

# Exploring Patterns as a Framework for Embedding Consent Mechanisms in Human-Agent Collectives

Stuart Moran, Ewa Luger, and Tom Rodden

Mixed Reality Lab,  
University of Nottingham, UK  
{firstname.lastname}@nottingham.ac.uk

**Abstract.** With ever increasing developments in computing technology, approaches to attaining informed consent are becoming outdated. In light of this ongoing change, researchers have begun to propose several new mechanisms to meet the emerging challenges of consent in pervasive settings. Unfortunately a particular problem arises when considering consent in the context of Human-Agent Collectives (HACs). These large-scale heterogeneous networks, of multiple co-operating humans and agents are particularly complex and it is difficult to know *what*, *where* and *how* to introduce these new mechanisms. In this paper we explore the potential of patterns of interactional arrangement as a framework for embedding consent mechanisms in HACs and other ubiquitous systems.

## 1 Introduction

The process of embedding agency within computing systems raises unprecedented challenges to human privacy. The requirement for consent, as the means by which we agree to invasion of our private selves, is an ethical principle enshrined within our contemporary social expectations. Simply put, we expect to have a choice over who does or does not have access to certain aspects of our lives and, increasingly, our digital selves. Our interaction with the digital has occurred principally through tangible interfaces, which have to some extent been informed by social expectations and familiarity. An example of this is the re-interpretation of physically signing a paper document [2] as a digital checkbox [3] as a means of giving consent. However, the emergence of embedded computing brings with it a new set of challenges on how we deal with consent, as “not in a single one of these dimensions is the experience of [pervasive computing] anything like that of personal computing” [1]. When these ubiquitous computing systems are implicated within our daily interactions, it will likely break social expectations, becoming less clear how, when and where consent should occur. As such, not only the mechanisms, but also the concept, of attaining consent must evolve and adapt [4]. This motivated researchers to develop a number of new techniques and modes of interaction for attaining consent in ubiquitous computing. However, even with these new consent mechanisms, in complex systems it is difficult to effectively know which to use and when. Human-Agent Collectives [36] (HACs) are one example of particularly complex ubiquitous systems, as they consist of large-scale real-time networks of multiple teams of humans and software agents working in

collaboration, with different degrees of autonomy. In this paper, we draw upon the notion of patterns [35] as an analytical lens through which to formalize the recurrent interactional themes in HACs. These themes contribute toward an emerging typology, which could simplify the selection and implementation of consent mechanisms in ubiquitous computing systems.

## 2 Consent in Pervasive Computing

With the ongoing advancements in mobile and embedded devices, we are rapidly approaching an era of computing where technology literally pervades every aspect of our lives [5]. Pervasive computing sees a fundamental change to the way humans perceive and interact with computers. No longer will a system be limited to one device per user, but with several, even hundreds of, devices serving each person throughout their daily lives. Interaction with these devices will vary significantly. Some will work in the background, hidden from view, while others will require more direct interaction. This new technology, with its wide range of configurations, has the potential to create highly novel services and social environments [6]. However, the current lack of real world deployments mean that the social considerations of this technology are not yet fully understood [34]. Furthermore, this computational trend stands to challenge many of the well-established practices associated with current technologies. For example, the number and frequency of user interactions with pervasive technology are likely to increase in comparison to existing technology; but more disconcerting is that so too will the numbers of implicit/unknown interactions. This is especially problematic for current approaches to attaining consent, with dynamic systems rendering the idea of consent at a single point redundant [7].

Subsequently, the complexity of the infrastructure, interaction and the available choices (or lack thereof) requires a rethinking of the idea of consent in pervasive computing [7]. Consent, in its most theoretical form, must be (a) voluntary; the user must be free to give consent, it should not be coerced or the result of fraud, (b) competent; the user should be capable of giving consent; for example, they should be an adult and should not be otherwise vulnerable, (c) informed, in that the user should be meaningfully and fairly furnished with sufficient information of the conditions of the agreement, and (d) comprehending, in that the user should fully understand those conditions in order that there is a shared understanding between the consent-giver and consent-seeker [8]. Only at the point that all these conditions are met should a signal of assent be secured, for example the signature or mark of agreement as currently represented by the check box. Whilst consent in the context of pervasive systems has received only limited recent attention, some studies have sought to deal with designing for consent explicitly in the context of more traditional online interactions.

