail, so to speak; these seemingly trivial snippets of information are made into so much more through the ways they are artfully enlisted. A caller’s name displayed on a phone’s display, for instance, comes stand in as a summons, not merely because of who is calling but because there is an implicative moral duty to answer some calls rather than others—e.g., your boyfriend’s over and above one from the so-called friend who didn’t return your call yesterday. The importance of a call and whether to answer it is worked out through the intermingling of the phone, its features, and its ongoing use. The likelihood of any intelligent algorithm(s) supplanting this interpretive work and making the correct judgements is small, to say the least.

It would seem that the designers of the Nomadic Radio have offloaded the wrong set of problems onto the system because they have assumed that context can be intelligently interpreted through computation. As a consequence, their heavyweight solution not only demands a lot from the system *vis-à-vis* its capacity to interpret context, it also demands a good deal from users in so far as they are expected to understand its complex workings, and subsequently relearn how to manage and interpret incoming information. Indeed, this is hinted at in their own early evaluation of the system that revealed users’ problems with understanding where they were when navigating through their messages. Users also found the auditory scheme complex and wished for either a simpler classification of notifications or more time to gain familiarity with it.

I’d like to suggest that my comments on the use of Nomadic Radio bring to the fore two useful resources with which to consider designing for augmented-awareness. For one, we’re reminded of the phrase *off-loading* that Buxton (1995) coined in discussing the problems of managing complexity when new functionality is made available in order to support context sensitive interaction. Working on the [Ontario Telepresence Project](#), Buxton and his colleagues demonstrated that through a user-centred approach and, specifically, by considering the primary decisions and actions that are made in using a technology, it is possible to determine what’s important for the user to know and control and what can be handled by computation. A central tenant was that the hard computational work should be

offloaded onto technology, but that users should remain in direct control of their actions. The project was specific about context-sensitive technologies not intruding on users and their actions and activities, but rather augmenting them (Riesenbach, et al., 1995, pp. 14-15). Intrusions were invariably seen to lead to changing a setting's demands and dynamics (and potentially the context) and thus complicating rather than easing the complexity of interaction.

The second resource to be brought to the fore is *intelligibility*. Introduced by Brown and Newman (1985) and later directed specifically at context-aware applications by Bellotti and Edwards (2001), intelligibility is closely tied to the notion of off-loading. The argument is that the work that is off-loaded onto the system, that is directly related to the users interaction with information, should be easily understood by users. As Bellotti and Edwards contend, systems that make inferences “must be able to represent to their users what they know, how they know it, and what they’re going to do about it” (ibid., p. 201). Importantly, Brown and Newman sought to raise the awareness of mutual intelligibility, whereby a system’s operations are jointly understood through users’ interactions with one another and not only with the system itself.

For augmented-awareness, the concept of off-loading gives us a starting point from which to consider what sorts of human awareness a system could attempt to augment without adding to the complexity of dealing with information. The concept highlights that augmented-awareness is something to be considered in light of the decision-making and actions humans find relatively easy to manage or want to be in control of. The work left to or off-loaded onto computation should be that which lessens the complexity of information that such decisions and actions need to be based on. Stated in these broad terms, these points raise questions about the sorts of information that should be sensed and captured by technology and what and how might it be best conveyed and presented to users? The intelligence of a system or environment, from this perspective, might be better thought of as how it allows for this. In the following, I will introduce two final design principles that aim to go some way towards addressing these questions.

2.3 A Conceptual Framework

This last example stands as a point of departure from the work I’ve presented in the preceding sections because it tackles the broader subject of a conceptual framework for designing context-aware, ubicomp applications. Even at this more abstract level, however, I hope to demonstrate that the concerns surrounding how context is conceived and the matters of off-loading and intelligibility remain central. I also aim to develop my argument around augmented-awareness by suggesting that one way in which it might be addressed is through the consideration of two design principles based on embodied practice: visibility and affordances.

This example, then, centres on a conceptual framework and its instantiation in the form of [context toolkits](#) designed to

facilitate the development and deployment of context-aware applications. Originally conceived of at the [Future Computing Environments](#) group at Georgia Tech’s College of Computing, the framework has since been formalised and applied in several application domains in order to demonstrate how ubicomp systems can be designed to properly handle the complexities of sensing, representing and reacting to context (see Dey, 2001; Dey, Abowd, & Salber, 2001; Dey, Mankoff, Abowd, & Carter, 2002).

