
Consenting Agents: Semi-Autonomous Interactions for Ubiquitous Consent

Richard Gomer

r.gomer@soton.ac.uk

m.c. schraefel

mc@ecs.soton.ac.uk

Enrico Gerding

eg@ecs.soton.ac.uk

University of Southampton
SO17 1BJ, United Kingdom

UbiComp '14, September 13 - 17 2014, Seattle, WA, USA
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3047-3/14/09...\$15.00.
<http://dx.doi.org/10.1145/2638728.2641682>

Abstract

Ubiquitous computing, given a regulatory environment that seems to favor consent as a way to empower citizens, introduces the possibility of users being asked to make consent decisions in numerous everyday scenarios such as entering a supermarket or walking down the street. In this note we outline a model of semi-autonomous consent (SAC), in which preference elicitation is decoupled from the act of consenting itself, and explain how this could protect desirable properties of informed consent without overwhelming users. We also suggest some challenges that must be overcome to make SAC a reality.

Author Keywords

consent, ubiquitous computing, semi-autonomous, agent

ACM Classification Keywords

- Security and privacy~Usability in security and privacy
- Human-centered computing~Interaction design theory, concepts and paradigms
- *Human-centered computing~Ubiquitous computing*

INTRODUCTION

How does one consent to being tracked by a garbage bin? A project in London, the Renew Bin, caused a furor (and ban) when it was found to be tracking passers-by via their mobile phone's WiFi MAC address, without their consent. While this project was shut down by the local council [4], it brilliantly typifies the challenge of data capture and consent in ubiquitous, pervasive computing environments.

It is exactly this kind of challenge – about how to ensure citizens have meaningful consent interactions – that the Meaningful Consent¹ project seeks to explore.

In the following note, we outline one approach we're exploring to support consent that would work in pervasive computing environments. Essentially, we propose a consent architecture that first, takes consent decisions out of the current just-in-time approach we see embodied by the nagging responses to the EU "Cookie Law" [5] and shifts that discussion to a dedicated activity in which users explain their preferences. Preferences are then enacted by semi-autonomous agents in response to consent requests from third parties.

In the following sections we provide further description of our nascent architecture and look forward to the input and comments of this workshop.

Background

The processes through which one party obtains consent to some activity from another, which we term Consent Interactions, are commonly encountered in screen-

based contexts. For instance, the now common (if ineffectual) notices posed by websites that use cookies in the UK or the software license agreements to which users must consent during software installation.

Previous work has looked at the components of consent in screen-based media, for instance Friedman et al [2] who described six components of consent: disclosure, comprehension, voluntariness, competence, agreement and minimal distraction.

However, effects such as user habituation and the cognitive "shortcuts" that are taken by users pose challenges to informed consent even in these screen-based interactions, and ubiquitous computing devices pose more challenges given the diverse everyday environments in which they might be found.

In the European Union, at least, consent is becoming an increasingly important part of data protection regulation – Potentially empowering users and improving commercial data handling practices by exposing those practices more directly to market forces driven by consumer preference.

"Smart" street furniture, such as the Renew recycling bin [4], which is able to track nearby individuals using the IDs of their WiFi devices neatly illustrates the problem that ubiquitous devices can pose to building consent interactions. As devices themselves and the associated requirements to obtain consent become more ubiquitous, how can consent interactions be constructed in a way that satisfies the requirements of regulators, citizens and businesses?

¹ <http://www.meaningfulconsent.org>

The makers of the Renew bin described their device as being similar to a cookie for the street. In the EU, the use of web cookies for non-essential functionality now requires user consent. Given the similarity between online tracking via cookies and the stateless tracking possible via devices like the Renew bin, and the public outcry over the bins (and the subsequent ban imposed by the City of London) we imagine how a similar consent requirement could be operationalised for such devices.

Models of Consent

Informed consent generally involves two broad components, which we describe here as information (in which a person is provided with information) and consent (in which they signal that they are agreeable to the request that is being made). In offline media this process could take the form of reading and signing a physical form, and on a conventional computing device it often involves reading a notice and clicking a button.

