Pervasive Intelligibility: Workshop on Intelligibility and Control in Pervasive Computing

Jo Vermeulen
Hasselt University – tUL -- IBBT
Expertise Centre for Digital Media
Wetenschapspark 2,
B-3590 Diepenbeek, Belgium
jo.vermeulen@uhasselt.be

Brian Y. Lim
Carnegie Mellon University
5000 Forbes Ave., Pittsburgh,
PA 15213, USA
byl@cs.cmu.edu

Fahim Kawsar
Computing Department,
Lancaster University,
Lancaster, UK
f.kawsar@comp.lancs.ac.uk

ABSTRACT
With this workshop, we seek to provide a forum for exchanging design principles, programming techniques, toolkits and insights derived from real world studies towards building intelligible and user-controllable pervasive computing systems. Drawing upon the state-of-the-art, our goal is to refine existing and identify new directions for research in intelligibility and control for pervasive computing that will foster further work in the community.

Keywords
intelligibility; feedback; explanations; control; end-user configuration; pervasive computing; ubicomp; context-aware.

1. BACKGROUND AND MOTIVATION
The technological challenges outlined by Weiser's original vision of ubiquitous (or pervasive) computing [18] are rapidly being overcome. However, core issues concerning the usability and user experience of pervasive computing systems still remain to be solved. Interacting with these systems is very different from working with applications on a single (desktop) computer. Over the years, well-understood and proven heuristics and solutions have emerged for allowing end-users to understand and control their desktop computing environments. However, Bellotti et al. [3] state that these existing solutions are rarely adequate for — typically highly dynamic and adaptive — pervasive computing environments.

Due to the proactive and complex behavior of pervasive computing environments, it is especially important that systems are *intelligible* to allow users to understand "what the systems know, how they know it, and what they are doing" [4]. Additionally, systems should be *controllable* to let users recover when the system makes a mistake [4, 6, 7, 8, 9]. Previous studies have pointed out that users might become frustrated and lose trust in a pervasive computing system when they are unable to understand or control it (e.g., [2]).

Researchers have been calling for the support of intelligibility and control for the past decade (e.g., [2,4, 9], and even recently [10, 11]), and consequently researchers have provided tools and frameworks to support these requirements (e.g., [1, 8, 12, 16]), have looked into different user interfaces for intelligibility and control (e.g., [7, 17]) , and studied the impact of intelligibility and control in pervasive computing (e.g., [5, 13, 14 , 15]). We would like to organize a workshop to bring together researchers active and interested in intelligibility and control to further develop and refine this body of work.

We believe the time is ripe for the pervasive computing community to (i) formally identify user needs for intelligibility and control; build systems that are intelligible and user-controllable, supporting these systems through (ii) developing tools, toolkits, and architectures; and through (iii) developing design principles for building systems that allow users to scrutinize and inspect them for explanations on their state and functionality and control them henceforth. Users should be able to do so efficiently, efficaciously, and in a user-friendly manner. This requires (iv) appropriate evaluation criteria to judge whether pervasive computing systems are sufficiently intelligible and users are given an adequate level of control.

In this workshop, we will bring together researchers from different backgrounds who are involved with the design and development of user interface and interaction techniques, system building, studies of user needs for intelligibility and control, and evaluation of existing use of intelligible and controllable pervasive computing applications. Our primary goals are to refine existing and identify new research directions for intelligibility and control issues in pervasive computing systems; and to foster relationships for future collaboration.

2. OBJECTIVES
The workshop is structured around four concrete agendas:

1. **Understanding *user needs* for intelligibility and control.** What should be explained in order to improve the intelligibility of a pervasive computing system from an end-user perspective? What control mechanisms should be in place to ensure that users feel in control and trust the system? What is known about the interplay between intelligibility and control? Should control mechanisms be tuned to the type of intelligibility that systems provide? Which provisions for intelligibility and control should be available in all pervasive computing systems, and which ones would depend on factors such as the usage context (e.g., critical or urgent situations)? We are also interested in the ethnographic, cognitive psychological and social science theories underlying how people understand how applications work and seek to control them. These can help answer questions such as: What social factors influence what questions users ask? How do users ask questions? Why do users ask for certain types of explanations?

   Expected outcome: To create a taxonomy or framework of information and system states that lead to better intelligibility and to define a set of concrete control constructs that should

be addressed to facilitate end-user control in a pervasive computing environment.

2. **Understanding *technical requirements* regarding intelligibility and control.** What can technically be explained or controlled? Is it feasible to explain the reasoning behind complex machine learning algorithms (e.g., neural networks, Support Vector Machines)? If not, are there compromise explanations that would be sufficient? What level of human control is attainable for these advanced algorithms? How can we allow developers to easily support intelligibility or control in their applications? Is it possible to balance the trade-off between better intelligibility or control and an increase in development effort?

   Expected outcome: To categorize current practices in building pervasive computing architectures or toolkits that expose their application logic for better understanding and control, and use this categorization as a basis to identify gaps in existing approaches. Additionally, we would like to gain more insight into how machine learning algorithms can be adapted so that they can explain their internal workings and allow users to control certain aspects of their behavior.

