
How to Inform Privacy Agents on Preferred Level of User Control?

Jessica Colnago

Hélio Crestana Guardia

Federal University of São Carlos

São Carlos, SP, Brazil

jessica.colnago@dc.ufscar.br

helio@dc.ufscar.br

Abstract

This paper presents an organized set of variables that can aid intelligent privacy agents in predicting the best and necessary moments to interrupt users in order to give them control and awareness over their privacy, avoiding information overload or over choice.

Author Keywords

Privacy agents, user control, decision making.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

Ubicomp/ISWC'16 Adjunct, September 12 - 16, 2016, Heidelberg, Germany

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4462-3/16/09\$15.00

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

Introduction

In the Internet of Things (IoT) a large number of connected devices request data from users or user devices. Given the ubiquity of requests that require user attention and consideration, to sustain a reasonable level of privacy new solutions have to be developed. Solutions that use a "notice and consent" approach hit barriers that human attention and cognition are limited and that an excess of choice can be detrimental to privacy [2]. Solutions that aim at full automation of decisions have to deal with the fluid nature of privacy preferences [6] and the effects of removing user awareness and control [18]. The first suffers from information overload and over choice; the second from information vacuum and lack of control.

A solution format that seems able to balance these traditional approaches is based on intelligent privacy agents. These solutions are close to an automated approach but also rely on machine learning, making them capable of adapting to changes and requiring some level of user involvement. Existing solutions avoid interrupting the user unless it is strictly necessary, i.e. low inference confidence [8, 31, 33].

Literature Review Methodology

The review was performed in three stages and based on an exploratory approach.

First step: examination of the past 10 years of relevant publications such as Communications of the ACM, IEEE Security and Privacy, and IEEE Pervasive Computing, and relevant conferences, such as CHI, Ubicomp and SOUPS. From these publications references were obtained.

Second step: review of follow-up literature and references found in the works from the first-pass.

Third step: broad search using “privacy”, “privacy solutions” and “interruption” keywords was performed using ACM and IEEE digital libraries in order to reduce the chance of having overlooked significant literature.

While Copigneaux [10] allowed users to identify how and when to be interrupted, this was not expanded nor further described. And so, this paper comes to present which aspects affect the “when” and briefly presents the result of a literature review from privacy and interruption research. This culminated in the organization and discussion of a more complex, and we believe complete, set of variables that could be used to inform privacy agents not only when it is possible to interrupt the user, but when it is necessary and *desired from both the system's and the users' perspective*. This would filter requests avoiding the depletion of human attention, while still giving the user enough control and awareness to feel comfortable participating in the IoT.

Variables

The variables presented were selected given their preeminence throughout the literature review (see sidebar). From this review we identified variables that answer two user-asked questions when dealing with interruptions: (a) can I be interrupted now, and (b) do I want to be interrupted? (see Figure 1). In this context, the first deals with issues of interruptibility and receptivity to interruptions and focuses on interruption research; the second, with privacy related aspects that can influence users' desire to have more or less control and awareness over privacy decisions. It is important to note that the references presented here is only a subset of existing literature given a space limitation.

Can I Be Interrupted?

These variables are related to aspects of interruptibility, having a higher focus on interruptions research. They are highly contextual and dynamic and are related to the user context, activity context and social context.

Mood: The user's emotional and internal state, here referred as mood, has been considered an influential aspect in a user's interruptibility and availability in several previous works [11, 19, 26, 28, 29]. For Dabbish and Baker [11] the relation is less direct and made through a connection of their observed variable of “interruption threshold” defined as varying in accordance to context and external cues. Ho and Intille [19] and Sarker et al. [29] mention “emotional state of the user”, “affect” and “stress” directly as variables that influence the definition of interruptibility and availability, respectively. Finally, Pejovic and Musolesi [28] list the set of emotional states they considered in their exploration. For privacy-related interruptions the presence of an altered state of mind may lead the user to make different-than-usual decisions [9]. So it may not be ideal to interrupt the user to offer feedback.

Frequency of Interruptions: Frequency of interruptions has been considered a factor that influences user interruptibility [19] as well as one of the challenges considered for privacy feedback [3, 27]. It may also influence users' privacy concern over sharing a piece of data [15, 21]. When we consider privacy-related interruptions in the context of IoT this becomes an even more important because without an intelligent agent to mediate reception and decision, the frequency of interruptions can be superior to what is acceptable. Furthermore, Pejovic and Musolesi [28] found that the recent exposure to an interruption can influence and determine the user's frustration.

