

Johnny Hartz Søraker and Philip Brey:

Ambient Intelligence and Problems with Inferring Desires from Behaviour

Abstract:

In this paper we will argue that many of the ethical problems raised by Ambient Intelligence stems from presupposing a behaviourist conception of the relation between human desires and behaviour. Insofar as Ambient Intelligence systems take overt, natural behaviour as input, they are likely to suffer from many of the same problems that have fuelled the widespread criticism of behaviourist explanations of human behaviour. If these limitations of the technology are not sufficiently recognized, the technology is likely to be insufficiently successful in supporting the needs and desires of human users. We will focus on four distinct challenges that result from this behaviourist presupposition, all of which ought to be taken into consideration at the design stage: reciprocal adaptation, bias towards isolated use, culture-specific behaviour, and inability to manually configure the system. By considering these issues, our purpose is to raise awareness of the ethical problems that can arise because of intelligent user interfaces that rely on natural, overt behaviour.

Agenda

What is Ambient Intelligence?.....	8
Problems with inferring desires from behaviour.....	9
Reciprocal adaptation.....	10
Bias towards isolated use.....	10
Cultural differences in behaviour.....	11
Inability to configure manually.....	11
Concluding remarks.....	11

Authors:

Johnny Hartz Søraker, Ph.D. Research Fellow:

Organization and contact address: Department of Philosophy, Faculty of Behavioural Sciences, University of Twente, Netherlands

Telephone, email and personal homepage: ☎ +31-53-489-4639, ✉ j.h.soraker@utwente.nl,
 🌐 http://www.utwente.nl/ceptes/research_staff/Soraker/

Prof. Dr. Philip Brey:

Organization and contact address: Department of Philosophy, Faculty of Behavioural Sciences, University of Twente, Netherlands

Telephone, email and personal homepage: ☎ +31-53-489-4426, ✉ p.a.e.brey@utwente.nl,
 🌐 <http://www.gw.utwente.nl/wijsb/medewerkers/brey/>

Ambient Intelligence is a vision in which computers play an increasingly pervasive yet unobtrusive part of our everyday lives. Whereas some hold that increased ubiquity alone will constitute a revolution in computing, others hold that it is not really a paradigmatic shift from more traditional forms of computing. In the words of pioneer Mark Weiser, "ubiquitous computing will produce nothing fundamentally new, but [make] everything faster and easier to do, with less strain and fewer mental gymnastics" (Weiser 1991:104). Although it is debatable whether ubiquitous computing introduces anything *fundamentally* new, it might come to exacerbate many of the ethical problems that arise as a result of our increasing dependence on computer technology. These problems include oft-debated issues such as invasion of privacy, identity theft, reduced autonomy and values-in-design.¹ Even if ubiquitous computing does not pose any unique problems, this is not a reason to ignore the phenomenon. To paraphrase Friedrich Engel's laws of dialectics, quantitative changes sometimes lead to qualitative changes. Our concern in this paper, however, is to argue that *Ambient Intelligence*, in virtue of adding Intelligent User Interfaces to ubiquitous computing, does introduce novel features that deserve special attention. Specifically, we will argue that AmI presupposes a behaviourist conception of the relation between human desires and behaviour. Insofar as we interact with AmI devices through natural, overt behaviour, we need to pay special attention to what kinds of behaviour these devices require, what kinds of desire-behaviour relations that are presupposed, and to what degree the required behaviour might be reinforced. Thus, rather than framing our discussion in terms of privacy, autonomy, risk or similar notions, we will focus on the functions and capabilities of Intelligent User Interfaces, in particular what kinds of behaviour they require and might come to foster. In doing so, we will propose and consider four distinct issues that signify when designers and engineers ought to pay special attention to the ethical and social impact of the behavioural requirements.

¹ See Brey (2006) and Tavani (2007:355-361) for an overview of some of the ethical issues that arise in connection with the Ambient Intelligence paradigm.

What is Ambient Intelligence?