3. **Investigating the *design space of user interfaces and interaction techniques* for intelligibility and control.** How can we effectively explain to users how a pervasive computing system works? How can non-technical and non-expert users be enabled to configure, personalize and correct their environment's behavior? In situations where no display is available (e.g., tangibles), where the available screen estate is limited (e.g., mobile devices, wearables), or where users' visual attention is required elsewhere (e.g., driving a car), it might be necessary to provide intelligibility through other modalities such as sound or touch.

   Expected outcome: To provide an overview of the current state of the art in user interfaces and interaction techniques for intelligibility and control, as well as develop ideas for novel user interfaces. With this discussion, we hope to contribute a set of design guidelines for intelligibility and control user interfaces.

4. **Defining a set of *evaluation criteria* for judging whether a system is sufficiently intelligible and puts enough power in the user's hands.** What do we expect will change when systems add support for intelligibility and control? How do we measure vague criteria such as trust or understanding? Although evaluation criteria such as effectiveness, subjective satisfaction, understanding or trust might seem to be logical choices, it is currently not clear whether these criteria are appropriate or if there other measures which are better suited. Moreover, it is necessary to determine how and under what conditions these evaluations should be performed (e.g., lab studies are unlikely to be sufficient).

   Expected outcome: To identify a set of measures and evaluation strategies that can be used to effectively evaluate a pervasive computing system's support for intelligibility or control.

We hope these four themes will provide a solid base for formulating future research directions and case studies for research on intelligibility and user-centric control in pervasive computing.

3. REFERENCES

[1] Assad, M., Carmichael, D.J., Kay, J. and Kummerfeld, B. 2007. PersonisAD: distributed, active, scrutable model framework for context-aware services. In *Proc. Pervasive 2007*. Springer-Verlag (2007), 55-72.

[2] Barkhuus, L. and Dey, A.K. Is Context-Aware Computing Taking Control away from the User? Three Levels of Interactivity Examined. In *Proc. Ubicomp '03*. Springer (2003), 149-156.

[3] Bellotti, V., Back, M., Edwards, W.K., Grinter, R.E., Henderson, A., and Lopes, C. Making sense of sensing systems: five questions for designers and researchers. In *Proc. CHI '02*, ACM (2002), 415-422.

[4] Bellotti, V. and Edwards, W.K. Intelligibility and accountability: human considerations in context-aware systems. *Hum.-Comput. Interact. 16*, 2 (2001), 193-212.

[5] Cheverst, K., Byun, H.E., 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*, 3-4 (2005), 235-273.

[6] Cheverst, K., Davies, N., Mitchell, K., and Efstratiou, C. Using Context as a Crystal Ball: Rewards and Pitfalls. *Personal Ubiquitous Comput. 5*, 1 (2001), 8-11.

[7] Coutaz, J. Meta-User Interfaces for Ambient Spaces. In *Proc. TAMODIA '06*, (2006), 1-15.

[8] Dey, A.K. and Newberger, A. Support for Context-aware Intelligibility and Control. In *Proc. CHI '09*, ACM (2009), 859-868.

[9] Edwards, W.K. and Grinter, R.E. At Home with Ubiquitous Computing: Seven Challenges. In *Proc. Ubicomp '01*, Springer-Verlag (2001), 256-272.

[10] Edwards, W.K., Newman, M.W., Poole, E.S. The infrastructure problem in HCI. In *Proc. CHI '10*, ACM (2010).

[11] Kyng, M. Making dreams come true: or how to avoid a living nightmare. In *Proc. Ubicomp '10*. 1-2. *Keynote*.

[12] Lim, B.Y. and Dey, A.K. 2010. Toolkit to support intelligibility in context-aware applications. In *Proc. of Ubicomp '10*. ACM (2010), 13-22.

[13] Lim, B. Y. and Dey, A. K. and Avrahami, D. Why and why not explanations improve the intelligibility of context-aware systems. In *Proc. CHI '09*. ACM (2009). 2119-2128.

[14] Stumpf, S., Rajaram, V., Li, L., Wong, W., Burnett, M., Dietterich, T., Sullivan, E., and Herlocker, J. Interacting meaningfully with machine learning systems: Three experiments. *Int. J. Hum.- Comput. Stud*. 67, 8 (Aug. 2009), 639-662.

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

[16] Vermeulen, J., Vanderhulst, G., Luyten, K. and Coninx, K. PervasiveCrystal: Asking and Answering Why and Why Not Questions about Pervasive Computing Applications. In *Proc. IE '10*, (2010).

[17] Vermeulen, J., Slenders, J., Luyten, K., and Coninx, K. I Bet You Look Good on the Wall: Making the Invisible Computer Visible. In *Proc. AmI '09*. 2009. 196-205.

[18] Weiser, M. The computer for the 21st century. *Scientific American 265*, 3 (1991), 66-75.