### 2.1 Designing for Consent

With particular reference to informed consent, Friedman et al. [9] offer a series of value sensitive design principles to shape the development of informed consent within

online systems; disclosure, comprehension, voluntariness (non-coercion), competence, agreement and minimal distraction. This latter point is considered important as undue distraction might undermine the act of consenting due to (a) the desensitizing effects and impact upon attention that could occur if a user is constantly notified to consent to minor issues, and (b) if the notification becomes too intrusive, the user may choose to bypass the consent process by ignoring or moving past the distraction altogether. Having made explicit the criteria of informed consent, the authors focus upon web-based interactions and seek to apply these criteria to the proposed redesign of both cookie notifications [10] and browsers in the belief that “informed consent provides a critical protection for privacy, and supports other human values such as autonomy and trust” [11 p.1]. Extending a value-driven view, the Principled Electronic Consent Management (ECM) is an approach suggested by Bonnici and Coles-Kemp [12]. The authors argue for a framework (consent theory, ECM norms, and manifestation of those norms) on the basis that (a) principled ECM addresses consent before and after the consent decision, (b) it considers a broad range of contributory factors such as both organizational and software processes, and (c) it builds upon theory in order to enhance consistency at the point of application [12].

Further studies also make suggestions in relation to how consent might be better designed within existing systems; though this is dominated by the field of bioethics, health and health data. Prasad and Kotz [13] make suggestions for solutions to the binary nature of consent management mechanisms through ‘privacy management interfaces’, which might incorporate the use of (a) privacy icons to support greater comprehension of risks to personal privacy (b) clearer ‘interface documentation’ designed to detail trade-offs/benefits related to each act of sharing, (c) the system could present beneficial recommendations based upon peer-behaviour, and (d) delegation of the decision to a specialist (e.g. in the case the doctor) [13].

In terms of the proposition informing consent, research has explored how the SMOG measure of text complexity can be used to assess whether terms and conditions can be understood by their readers [14]. A similar approach has also been developed to visualize such documents [15]. Whilst interesting, such approaches are not customizable to user preferences. More interestingly, it may be possible to automatically parse the consent-supporting text by machine and customize the output to suit the consent-giver (e.g. [16]). Given that consent is contextual [17], this could also be tailored to support differentiated consent scenarios, using the system to adapt the content and means of attaining consent. Hence, rather than undermining consent, system agency could be employed to support user agency in the consent process.

With much potential for adaptive technology, the question arises: what if a machine were able to act on our behalf and give consent? If we conceive of a raft of future systems, all interacting in a variety of different ways with different information, it is hard to imagine that humans will be able to process this information unsupported [18]. Hence it will become necessary to distribute some aspects of control and autonomy to intelligent software agents. Whilst seemingly dystopic, this is merely an extension of current practices of distributed cognition, such as using a computer for complex numeric processing or for setting reminders. The emerging field of human-agent collectives (HACs) aims to explore this type of interactional arrangement

in large-scale pervasive networks of multiple co-operating software agents and humans; which brings its own complexity to consent attainment.