Ambient Intelligence (AmI) is an approach that combines two major technologies: Ubiquitous Computing and Intelligent User Interfaces (IUI). In Ubiquitous Computing, computers do not appear as distinct objects, but are embedded into everyday working and living environments in an invisible and unobtrusive way. They make information, media and network access constantly and transparently available.² To the Ubiquitous Computing approach, AmI adds the technology of Intelligent User Interfaces. These interfaces, which are based on human-computer interaction research, go beyond traditional interfaces like the keyboard, mouse and monitor. They aim to make information technology easier to use by making interactions with it more intuitive, efficient, and secure; by "dissolving design in behaviour" (Greenfield 2006:26). As such, they are designed to allow the computer to know a lot more about users and the user environment than traditional interfaces can. Intelligent User Interfaces have two key features: profiling and context awareness. Profiling is the ability to personalize and automatically adapt to particular user behaviour patterns. Context awareness is the ability to adapt to different situations. Profiling and context awareness depend on sensors to record aspects of the environment and of user behaviour, and intelligent algorithms to make inferences about situations and users. IUIs are capable of creating a perceptive and proactive computer environment, rather than a passive one that relies on active and comprehensive user input.

One of the most interesting and novel aspects of Ambient Intelligence is the way human-computer interaction is redefined. The user interfaces of AmI seek to radically change the way we interact with computer technology – primarily by means of letting the computer infer our desires on the basis of overt and natural behaviour. The traditional way of issuing commands to a computer is by means of specially adapted peripherals such as mouse, keyboard or joystick. These traditional interfaces are limited in the sense that they require what we could refer to as "digital" behaviour – that is, discrete, non-natural actions that can easily be converted to digital input.

² Tavani (2007:356) makes a helpful distinction between pervasive *computing* and ubiquitous *communication*, but for the purposes of this paper we have subsumed these under the heading 'ubiquitous computing'.

Consequently, our behaviour in front of the computer is usually different from our behaviour away from the computer, which also means that we can easily distinguish between human-computer interactions and other actions.³ If we go beyond the traditional human-computer interfaces, there are primarily three different ways in which more natural, less discrete behaviour can be used to control computers. We will refer to these as different behaviour-desire relations – that is, different ways of inferring our desires (what we want the computer to do) on the basis of our behaviour:⁴

Pre-configured behaviour-desire relations:

The device can be manufactured in such a way that specific non-peripheral behaviour leads to the desired results. One simple example is the infamous “Clapper” technology, which allows the user to turn the lights on and off by means of clapping in a determinate way.

User-configured behaviour-desire relations:

The device can be manufactured in such a way that the users themselves can configure it to respond to specific behaviour. For instance, many mobile phones allow the user to record voice commands that correspond to specific functions.

User-adaptive behaviour-desire relations:

More advanced forms of user interfaces, and one of the cornerstones of Ambient Intelligence, is to let the device observe your natural behaviour and infer how your behaviour relates to your desires. For instance, Mark Weiser gives an example of an IUI in your bedroom that interprets restless rolling in the morning as an (imminent) desire for coffee (Weiser 1991:101).

One device can of course employ more than one of these interfaces, but the most interesting and unique challenges, and advantages, of Ambient Intelligence stem from *user-adaptive* systems. Having a computer system adapt to our behaviour

means that we do not have to configure it ourselves, which ensures that the technology disappears in the background. In order to become a transparent, unobtrusive technology that will effortlessly blend into our everyday lives, Ambient Intelligence depends on the successful implementation of user-adaptive interfaces. This is also where the unique challenges posed by Ambient Intelligence begin.

Problems with inferring desires from behaviour

Ambient Intelligence differs from traditional IT in the sense that we no longer consider what our desires are and interact with the device (behave) accordingly. Instead, we leave it up to the device itself to infer “what we really want” on the basis of our natural behaviour. In order for AmI to function optimally, it must therefore be possible to reliably infer certain human desires by way of observing behaviour alone. As such, AmI presupposes that behaviourist accounts of human behaviour are valid, at least for the application domain in question. This raises one of the most discussed issues in philosophy of mind and psychology: can desires be reliably inferred on the basis of behaviour alone? The near-consensus in psychology and in philosophy of mind is that this is not the case (see e.g. Fodor 1975; Searle 2001). The common view is that it is not single beliefs or desires that can be correlated with particular behaviours, but only complex webs of mental states. If I *want* coffee, for example, I may *take* the coffee in front of me, but only if I *believe* that the black liquid in the cup is coffee, I do not *believe* that the coffee is poisoned, and I do not *fear* that it is so hot I will burn myself, etcetera. Conversely, my coffee-taking behaviour may be caused by a desire for coffee, but also by a desire for the cup itself, a fear that a nearby child will spill the coffee over itself, or a belief that the cup contains tea, which I happen to desire. In spite of these kinds of problems, fully developed IUIs seem to presuppose a classical behaviourist account of the behaviour-desire relation in which desires can be reliably inferred from behaviour. This behaviourist underpinning gives rise to four challenges.⁵

³ To put it bluntly, when away from the computer we do not press our left finger twice when we open a document or tap our fingers on plastic keys when we communicate.