Building on the established reputation of the Future Computing Environments group, Dey and his colleagues have set out six ‘requirements’ that should be met by a conceptual framework—requirements that they contend must be addressed for designers to more easily deal with context. Broadly, the six requirements presented build on the principles of abstraction, reuse and component-based architectures to ease system development and enable the building of lightweight applications. Drawing on these principles, the requirements have been translated into a conceptual framework with five categories of components that, in turn, allow for instantiations of the framework, or Context Toolkits, to be configured. Figure 5 presents an example of a toolkit with the aggregator, discoverer, interpreter, service and context widget components. For the specifics behind these categories and the framework, I refer the reader to Dey et al.’s anchor article in the special issue of Human-Computer Interaction on context-awareness (Dey et al., 2001).

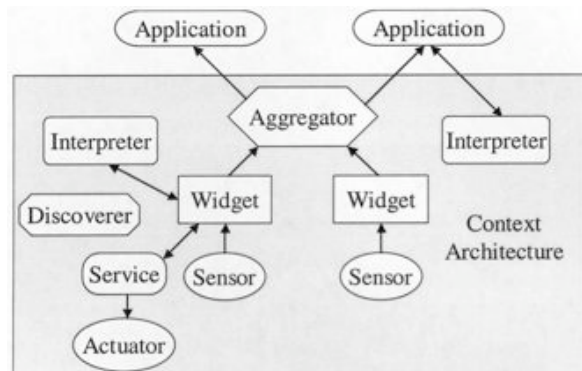


Figure 5. Example configuration of Context Toolkit.

Leaving the detail to one side, first, let us see, yet again, how context has been configured as a representational problem through this example. Above we’ve seen how Dourish characterised the representational conceptualisation of context as something that equates to information. From their definition of context, it is clear that Dey et al. view this as the case:

Context: any information that can be used to characterize the situation of entities (i.e., whether a person, place, or object) that are considered relevant to the interaction between a user and an application, including the user and application themselves. Context is typically the location, identity, and state of people, groups, and computational and physical objects. (Dey et al., 2001, p. 106)

Through their framework, Dey et al. do a good deal of work to separate out the forms of context that technology is suited to sense and encode, and, as they express it, aggregate and interpret. This again raises the same old division between the representational and interactional problems of conceptualising context. It presupposes that the problem with context for technology has not been with what it is, but rather how to capture and encode it.

The omission that is most striking in the framework Dey and his colleagues have developed is any representation of human involvement in sensing, interpreting and acting on contextual information (and producing context). Because they have conceived of context as a representational problem, they have worked from the basis that contextual information can be extracted from a setting and fed into an architecture of interacting computational components. As a consequence, the framework offers no resources to contemplate how the aggregation and interpretation of information from a setting might best be done in relation to the users of context-aware applications. Indeed, the separation of applications from the sensing, inferencing and service components in the framework threatens to further remove users from a sufficient understanding the workings of a system, making its use potentially more complex. Commenting on the framework's component-based architecture, Bellotti and Edwards (2001, p. 198) have remarked that designers who apply the framework run the risk of "building a context-aware application that is ignorant of and thus hides the nature of its own sensing and inferencing..." leading to an absence of the sensed and inferred information that a user might need to operate a system.

We're reminded here of the Nomadic Radio where it's not clear what benefits would be achieved by off-loading certain sorts of work onto technology and how a system's actions should be made available and intelligible to people who come into contact with it. I suggested that off-loading the interpretive work onto the system led to users' confusion over notifications because the filtering was hidden and thus unintelligible. I also suggested that this interpretive work of judging availability and the importance of incoming information is something people routinely do and do so in ways that fit the dynamic ongoing production of context.

Returning to the framework, the question that arises in criticising the glossing over of human involvement, and specifically the issues surrounding off-loading and intelligibility, is what can be done with respect to design to consider the dynamic and interpretive character of human interaction. To address this question, I introduce two further interrelated design principles, *visibility* and *affordances*, that go some way towards tackling the considerations of off-loading and intelligibility.