### 3 Human-Agent Collectives

The idea of software agents has been around for a long time, and their relationship to humans has been explored in the field of HCI [19]. However, to date, there has been a tendency to focus on a limited scale one-to-one relationship between humans and agents. This is due to a focus on specific tasks and problems where agents can most effectively support humans. As the potential of the technology becomes increasingly evident, we can begin to explore new and more complex challenges that require us to move beyond the one-to-one relationship. Subsequently, a body of research has grown within the study of large groups of humans and software agents, in a variety of configurations, acting in what has been termed ‘Human-Agent Collectives’ (HACs) [36]. It is intended that these configurations will help solve, support and manage problems that humans cannot easily complete alone, such as extremely dynamic, cognitively demanding, and timed constrained tasks [37]. Disaster response is one example of a complex socio-technical problem where agents could assist humans in data collection within inaccessible areas (e.g. scouting aerial drones), and decision-making (e.g. knowledge of the ‘bigger’ picture that is incomprehensible to humans) [20]. HACs are envisaged as large, complex networks of real time interactions. These networks might consist of professionally trained personnel carrying out a task alongside members of the public who volunteer to help during a difficult time. Unlike current technologies, in HACs there is an intriguing inter-play between different configurations of professionals, members of the public and software agents and the distribution of agency amongst these actors [23]. Conceptually, consent is also predicated upon the idea of agency (the power to act) in the decision making process. *So, if HACs are to reconfigure the locus of agency, what does this mean for consent?* With such complex adaptable and dynamic networks of interaction, the difficulty and need to explore how consent to participation or exchange of information is attained or promoted becomes imperative; particularly so if agents are to act on our behalf.

#### 3.1 Spectrum of Control, Visibility and Complexity

With intelligent, adaptable and dynamic software agents come different degrees of autonomy; some agents may mimic the contribution of a human, and take control of a task entirely. The main point for consideration is that there is a spectrum of human-to-agent control during completion of a task. What makes this relationship particularly interesting is that the level of autonomy and agency may change during the course of an interaction, with humans handing over and taking control as and when is needed. This has significant implications for attaining and maintaining consent, as any change in levels of autonomy and agency need to be appreciated and understood by the

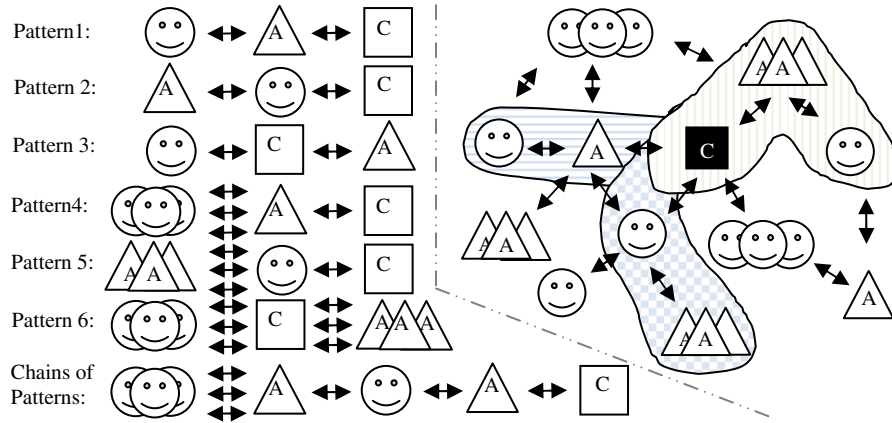
human; particularly if a machine is autonomously taking control. These different exchanges can also lead to agents invisibly performing processes in the background, while a human completes a task, unaware of that processing. Equally, an agent could be completely transparent about its actions, explicitly showing the user its actions as they occur. Software systems may act on a user's behalf in ways that they may not fully be aware of at the point of action. For this reason, it is necessary to consider the ways users will be kept fully informed about the system they are a part of, in addition to ensuring that they understand the choices they make and the actions they are offloading to machines. The ideas of a flexible spectrum of autonomy, and visibility of the processes involved, potentially add even further complexity to the system [36].

There is a clear need to understand and explore how individuals (and groups) might give consent to participate or exchange information in a HAC and also how they might use an agent to do this on their behalf. The difficulty in examining these questions is to know how best to make sense of the complex interactions, and the implications of the propagation of their effects. We suggest that what is needed is a framework with which to effectively capture and analyze the points of interaction and changes in autonomy. This framework must be scalable in order to deal with the different configurations of humans and agents. As such, in this paper, we propose to operationalize HACs into a series of discrete core configurations, or patterns, to allow consent-based mechanisms to be embedded in the interaction.