Friedman et al [2] describe six components of informed consent as: Disclosure (providing adequate information), Comprehension (the individual having sufficient understanding of the provided information), Voluntariness (the ability for the individual to reasonably resist participation), Competence (the individual possessing the requisite mental, emotional and physical capabilities), Agreement (a reasonably clear opportunity to accept or decline participation) and Minimal Distraction (the consent process itself not being so overwhelming as to cause the individual to disengage from the process).

Another, more general, model of some relevance to the discussion of consent interactions in ubiquitous

contexts is the Communication-Human Information Processing (C-HIP) model, proposed by Wogalter et al [1] in the domain of risk communication. C-HIP has four main stages: Source (the originator of the risk information), Channel (the way that the information is transmitted from the source to the receiver), Receiver (the individual that is receiving the information) and Behavior (the response of the receiver to the information).

Implicit in Friedman et al's model of the components of consent is the need to get the user's attention in order to disclose information to them and to obtain their agreement. This requirement for user attention is explicit in the C-HIP model of information processing, where attention switch and maintenance is part of the Receiver stage.

The Role of Automation

Ubiquitous computing moves the context of consent decisions from computer-based activities to diverse new contexts such as walking down the street or entering a supermarket. As the need to make these consent decisions becomes more ubiquitous, so must our ability to do so.

Individuals will not accept violation of the "minimal distraction" principle as they go about these everyday tasks and so requiring them to make regular interactive consent decisions as soon as consent is required is not tractable. Equally, though, businesses will not simply accept that these new opportunities to understand customers are left untapped and consumers may also feel aggrieved at the opportunity cost of not pursuing these possibilities.

We suggest that semi-autonomous consent decisions, or even negotiations, present a compelling avenue for research, offering the potential to balance the possible economic value from new and innovative data uses with the consent requirements that policymakers and citizens want to see.

This approach uses a semi-autonomous software agent [3], acting on behalf of the user, to make consent decisions when consent is requested by a third party.

Our semi-autonomous consent (SAC) approach has three main phases: 1: Preference setting phase, in which a user expresses their preferences to the agent. 2: Consent phase, in which the agent responds to consent requests on behalf of the user, and 3: Review, during which the user can review the consent decisions that have been made, refine their preferences and provide additional information to their consent agent.

Fundamentally, SAC decouples the act of a user determining their preferences (disclosure, comprehension and competence) from the act of making a consent decision (voluntariness and agreement) and thereby protects the principle of minimal distraction.

Figure 1 shows the three stages involved in a semi-autonomous consent agent approach to ubiquitous consent, and shows how the factors involved in the C-HIP model of information processing and the six components of the Friedman model relate to it.

Phase 1: Preference Setting

During this phase, the user expresses their preferences to their consent agent. In the context of privacy and

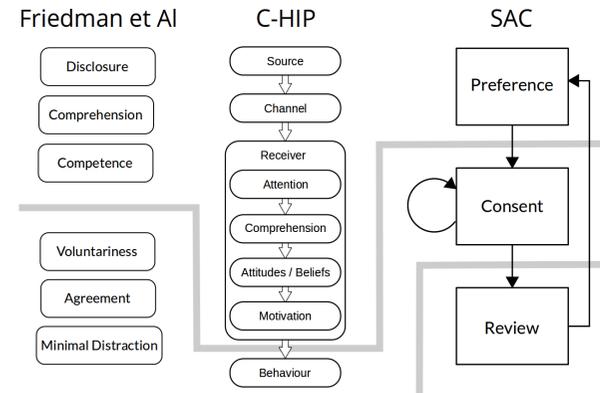


Figure 1: Relationship between Friedman model, C-HIP and SAC

personal data this could take the form of providing “pre-consent” to particular data uses or handling practices. The user might, for instance, express that they are happy for a store to attempt to record how often they pass by, but not to identify them individually or attempt to collate that information with their purchases. There are many conceivable models through which preferences could be expressed, ranging from simple enumeration of common consent scenarios to rule-based access-control or machine learning based on interactive decisions made by the user. One might even choose to defer decisions to another party (or their agent), perhaps a trusted friend.