⁴ One could add highly advanced brain-computer interfaces to this list, which raises even more profound questions with regard to the relation between our desires and observable brain signals. We are still far away from seeing these kinds of technologies in widespread use, however.

⁵ Some of these challenges can be described as constraints on our autonomy. For a discussion in these terms, see Brey (2006).

Reciprocal adaptation

In a perfect world, we could envision intelligent user interfaces that reliably and accurately infer our desires from our behaviour, but this is not the case, neither when humans nor computers try to do so. To use a common example, a desire to escape pain does not necessarily lead to pain-aversive behaviour, and pain-aversive behaviour does not necessarily signify a desire to escape pain. As a result of this basic problem with behaviourism, successful interaction between humans and user-adaptive systems requires some adaptation on the human's part as well; we need to act in such a way that our desire becomes evident and predictable. In many ways, making a computer system adapt to your desires is similar to making a *pet* adapt to your desires. In order to properly train and command a pet animal, your behaviour must be discrete, predictable and overt, as opposed to vague, random and subtle. Since the artificial intelligence that underpins these user interfaces is unlikely to exceed the intelligence of most pet animals, we must adapt our behaviour in a similar fashion in our interactions with user-adaptive systems. Consequently, AmI is likely to make us change our natural behaviour to accommodate its limitations.

In this connection, it is also interesting to note that behaviourism is not only a theory of how to *explain* human behaviour. Although behaviourism has been largely discredited as an *explanatory* framework, its continued influence in psychology primarily stems from its ability to prescribe and predict how behaviour can change as a result of conditioning. Through concepts like positive and negative reinforcement, avoidance learning and habituation, behaviourism yields insight into how certain stimuli can lead to dramatic changes in our behaviour. Thus, if a user-adaptive system yields some kind of visual, auditory or tactile stimulus apt for conditioning, our behavioural adaptation to the system could become more entrenched, instinctive and even transferred to situations where we do not interact with the system at all. If behaviourism is correct in assuming that these mechanisms are particularly powerful with children, we should be especially aware of AmI devices that can reinforce behaviour in children.

In other words, not only the computer system will come to adapt its "behaviour" according to ours, it is also likely that we come to change our behaviour in order to effectively make the user-adaptive system comply with our desires. It should be noted that this is a problem with many other technologies as well. For instance, in order to watch TV, the user needs

to be located relatively still in front of the television set. As a consequence, TV does not only require immobility but the more it becomes a part of our lives the more it comes to *foster* that behaviour – what is sometimes referred to as the couch potato syndrome. If we add the hypothesis that couch potato behaviour is responsible for an alarming increase in obesity in many countries, then it becomes clear that behaviour fostered by technology can have profound implications. Although these kinds of affordances can be found in many technologies, AmI not only implicitly, but explicitly requires particular forms of behaviour. This is the reason why the behaviourist presuppositions of AmI deserve special attention. With a technology that is both designed to become a part of our everyday life *and* that explicitly requires certain forms of behaviour, we should be particularly aware of what kinds of behaviour such systems require and therefore might come to foster.

Bias towards isolated use

One design problem with AmI devices is that user-adaptation sometimes becomes difficult when multiple users interact with the same device. For instance, when your AmI-enabled TV has perfectly adapted to your desires and can anticipate your preferences after having observed your behaviour for a long time, you run the risk of losing that adaptation if someone else starts using it. Thus, the optimal adaptation of AmI devices often requires interaction with only one person, which in turn means that each user needs an individually tailored device. If we return to the previous analogy, a television set fosters sitting still in front of it, but it does not discriminate between watching it alone or together with other people. An AmI-enabled TV, on the other hand, might foster sitting still in front of it *alone*.⁶ It should be noted that this is not a general feature of all AmI devices. Compromises can often be found when the device manipulates variables that form a continuum, as in temperature regula-

⁶ Such individualized profiling also raises many of the same issues that Cass Sunstein has raised with regard to profiling on the Internet. According to Sunstein, social interaction and external deliberation is related to having had mutual experiences that can be the source of discussions in public fora and "around the water cooler". These mutual experiences, Sunstein claims, would be diminished if we all have individually tailored sources of information (Sunstein 2001).

tion, or when one device allows for multiple profiles. However, *if* an AmI device works best when used in isolation, it is likely to foster use in isolation as well. A similar worry is expressed in the recommendations of the Information Society and Technology Advisory Group (ISTAG). ISTAG stress that AmI should facilitate community building and provide “flexible participation in ... family/social interactions” (ISTAG 2003:10). For some AmI devices, the behavioural requirements will make it difficult to live up to this standard.