#### 4 Patterns of Consent

Many real-life applications can be modeled as a system of interacting actors. By using the idea of patterns [35] as a theoretical lens, examination of any such system could allow us to identify discrete interaction points; and thus highlighting instances within which consent might be embedded. Furthermore, breaking down and modeling the system in this way could also potentially allow us to visualize information trajectories (in the form of interactions and exchanges of information). These trajectories could then be translated into 'consent trajectories' that model the consequences of giving or withholding consent throughout the entire HAC. To begin to utilize such a framework, we must examine the most fundamental patterns of human-agent interaction. In order to identify core-interactive arrangements between software agents and humans, we examine a number of existing systems and their patterns of interaction, and consider what primary class of entity makes up a pattern.

In Fig.1. circles represent Humans, which are the users, the people who interact with the software/embodied agent or the core system. Triangles represent an Agent, which is an often semi-autonomous software/embodied agent which can act standalone or as an interface to the core system. Finally, squares represent the Core System, which is the technology that serves as the foundation of (or direct solution to) the task at hand. This could be in the form of storage, processing, operating, routing or even the main tool for the task (e.g. a plane). This is one aspect of the entire system



**Fig. 1.** Patterns of interactional arrangement (left), and example human agent collective (right)

which would include all of the above components. Initially we explore agent-interactions that are typically found in computing systems developed in recent years, focusing on an individual user and individual agent.

#### 4.1 Pattern 1: Intermediary/Proxy Interaction

An intermediary interaction is the most basic pattern (see Fig. 1.) where a human interacts with the core system through a software agent acting as an intermediary. This agent is essentially an intelligent user interface, and could take on a variety of forms, ranging from a purely virtual agent that supports tasks on a computer, to a physically embodied agent (e.g. a robot) through which information is exchanged. The point of interaction with this core system is via the agent, and in some ways the agent will be a representation and spokesperson for the human. The interaction could take different forms, including (a) the exchange and propagation of instructions to the core system, (b) the agent observing/monitoring the activities of the human, or (c) after an initial interaction the agent acts on the human's behalf. Consider a first responder during a disaster interacting with an unmanned ground vehicle (UGV) (e.g. [21]). This UGV relays richer information between HQ and the first responder than that which is possible using current approaches. Some other existing examples of this type of pattern/configuration include human-robot interaction [20], e-commerce systems [22] and mixed reality games [23]. The implications of this pattern in terms of consent centre on whether the agent supports or undermines the agency of the user. For example, the user may choose not to interact with the agent, but simply to perceive the agent as a standalone technology, rather than a gateway. In this way, user knowledge of the core system is undermined, raising a number of questions (related to awareness) that need to be addressed when users are exposed to this pattern of interaction: (a) *Does the user have control over whether the agent acts on their behalf?* (b) *Do users have an appreciation of the trajectories of their information?* (c) *Are*

*users aware of the wider-consequences of their consent-based choices in ‘the-moment’?*

#### **4.2 Pattern 2: Supervised/Advisory Interaction**

A supervisory interaction, while similar to intermediary, places the user in direct contact with the core system, while a software agent supervises/observes (see Fig. 1.). The interesting aspect of this pattern is the (possible) disconnect between the agent and the core system, with the agent becoming an almost independent entity. The agent may act as a prompter giving advice to the user as they complete their work. Equally, it may take on a more proactive role acting as a quality assurance. In some cases, the agent may even be a safety backup system that is programmed to act when the situation is deemed unsafe for the user. Consider a user being advised by a SatNav on directions to their specified destination. The SatNav is collecting data about the user’s actions (i.e. current location) and making recommendations for future actions to achieve the user’s goal. The user interacts directly with the core system (the car) which the SatNav has no control over (in this instance). Other examples might include agents which take control of planes [24], manufacturing systems [25] and intelligent surveillance systems that assess suspicious behaviors [26]. With the agent being a more independent entity this will require a degree of user approval prior to/during interaction, in addition to interaction with the core system. Where the agent intervenes in critical situations, one could imagine a model of consent where the user implies or infer consent as a default but inbuilt explicit moments of revocation? This unique arrangement highlights a number of key points for the process of consent in respect of voluntariness and agency: (a) *At what point does a user consent to the core system taking control? Can users take back control of the core system, after an agent has taken it?* (b) *Where does the accountability lie when an agent is instructing a human?* (c) *Do users have a choice whether they are advised?*