As I suggested at the beginning of this paper, much of our everyday action is embodied. We rely on a host of properties of material things to make judgements about how we should interact with them and how they work. Although it

arguable failed to account for the dynamic nature of context production, the MediaCup capitalised on this sense of embodied (inter)action by making information available through the physical proximity of objects and spaces. Indeed, as Don Norman (1988) has pointed out, a cup or, more specifically, mug provides an illustration of how the interactional features of a device can be made visible and afford being used; its handle makes a mug's interactional features clearly visible and the handle's placement and shape afford holding and gripping.

Drawing on a paper I've already cited by PARC's Victoria Bellotti and Kieth Edwards (2001), my argument is that ubicomp technologies have the potential to reduce the burden on users dealing with multifarious forms of information by making it intuitively visible and available in everyday action. From what we've seen, the solution to handling information is not to simply off-load the work of interpreting context onto computation. As Bellotti and Edwards (*ibid.*, p. 199) rightly point out, "the more we try to get systems to act on our behalf, especially in relation to other people, the more we have to watch every move they make." Suggesting a way forward, they conclude that "a set of design principles are needed that will enable users to be able to reasonably and rationally make their own inferences based on the state of the world as perceived by them, and as reported by the system" (*ibid.*, p. 201).

It is my contention that the dual notions of *visibility* and *affordances*—that have been central to other programmes of research in HCI—provide some useful building blocks from which to consider how ubicomp technologies might be made available as resources in the ongoing production of context. What the notions of visibility and affordances allow for is a shift in emphasis away from technologies attempting to separate out and accurately interpret context and towards a 'model' of ubicomp technologies making available the resources with which people can make judgements, themselves, about their actions. In a somewhat unintuitive way, the tenants of invisibility and pervasiveness that the ubicomp project turn on are achieved by making certain aspects of interaction 'seamful' (Weiser, 1994), rather than seamless, and thereby making the most of a device's computational and physical properties (MacColl et al., 2002). That is to say, visibility and affordances offer two design principles from which to consider the features of a solution should be geared towards this 'seamfulness' and the form in which information is made available to make it unobtrusive and yet at-hand.

Visibility relies on a fairly basic premise. That is, it purports that the interactional features of a technology and how they have been designed to operate in real-world practices should be visible to the user. The established design principle of feedback can offer a fairly simple example of good visibility in using technology. For example, to make a call on a standard landline phone we're presented with several forms of feedback that help make the operations *visible* (N.B. I talk here of visibility in the general sense of the

word rather than limited to the visual modality). First, when we pick up a receiver we hear a dial tone that informs us of the readiness of the device. Admittedly, the sound of the dial tone is something that we must learn to understand, but once learnt it offers feedback that is sufficiently simple to comprehend every time we pick up the receiver. We then hear direct auditory feedback in dialling a number and pressing the number keys. If skilled, the caller even has some notion of the relative ordinal value of the numbers pressed based on the tone made by each key press. Finally, the ring tone sounds almost instantaneously once the correct numbers have been entered. The system is sufficiently intelligent to work out a correct number has been dialled and provides immediate feedback in the form of a ring tone.

Affordances are closely tied to this idea of visibility. Emerging from research into visual perception in the mid to late seventies (Gibson, 1979), the term affordances has found favour amongst the design community as a means of expressing the ways in which artefacts afford particular understandings and actions through their physical and interactional properties (e.g., Gaver, 1991; Norman, 1988; Sellen & Harper, 2002). Returning to the example of using a landline phone, a standard phone's handset could itself be said to afford being picked up and grasped because of its visual appearance. Thus, even if we were totally unfamiliar with the device, it wouldn't take long to figure out how to interact with the handset if a phone were to ring. The grip affords being held; upon hearing a disembodied voice coming from the earpiece, the ear and mouthpiece afford being placed to the ear and mouth; and the wire between the handset and phone's base constrain, to some extent, what way up to hold the handset.

The same principles of visibility and affordances can, I argue, be similarly applied when considering augmented-awareness in ubicomp environments. Although the systems might potentially have more flexibility, the importance of visibility and how (inter)actions are afforded remain key. To demonstrate this and tie these principles back to the points I have made throughout this report, in conclusion I want to briefly reconsider the design of the Nomadic Radio and propose a number of possible modifications.