Activity Engagement: Activity engagement represents a combination of two relevant contexts: social and cognitive [14, 16]. In this work, the user's cognitive context, frequently considered in the

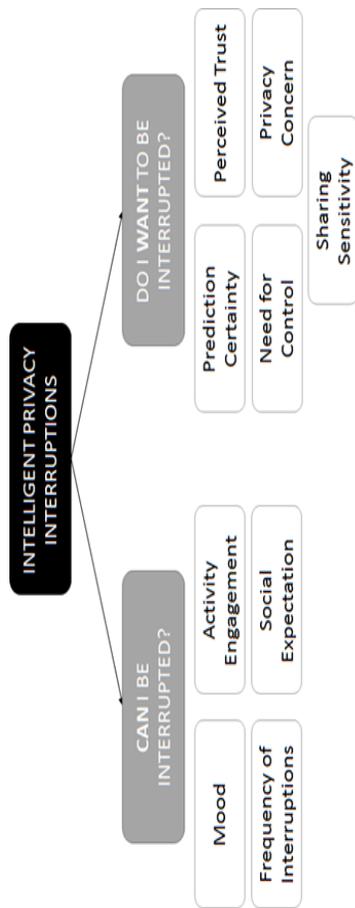


Figure 1. Intelligent Privacy Interruptions

literature of interruptibility [19, 26, 29], is defined as to “encompass the interruptee’s cognitive level of involvement in tasks and how it affects task performances” [16]. However, the definition of social context as “the interruptee’ physical environment as understood in a social sense” - relating to the place the user is in, the people around him/her and the nature of the task [16] – is closer to “social expectation”. Hence, the activity engagement variable considers only the social engagement of the user in his/her current task [17, 19]. The activity engagement variable is related to what has been previously considered as interruption threshold [11, 13], which also relates to the current activity type [29]. Ercolini and Kokar’s [13] four levels of interruptibility threshold can be combined to the four levels of activity engagement, when considering workload and social engagement as having two levels (see Figure 2). Lastly, in InterruptMe [28] one of the factors was named activity engagement. However, the factors considered were mainly related to the user cognitive context and aspects of the activity, not considering the user’s social involvement.

Social Expectation: This variable reflects the known importance that the social environment has on user behavior and decisions. From the perspective of user-technology interaction, the Technology Acceptance Model (TAM) [12] has added the factor of subjective norm since it became apparent that the acceptance of technology seems to be influenced by others [32]. The influence of the broad social context includes culture and social norms [5, 7, 34], the presence and opinions of others [5] and social interactions [26]. This has been studied in regard to the effect on user behavior, decisions and expectations towards privacy, including how much control over data sharing may be desired

[23], as well as its effects on the user’s interruptibility [9, 19, 25, 28] and the decision to interrupt somebody else [17].

Do I Want to Be Interrupted?

These variables explore the necessity and desire of the user to be interrupted to make a privacy-related decision. As such, it considers aspects from the data requests perspective, system perspective, user characteristics and system dependent characteristics.

Prediction Certainty: Prediction certainty relates to the decision for a particular interaction. It has been the main variable considered when deciding to interrupt the user for further input [8, 31, 33] and is defined solely by the system’s ability to correctly infer the user’s privacy preferences. Depending on the user and context its role may vary immensely. Users may not require a high level of certainty in exchange for a lower level of interruption in a social context that may not afford interruptions. But if users can be interrupted, they may desire a higher level of certainty on the prediction or make the decision by themselves with complete certainty. It is hard to consider this variable by itself because in isolation it is not expressive of the user’s preferences, only of the system’s needs. However, it could still be used as an indicator that an interruption may be necessary.

Perceived Trust: Trust has long been studied in the area of human-computer interaction (HCI). It is an important aspect of technology acceptance and, when considering automated systems in which the automated task is one the user can perform, the influence of trust increases significantly [20]. As highlighted by Bainbridge [4], the perception of the computer’s

		Social Engagement	
		Low	High
Workload	Low	At home relaxing (Bored)	Talking with a colleague (Flexible)
	High	Studying (Busy)	Giving a presentation (Do not disturb)

Figure 2. Activity Engagement representation with the combination of the two composing variables: social engagement and workload. The definition of where an activity fits in this matrix and the associated threshold level may vary from user to user.

abilities in automated systems influences the user's decision to allow the automation to continue or to override it. However, trust is a complex notion with many influencing factors [20]. In this work, we do not aim at defining what influences trust and how to quantify it; trust is viewed as a subjective user-dependent and dynamic variable which is perceived by the user in the moment of the interaction. Similar to TAM, the users' perceived trust in technology is what we consider for the evaluation of this variable.