#### **4.3 Pattern 3: Remote/Tele-interaction**

A tele-interaction explores how a user can observe the actions and assume control of an agent (see Fig. 1.). The human uses the core system to remotely operate and receive information from an agent in the field, or at a significant distance. This type of pattern is typical in reconnaissance or in environments that are dangerous or impossible for humans to reach. The user may not necessarily directly control the agent, and may simply observe and analyze the information it transmits. The point of interest here is that the interaction occurs at a distance, with a greater disconnected experience with the agent. Consider two soldiers tele-operating a UAV drone. They have a variety of different information streams, and make decisions in real-time. This serves to allow soldiers to reach targets in difficult terrain, but also promote the safety of the lives of the soldiers themselves. Other examples of this type of system include: tele-robotics such as NASA mars landers and probes [27] and submersible mapping systems [28]. With a significantly greater disconnect between human and agent (rather than agent/core system), a series of important considerations should be made for

consent. These relate most directly to the issue of competence, autonomy and accountability: (a) *When should a user take control of an agent (or agent become autonomous)?* (b) *Who is accountable for the actions of the agent in situ; both during agent and human control?* (c) *How might humans in situ approve/engage with interaction with the agent?* (d) *What affect does tele-operation have on human judgment; during both automation and control?* Following on from a single user and agent, we now extend the patterns to include multiple agents and multiple humans. Applications that involve these numerous agents include crowd sourcing, crowd management, and manufacturing/assembly line computing systems.

#### 4.4 Pattern 4: Consolidated Intermediary Interaction

Consolidated intermediary interaction expands the intermediary pattern to include additional humans, which provide a larger corpus of information to an agent than an individual would (see Fig. 1.). Unlike the intermediary pattern, this agent has additional processing requirements in order to distinguish between the humans it interacts with, and the range of data types it receives. The main difference in this pattern is the unique exchange that may occur between the humans prior to interaction with the agent. Consider a busy pedestrian cross roads: the traffic light systems used to control thoroughfare during peak and off peak times are the same agent interacted with by both drivers and pedestrians (each who will act slightly differently, but generally conform to a series of rules/laws). The coordinating core system communicates information to different embodied agents, which then relay that information in real time. Users can interact with the core system through embodied agents that take the form of road crossing buttons. This pattern can also take other forms such as crowd sourcing [29], intelligent crowd control [30] or scheduling of shipping operations [31]. The same considerations as intermediary are applicable here (Pattern 1), but in addition we must also consider the factors brought about by the multiple human interactions, such as: (a) *How might a user exercise control of how their data is shared with others?* (b) *What information is exchanged between users?*

#### 4.5 Pattern 5: Consolidated Advisory Interaction

Consolidated advisory interaction adds multiple agents to the advisory pattern, which can provide different sources of information or even different interpretations of data (see Fig. 1.). The agents may take different physical and virtual forms, or interact in different ways (via audio, actuation or visually). In this pattern, the user has potentially an additional cognitive load added by the increased number of software agents observing/providing them with information. Consider a CCTV operator receiving multiple streams of information from different agents (cameras). This information is then used to interact with the core system, intended to relay important and timely information on to the appropriate persons. Equally, multiple agents may work to take control of different parts of the core system – e.g. an aeroplane may have multiple agents controlling multiple facets. Many of the same considerations apply from advisory interaction (Pattern 2), but the following should also be considered in respect of



consent: (a) *Is the user aware of the information exchanged between agents?* (b) *Does the user need to give consent to each agent (local vs global consent)* (c) *Can the user make sense of the information exchanged and produced from multiple agents?*