3. CONCLUSIONS

Returning to the central focus of this paper, intelligence, my hope is that the five concepts I've outlined above (i.e., augmented-awareness, off-loading, intelligibility, visibility and affordances) provide a means to reconsider what it is to design intelligence into our environments. What I hope I've demonstrated is that an environment's intelligence and the intelligence of those objects within it might be judged on how they enable *us* to be aware of the world and act in it. The job of computation should not, then, be that of replacing human endeavour, but rather how best to intelligently support it. Design, thus, ought to be oriented around the opening up of possibilities people have to exploit the world in intelligent ways. Relying on the concepts of intelligibility

and affordance, the design problem centres on how to make this exploitation possible in intuitive ways. This, as I've suggested, is where a concern for embodiment comes into play; a concern for embodiment, and the physical properties of things we are familiar with, allows us to build on people's established expectations of how things work.

To wrap up, I want to express these ideas though example by returning to the Nomadic Radio system described earlier. In considering the design of a wearable, audio-based notification system, I'm reminded of how I operate my personal stereo when I'm out and about and find myself talking with passers-by, shop assistants, people calling my mobile phone, etc. To handle these brief encounters, I simply remove one earpiece. As an almost instinctive move this offers a very good solution to handling what I believe at the time will be a short interruption. At a purely attentional level, the removal of the earpiece decreases the apparent volume of the music I'm listening to so that I can attend to the talk at hand. When talking to someone face to face, the solution is also useful at an interactional level. It serves as a demonstrable cue to my interlocutor that I am open to talk. Like other gestures we make, such as moving one's ear towards someone speaking or opening the palm of the hand around the ear, it offers a cue to indicate *I-am-listening*. It also avoids the fuss and fiddling and slightly uncomfortable pause in talk that often ensues when I try to retrieve the player from a pocket or bag and depress the pause button or some other equivalent.

This method can also be easily extended to talk that I might foresee to be slightly longer or more important to me; in such cases the removal of both earpieces decreases the apparent music volume to an inaudible level and presents me as wholly committed to the conversation. Whatever the case, what should be apparent here is that I've adopted my own method for dealing with interruptions that is highly embodied and fits closely to the concepts I've discussed above. With respect to off-loading, the personal stereo is left to handle what it is good at—playing music—whereas I take control of setting the apparent volume in response to the contexts I find myself in. This is an intelligible method to me because it's clear that the volume of music is associated with the proximity of the earpiece to my ear. The relationship between earpiece and volume is clear or visible, and it's this very visibility that affords the removal of the earpiece(s) rather than any interaction with the device tucked away in a pocket or bag.

My point in describing this method is to demonstrate that the features of Nomadic Radio could be improved if its designers were to take into account how users currently deal with interruptions in practical and embodied ways. My suggestion is that Nomadic Radio should retain its capacity to augment awareness by providing notifications of incoming messages, but that it should also leave users to control the notification levels through interactions that are both visible in and afforded by the system's design (i.e., intelligible). A simple design solution could be to replicate the

earpiece method so that users can dictate what level of notification they receive through the way they wear their headsets/earpieces. If both earpieces were in place, the notifications would be made in full, but if one or both were removed the notification levels would drop accordingly. More elegant solutions could no doubt be produced with input from interaction designers, but the point made here remains the same: design should be reasoned from how people organise and arrange information in material and embodied ways, and any system that augments awareness should continue in this vein. It's also worth pointing out that this change removes from the system the job of interpreting a user's openness to disruption. Instead, this interpretive work—that we've seen is so hard for computation to successfully pull off—remains in users' hands, supported through the material and embodied features of the system.

In a very roundabout way, this commentary on the Nomadic Radio gets me to the crux of this paper. I propose that we think about the intelligence we design into our environments in a different way from that implied through context-awareness. An environment's smartness should not be measured by how much and what variety of data it can sense, or by the complexity of things it can do with that data. Rather, we might imagine an environment's intelligence to be about how clearly the environment's workings are revealed to us (without making the world more complex), and how easily such workings can be harnessed. The business of designing intelligence into an environment should thus be concerned with building things that we can get to work intelligently for us.