Need for Control: Control is also an important aspect of HCI. Its lack increases anxiety and stress when dealing with computational systems [18]. Moreover, the perceived control over the interrupting device has been identified by Ho and Intille [19] as a factor that influences the user's perceived burden of the interruption. From the user's perspective, control is a "personality trait that reflects individual differences in the appreciation of choice in life" [18]. Some people have a stronger desire and need for control than others, and previous research has identified these nuances in particular with agent interactions [30].

User's Privacy Concern: The user's privacy concern should play a significant role in determining not only the disclose decision, but on whether or not the user should be interrupted for further input. As a way to abstract privacy concern and associate it with possible behaviors we consider Alan Westin's privacy indexes and categorizations [22] which classify users as fundamentalists, pragmatists and unconcerned in decreasing level of privacy concern and openness to external factors. Other tools for privacy concern classification, such as IUIPC [24], may be more appropriate. However, the simplicity of Westin's

categorization is appealing given the overall complexity of the variable set. Finally, because of the more static aspect of this variable and the influence it may have on the weight given to other variables, it could be used to inform the behavior in unknown situations on top of informing the decision to interrupt.

Sharing Sensitivity: Interruption literature has reported that interruptions that provide useful information are viewed more positively [11, 19, 25, 28]. This has also been noticed in the context of privacy decisions: if the reward outweighs the cost of sharing, it seems rational to do so [1]. Even though every privacy-related interruptions should be considered important and worthy of an interruption, some requests can have a higher risk/costs than others. Thus being perceived as more important to be dealt with personally. For this reason, utility and importance of interruptions was adapted to sharing sensitivity. This work associates higher privacy concerns when sharing data with a higher sharing sensitivity. A literature review related to user's data sharing privacy concern has elicited a large number of variables – organized as *what*, *who*, *why*, *when*, *where*, *how* and *context*. Because *when*, *where* and *context* can be better identified by the user and *who* is requesting the data, *why* they need it and *what* they need, have had a higher impact on the user's concern when sharing data [9, 23, 26], these variables can be used to classify sharing sensitivity. This subset is used in notice and consent in mobile platforms (Figure 3).

Conclusion

The fact that this set of variables was extracted through observation of several and diverse literature, affords it a higher generality. However, this is an initial

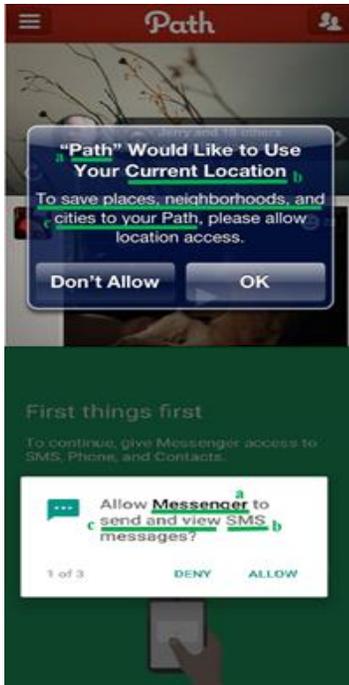


Figure 3. Permission request notification identifying (a) the app requesting information (who); (b) what data they need access to (what); and, (c) why they need access to it (why). (top) iOS 6. (bottom) Android 6.0.

exploration of relevant concepts and it has not been tested in the scope of a real privacy agent in a real IoT environment. Nevertheless, we believe that this issue's discussion and the existence a reference material can be of assistance to the developers of such systems.