#### 4.6 Pattern 6: Common Pool

With the core patterns identified, we now consider the final more complex pattern, which is the common pool: multi-human multi-agent collectives working together via a core system (see Fig. 1.). This could take on a variety of different levels of autonomy and control, with humans and agents working independently, for or via each other. Consider multiple humans and multiple embodied agents working together to complete a task, such as in a car manufacturing plant, where there is a need for both humans and machines; they work together via the core system (the factory line). Some of the machines are autonomous, while others will be supervised/controlled by humans. Other examples of this pattern might include industrial tasks supported by robots [32], astronauts [33] and the envisaged applications of HACS in disaster response. This pattern is on the edge of succumbing to complexity. All the teleoperation considerations are important (Pattern 3), but also the social relationships between the people like in the advisory interaction (Pattern 2).

#### 4.7 Chains of Patterns

Having explored some of the possible fundamental patterns, we can see how these can join together in chains of patterns (see Fig. 1.). Patterns may repeat, or may connect via a core system, agent or human as the networks of interaction grow. Consider a pattern that joins remote telecommunication with intermediary interaction. A field responder might interact with a software agent that is controlled by a human, this combines these two patterns. It may be possible that within a pattern, the nature of the interaction can change, and indeed the pattern as well (depending on the role of the agent and its situation). For example, one moment a UGV could be relaying information between HQ and the user, the next it could make advisory points about the situation independent of the core system, then a human could control it.

## 5 Using Patterns

With a series of fundamental, but not necessarily exclusive, patterns identified we now have a means to explore the operationalization of complex large-scale HACs in order to embed consent mechanisms in the network. For example, a complex HAC, which consists of multiple humans and agents interacting, centered around a core system presents considerable challenges (see Fig. 1.). To explore aspects of consent in these collectives we must first break the HAC down into as many ‘major’ patterns as possible. There are likely to be individual and consolidated versions of intermediary and advisor patterns, all existing within a larger common pool pattern, including: Intermediary/Proxy Interaction (Horizontal Line), Agent-Consolidated Intermediary/Proxy Interaction (Vertical Lines), and Agent-Consolidated Advisory Interaction (Checkered) (see Fig. 1.).

### 5.1 Incorporating Consent Mechanisms

Having identified a number of patterns in the HAC, we examine each in more detail and attempt to resolve some of the considerations outlined above.

*Intermediary/Proxy Interaction (Horizontal Line):* The key things for consideration in this pattern are whether or not the human understands what information is collected and how it is used; in addition to whether or not a user has any control over the agent. If the core system is part of a planned interaction that occurs regularly, then ensuring the human is informed about and understands the purpose of the agent and the data it collects is adequate. However, if the interaction is momentary or ad hoc, then one recommendation would be to have the agent express its purpose and intent (in an accessible way), and to proceed only when a human has given approval.

*Agent-Consolidated Intermediary/Proxy Interaction (Vertical Lines):* In this pattern we see multiple consolidated agents acting as an intermediary to the main system. Again, consideration should be made about how aware the user is of the agents and whether or not the interaction is fleeting. Unlike the example given above this sees a single user interacting with multiple agents. With agents for different purposes it will likely be difficult for a user to fully comprehend the implications of their actions. For this reason, it would be beneficial to instigate a ‘representer-agent’, such that a user has a single point of control/interaction. This would reduce the cognitive burden on the user, and also create a more manageable situation. The agent may have periodic exchanges with the user about changes to the wider system, or a need for permission from one of the many agents in the pattern.

*Agent-Consolidated Advisory Interaction (Checkered):* This pattern sees multiple agents giving advice/information to a single user. It is questionable whether a user is able to fully comprehend the way that the agents exchange information amongst themselves, or the impact of providing consent to a single agent amongst others. Again, there is a clear need for an interface agent, one that can consolidate the advice from multiple agents and manage the exchange between users. There are also a number of questions with respect to agents taking control of the situation, with perhaps each agent taking control of different parts of the system. The way this information is expressed is critical, as is whether or not a user is able to take control back.