This phase requires the user's attention, comprehension and competence and corresponds to the majority of the C-HIP model.

Phase 2: Consent phase

In this phase, the agent receives requests for consent and, based on the users expressed preferences, either grants or denies consent. This could happen many times with many different parties. The agent could conceivably use previous decisions to influence later ones, for instance by limiting the number of parties that are allowed to track the individual in a particular area on a first-ask-first-consent basis.

In this phase, the voluntariness of the consent interaction is enacted and agreement (potentially) reached. At this point, since the agent is acting on the user's behalf, the user is not required to shift their attention from the task at hand to deal with the consent request, and so the principle of minimal distraction is maintained. This phase corresponds functionally to the behavior phase of the C-HIP model, since it is the point at which the user's preferences and predefined course of action is enacted.

Phase 3: Review

In this phase, the user can review the consent decisions made by their agent and alter their expressed preferences. This could involve viewing how changes to their expressed preferences would have altered the decisions previously made by the agent.

The agent could request additional user input to resolve decisions that it had been unable to make during the consent phase – perhaps in a rare scenario that the user had not yet expressed a preference for.

Review may also encompass an aspect of auditing. Requesting parties could be required, through some

mechanism, to demonstrate their compliance with the policy that the agent had consented to.

Challenges

By decoupling, temporally, the act of informing and the act of consenting, we predict that SAC can protect properties of consent such as attention and minimal distraction even in ubiquitous environments. However, there are broad open questions, even beyond the agent-based technology that would underpin such a system, that require research before semi-autonomous consent can be operationalised. We present some of those challenges here, but look forward to identifying additional challenges with the other workshop participants.

Legal Issues

We expect that SAC will raise legal issues. Relying on an agent to determine whether a scenario is compatible with a user's consent preferences may, for instance, raise issues of liability in the event that the user perceives the decision to be wrong or detrimental to them. Determining where fault lies – with the user, agent or requesting party – could be non-trivial. Audit trails may help to settle these questions and allow retrospective introspection of the semi-automated decisions.

Abstraction

As with other consent scenarios, it is necessary to understand the correct level of abstraction to use when explaining consent choices to users, and hence their agents. An explanation such as “we will store a text file on your device” may explain the technical aspect of the choice but fail to convey the intent behind doing so; users are likely to be more concerned about whether

their movements will be logged (and for what purpose) than whether a negligible amount of storage space on their device is used to do so. Defining an ontology or taxonomy of consent that users can understand and which reflects their concerns may have implications for the representation at a technical level as well as the interaction level.

Interaction

As mentioned previously, a range of interaction models could be used to support the preference and review stages of SAC. Given the importance of the preference stage, designing interactions that are understandable and not undermined by the same shortcomings - such as cognitive shortcuts or inattention - as other consent scenarios remains a challenge. We are confident, though, that moving preference elicitation into a dedicated activity can provide benefits in this regard over the just-in-time model that is implicit in existing notice-and-consent models.

Conclusions

In this note we have outlined a semi-autonomous model of consent interactions that offers advantages over the 'just-in-time' interactions that we see today, especially in the context of ubiquitous computing which, given the social and regulatory context, promises (or threatens) to move consent interactions from screen based media to a diverse range of everyday situations.

We have also outlined some challenges - we are hopeful that by participating in this workshop we will begin to arrive at resolutions to these challenges, as well as identifying any other issues that must be tackled in order to make SAC possible.

References

- [1] Conzola, V. and Wogalter, M. A communication-human information processing (C-HIP) approach to warning effectiveness in the workplace. *Journal of Risk Research*, July 2001 (2001), 37-41.
- [2] Friedman, B., Lin, P., and Miller, J. Informed consent by design. In *Security and Usability*. 2005, 503-530.
- [3] Jennings, N.R. An agent-based approach for building complex software systems. *Communications of the ACM* 44, 4 (2001), 35-41.
- [4] Miller, J. City of London calls halt to smartphone tracking bins. *BBC News*, 2013. <http://www.bbc.co.uk/news/technology-23665490>.
- [5] DIRECTIVE 2009/136/EC OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL. European Union, 2009.