4. REFERENCES

1. Beigl, M. (1999). *Using spatial co-location for coordination in ubiquitous computing environments*. HUC '99, Karlsruhe, Germany.
2. Beigl, M., Gellersen, H.-W., & Schmidt, A. (2001). MediaCups: experience with design and use of computer-augmented everyday artefacts. *Computer Networks*, 35, 401-409.
3. Bellotti, V., & Edwards, K. (2001). Intelligibility and accountability: human considerations in context-aware systems. *Human-Computer Interaction*, 16, 193-212.
4. Brown, J. S., & Newman, S. E. (1985). Issues in cognitive and social ergonomics: from our house to Bauhaus. *Human-Computer Interaction*, 1(4), 359-391.
5. Buxton, B. (1995). *Proximal Sensing: Supporting Context Sensitive Interaction*. UIST '95.
6. Clark, A. (2003). *Natural-born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford: Oxford University Press.
7. Dey, A. K. (2001). Understanding and using context. *Personal and Ubiquitous Computing*, 5(1), 4-7.
8. Dey, A. K., Abowd, G. D., & Salber, D. (2001). A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction*, 16, 97-166.
9. Dey, A. K., Mankoff, J., Abowd, G. D., & Carter, S. (2002). *Distributed mediation of ambiguous context in aware environments*. UIST 2002, Paris, France.
10. Dourish, P. (2004). What we talk about when we talk about context. *Personal and Ubiquitous Computing*, 8, 19-30.
11. Gaver, W. W. (1991). *Technology affordances: use of familiar things in the design of interfaces*. CHI '91, New Orleans, Louisiana.
12. Gellersen, H.-W., Beigl, M., & Krull, H. (1999). *The MediaCup: awareness technology embedded in an everyday object*. HUC '99, Karlsruhe, Germany.
13. Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. New York: Houghton Mifflin.
14. Grudin, J. (2001). Desituating action: digital representation of context. *Human-Computer Interaction*, 16, 269-286.
15. Hutchins, E. (1995). *Cognition in the Wild*. Cambridge: MIT Press.
16. Laurier, E. (2001). Why people say where they are during mobile phone calls. *Environment and Planning D: Society and Space*, 19(4), 485-504.
17. MacColl, I., Chalmers, M., Rogers, Y., & Smith, H. (2002). *Seamful ubiquity: beyond seamless integration*. Ubicomp 2002, Workshop on Models and Concepts for Ubiquitous Computing, Gothenburg.
18. Norman, D. A. (1988). *The Psychology of Everyday Things*. New York: Basic Books.
19. O'Connell, B., & Frolich, D. (1995). *Timespace in the workplace: dealing with interruptions*. CHI '95, Denver, CO.
20. O'Hara, K., & Sellen, A. (1997). *A comparison of reading paper and on-line documents*. CHI '97, Atlanta, GA.
21. O'Hara, K., Taylor, A. S., Newman, W., & Sellen, A. J. (2002). Understanding the materiality of writing from multiple sources. *International Journal of Human Computer Studies*, 56(3), 269-305.
22. Riesenbach, R., Buxton, B., Karam, G. M., Chattoe, J., & Leach, P. (1995). *Ontario Telepresence Project: Final Report*. Information Technology Research Centre: Telecommunications Research Institute of Ontario.
23. Sawhney, N., & Schmandt, C. (1999). *Nomadic Radio: Scaleable and Contextual Notification for Wearable Audio Messaging*. CHI '99.
24. Sawhney, N., & Schmandt, C. (2000). Nomadic Radio: Speech and audio interaction for contextual messaging in nomadic environments. *ACM Transactions on Computer-Human Interaction*, 7(3), 353-383.
25. Sellen, A. J., & Harper, R. (2002). *The Myth of the Paperless Office*. Cambridge, Mass.: MIT Press.
26. Weilenmann, A. (2003). Location and availability in mobile conversations. In *Doing Mobility: Doctoral Dissertation* (pp. 131-152). Göteborg: Department of Informatics, Göteborg University.
27. Weiser, M. (1994). *Creating the invisible interface: (invited talk)*. UIST94.