## 6 Conclusions

Consent is an important part of everyday life. With the ever increasing power of mobile computing systems the traditional approaches to attaining consent are no longer applicable. This has motivated researchers to explore new more appropriate consent-based mechanisms for the technology. However, complex large-scale heterogeneous networks of interacting software agents and humans pose new sets of problems. As such, we have proposed a nascent, exploratory use of patterns as a possible framework for breaking down complex systems into more easily examinable parts. Even our rudimentary overviews begin to offer some, key insights into the challenges arising. Examining each pattern more closely can help select the most appropriate approach to attaining consent, but also highlight many of the other socio-technical issues

that arise. The next steps for this work are to continue to explore and refine the core patterns of interactional arrangement, to develop a pattern language for HACs in order to cross-compare systems, and to empirically test the framework by introducing consent mechanisms to a real world deployed HAC.

## References

1. Greenfield, A.: *Everyware: The dawning age of ubiquitous computing*. New Riders, Berkeley (2006)
2. Wendler, D., Rackoff, J.E.: Informed consent and respecting autonomy: What's a signature got to do with it? *Ethics & Human Research* 23(3), 1–4 (2001)
3. Averitt, J.: Legal Ethics and the Internet: Defining a Lawyer's Professional Responsibility in a New Frontier. *J. Legal Prof.* 29, 171 (2004)
4. Luger, E.: Consent reconsidered; reframing consent for ubiquitous computing systems. In: *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, pp. 564–567 (2012)
5. Weiser, M.: The Computer for the 21st Century. *Scientific American* 265(3), 94–104 (1991)
6. Callaghan, V., Clarke, G., Chin, J.: Some socio-technical aspects of intelligent buildings and pervasive computing research. *Intelligent Buildings International* (1), 56–74 (2009)
7. Luger, E., Rodden, T.: Terms of Agreement: Rethinking Consent for Pervasive Computing. *Interacting with Computers* 25(3), 229–241 (2012)
8. Faden, R., Beauchamp, T.: *A History and Theory of Informed Consent*. Oxford University Press (1986)
9. Friedman, B., Lin, P., Miller, J.K.: Informed Consent by Design. In: Cranor, L.F., Garfinkel, S. (eds.) *Security and Usability*, pp. 503–529. O'Reilly Media Inc. (2005)
10. Millett, L.I., Friedman, B., Felten, E.: Cookies and Web browser design: Toward realizing informed consent online. In: *Proc. CHI 2001*, pp. 46–52. ACM Press (2001)
11. Friedman, B., Howe, D.C., Felten, E.: Informed Consent in the Mozilla Browser: Implementing Value Sensitive Design. In: *Proc. HICSS 2002*, vol. 8, pp. 247–257. IEEE Computer Society, Washington, DC (2002)
12. Bonnici, C.J., Coles-Kemp, L.: Principled Electronic Consent Management: A Research Framework. In: *Proc. 2010 International Conference on Emerging Security Technologies*, pp. 119–123. IEEE (2010)
13. Prasad, Kotz, D.: Can I access your data? Privacy management in mHealth. In: *Proc. USENIX Workshop on Health Security and Privacy*. USENIX Association (2010)
14. Luger, E., Moran, S., Rodden, T.: Consent for All: Revealing the Hidden Complexity of Terms and Conditions. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2687–2696 (2013)
15. Kay, M., Terry, M.: Textured Agreements: Re-envisioning Electronic Consent. In: *Proc. SOUPS 2010*, pp. 1–13 (2010)
16. Carreira, R., Crato, J.M., Gonçalves, D., Jorge, J.A.: Evaluating adaptive user profiles for news classification. In: *Proceedings of the 9th International Conference on Intelligent User Interfaces* (2004)
17. Luger, E., Rodden, T.: An Informed View on Consent for UbiComp. In: *Proc. UbiComp 2013*. ACM (2013)
18. Gonçalves, D.J., Jorge, J.A.: Ubiquitous access to documents: Using storytelling to alleviate cognitive problems. In: *Proceedings of the Tenth International Conference on Human-Computer Interaction*, pp. 374–378 (2003)