Cultural differences in behaviour

The most advanced AmI research and development centres are spread across the world, and we are likely to see AmI devices from both Western and East Asian countries. In order for AmI to function optimally, it is important that the behavioural input is natural and highly indicative of the underlying desire. However, what is seen as natural behaviour and how certain forms of behaviour relate to underlying desires depends to some degree on our cultural background. Behavioural indicators such as the range and importance of gesticulation, facial expressions and body language can differ radically from one culture to another. Problems regarding culture-specific forms of human-computer interaction is already an important issue in computer ethics (cf. Ess 2002), and these problems are likely to become more pressing as our interactions become more pervasive, ubiquitous and requiring reciprocal adaptation. For instance, some AmI-devices might discriminate against certain culture-specific forms of behaviour. Returning to reciprocal adaptation, globalization researchers have expressed concern over homogenization of cultural expressions as a result of technology being transported from one culture to another. AmI devices that require users to adapt to culture-specific forms of behaviour is one way in which such homogenization might occur.

Inability to configure manually

A common response to many objections raised against AmI is to simply include the possibility to override the user-adaptations and reset or configure the system manually if it misbehaves. This is somewhat question-begging, since the purpose of AmI is to make our interactions transparent and seamless, which is undermined if we constantly have to manually reconfigure the device in question. More to the point, given that many people are unable or unwilling to configure devices such as video recorders or mobile phones, it is a legitimate concern that many

will simply go along with whatever behaviour the AmI device requires. If we are dealing with AmI that targets multiple users, the ability to adjust the system individually could also mean that savvy users will have more influence on the system than others. In other words, a digital divide could arise between those who simply adapt to the required behaviour and those who are savvy enough to configure it manually.⁷

Concluding remarks

The purpose of this paper has not been to show that Ambient Intelligence necessarily leads to unwanted behaviour, nor that the fostering of certain kinds of behaviour is necessarily wrong.⁸ Rather, the purpose has been to show that insofar as an AmI device infers our desires based on natural, overt behaviour, designers and engineers need to pay special attention to what kinds of behaviour it requires – and to what extent it can reinforce this behaviour. This is especially the case if it 1) requires reciprocal adaptation, 2) has a bias towards isolated use, 3) requires culture-specific behaviour, or 4) cannot easily be configured or reset manually. These considerations become especially important when dealing with AmI devices targeted at children, given that they are more susceptible to reinforcement. If these and similar considerations are taken seriously at the design stage, we could avoid many of the societal and ethical implications that can arise from Ambient Intelligence.

References

Brey, Philip (2006). *Freedom and Privacy in Ambient Intelligence, Ethics and Information Technology* 7(3), 157-166

⁷ A related concern, especially with AmI systems used for public services, is that there could be no easy way of opting out of the required behaviour either – at least not by other means than giving up the service entirely (cf. Greenfield 2006:246-247).

⁸ An argument could be made to the effect that the explicit and pervasive behavioural requirements of AmI could be seen as wrong in-principle, in virtue of reducing our autonomy and/or constituting a questionable way of forcing the designers' values upon the end-user.

- Ess, Charles (2002). Cultures in Collision: Philosophical Lessons from Computer-Mediated Communication, Metaphilosophy 33 (1&2), 229–253*
- Fodor, Jerry (1975). The Language of Thought. Cambridge, MA, Harvard University Press*
- Greenfield, Adam (2006). Everywhere: The dawning age of ubiquitous computing. Berkeley, CA, New Riders Press*
- ISTAG (2003). Ambient Intelligence: from vision to reality (Working Group Report, September 2003). Retrieved Nov 1, 2007, from <http://cordis.europa.eu/ist/istag-reports.htm>*
- Searle, John (2001). Rationality in Action, Cambridge, MA, MIT Press*
- Sunstein, Cass (2001). Republic.com, Princeton, NJ, Princeton University Press*
- Tavani, Herman (2007). Ethics & Technology: Ethical Issues in an Age of Information and Communication Technology. Hoboken, NJ, John Wiley & Sons*
- Weiser, Mark (1991). The Computer for the 21st Century. Scientific American 265(3):94-104*