19. Green, S., Hurst, L., Nangle, B., Cunningham, P.: Software agents: A review (1997)
20. Fischer, J.E., Jiang, W., Moran, S.: AtomicOrchid: A Mixed Reality Game to Investigate Coordination in Disaster Response. In: Mobile Gaming Workshop (MOGA) 2012 as a part of Proceedings of the 11th International Conference on Entertainment Computing (ICEC 2012), pp. 572–577 (2012)
21. Murphy, R.R.: Human-robot interaction in rescue robotics. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* 34(2), 138–153 (2004)
22. Wang, N.D., Shen, X., Georganas, X.: A fuzzy logic based intelligent negotiation agent (FINA) in ecommerce. In: Electrical and Computer Engineering CCECE 2006, pp. 276–279. Canadian (2006)
23. Moran, S., Pantidi, N., Bachour, K., Fischer, J., Flintham, M., Rodden, T.: Team reactions to voiced agent instructions in a pervasive game. In: International Conference on Intelligent User Interfaces, pp. 371–382 (2013)
24. Vaščák, P., Kováčik, J., Betka, P., Sinčák, F.: Design of a Fuzzy Adaptive Autopilot. In: The State of the Art in Computational Intelligence, pp. 276–281 (2000)
25. Shen, W., Hao, Q., Yoon, H.J., Norrie, D.H.: Applications of agent-based systems in intelligent manufacturing: An updated review. *Advanced Engineering INFORMATICS* 20(4)
26. Hengstler, H., Prashanth, S., Fong, D., Aghajan, S.: MeshEye: a hybrid-resolution smart camera mote for applications in distributed intelligent surveillance. In: 6th International Symposium on Information Processing in Sensor Networks, p. 360 (2007)
27. Fong, T., Thorpe, C.: Vehicle teleoperation interfaces. *Autonomous Robots* 11(1), 9–18 (2001)
28. Williams, H., Newman, S.B., Dissanayake, P., Durrant-Whyte, G.: Autonomous underwater simultaneous localisation and map building. In: Conference on Robotics and Automation, vol. 2, pp. 1793–1798 (2000)
29. Savage, N.: Gaining wisdom from crowds. *Communications of the ACM* 55(3), 13–15 (2012)
30. Jin, J., Xu, X., Wang, J., Huang, C.C., Zhang, S.: Interactive control of large-crowd navigation in virtual environments using vector fields. *Computer Graphics and Applications* 28(6), 37–46 (2008)
31. Steenken, D., Voß, S., Stahlbock, R.: Container terminal operation and operations research—a classification and literature review. *OR Spectrum* 26(1), 3–49 (2004)
32. Heyer, C.: Human-robot interaction and future industrial robotics applications. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 4749–4754 (2010)
33. Carignan, D.L., Akin, C.R.: Using robots for astronaut training. *IEEE Control Systems* 23(2), 46–59 (2003)
34. Moran, S., de Vallejo, I.L.: Introduction to the special section on the social implications of embedded systems. *Interacting with Computers* 25(3) (2012)
35. Alexander, C., Ishikawa, S., Silverstein, M.: A pattern language. Ch Alexander 60 (2006)
36. Jennings, N.R., Moreau, L., Nicholson, D., Ramchurn, S.D., Roberts, S., Rodden, T., Rogers, A.: On human-agent collectives. *Communications of the ACM* (in press, 2014)
37. Fischer, J.E., Jiang, W., Kerne, A., Greenhalgh, C., Ramchurn, S.D., Reece, S., ... Rodden, T.: Supporting Team Coordination on the Ground: Requirements from a Mixed-Reality Game. In: Proc Intl. Conf. on Design of Cooperative Systems, COOP (2014)