References

1. Alessandro Acquisti. 2013. Complementary Perspectives on Privacy and Security: Economics. *IEEE Secur. Priv.* 11, 2, 93–95.
2. Alessandro Acquisti, Idris Adjerid, and Laura Brandimarte. 2013. Gone in 15 Seconds: The Limits of Privacy Transparency and Control. *IEEE Secur. Priv.* 11, 4, 72–74.
3. Hazim Almuhammedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your Location has been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proc. ACM Conference on Human Factors in Computing Systems (CHI '15)*. 787–796.
4. Lianne Bainbridge. 1983. Brief paper: Ironies of automation. *Automatica* 19, 6, 775–779.
5. Louise Barkhuus. 2012. The mismeasurement of privacy: using contextual integrity to reconsider privacy in HCI. In *Proc. ACM Conference on Human Factors in Computing Systems (CHI '12)*. 367–376.
6. Victoria Bellotti and Abigail Sellen. 1993. Design for privacy in ubiquitous computing environments. In *Proc. European Conference on Computer-Supported Cooperative Work (ECSCW'93)*. 77–92.
7. Michael Boyle and Saul Greenberg. 2005. The Language of Privacy: Learning from Video Media Space Analysis and Design. *ACM Trans. Comput. Interact.* 12, 2, 328–370.
8. Christian Bunnig and Clemens H. Cap. 2009. Ad Hoc Privacy Management in Ubiquitous Computing Environments. In *Proc. International Conference on Advances in Human-Oriented and Personalized Mechanisms, Technologies, and Services*. 85–90.
9. Sunny Consolvo, Ian E. Smith, Tara Matthews, Anthony LaMarca, Jason Tabert, and Pauline Powledge. 2005. Location Disclosure to Social Relations: Why, When & What People Want to Share. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. 81–90.
10. Bertrand Copigneaux. 2014. Semi-autonomous, context-aware, agent using behaviour modelling and reputation systems to authorize data operation in the Internet of Things. In *Proc. World Forum on Internet of Things (WF-IoT)*. 411–416.
11. Laura A. Dabbish and Ryan S. Baker. 2003. Administrative assistants as interruption mediators. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems (CHI EA '03)*. 1020–1021.
12. Fred D. Davis, Richard P. Bagozzi, and Paul R. Warshaw. 1989. User acceptance of computer technology: a comparison of two theoretical models. *Manage. Sci.* 35, 8, 982–1003.
13. Deborah Ann Guerrera Ercolini and Mieczyslaw M. Kokar. 1997. Desktop Agent Manager (DAM): Decision Mechanism. *Int. J. Hum. Comput. Interact.* 9, 2, 133–149.
14. Robert Fisher and Reid Simmons. 2011. Smartphone Interruptibility Using Density-Weighted Uncertainty Sampling with Reinforcement Learning. In *Proc. International Conference on Machine Learning and Applications and Workshops (ICMLA '11)*. 436–441.
15. Nirmal Gaud, Anjana Deen, and Sanjay Silakari. 2012. Architecture for discovery of context-aware web services based on privacy preferences. In *Proc. Conference on Computational Intelligence and Communication Networks*. 887–892.
16. Sureshini Grandhi and Quentin Jones. 2010. Technology-mediated interruption management. *Int. J. Hum. Comput. Stud.* 68, 5, 288–306.
17. Rikard Harr and Victor Kaptelinin. 2012. Interrupting or not: Exploring the effect of social context on interrupters' decision making. In *Proc. Nordic Conference on Human-Computer Interaction Design (NordiCHI '12)*. 707–710.

18. Hans Van Der Heijden. 2003. Ubiquitous computing, user control, and user performance: conceptual model and preliminary experimental design. In *Proc. Research Symposium on Emerging Electronic Markets*. 107–112.
19. Joyce Ho and Stephen S. Intille. 2005. Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. In *Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*. 909–918.
20. Kevin Hoff and Masooda Bashir. 2013. A Theoretical Model for Trust in Automated Systems. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*. 115–120.
21. Jason I. Hong and James A. Landay. 2004. An architecture for privacy-sensitive ubiquitous computing. In *Proc. Conference on Mobile systems, applications, and services (MobiSys '04)*. 177–189.
22. Ponnurangam Kumaraguru and Lf Cranor. 2005. Privacy indexes: A survey of westin's studies, Pittsburg, PA.
23. Scott Lederer, Jennifer Mankoff, and Anind K. Dey. 2003. Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. In *CHI '03 Extended Abstracts on Human Factors in Computing Systems (CHI EA '03)*. 724–725.
24. Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Inf. Syst. Res.* 15, 4 (2004), 336–355.
25. Abhinav Mehrotra, Mirco Musolesi, Robert Hendley, and Veljko Pejovic. 2015. Designing content-driven intelligent notification mechanisms for mobile applications. In *Proc. ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. 813–824.
26. Antti Oulasvirta and Antti Salovaara. 2004. A cognitive meta-analysis of design approaches to interruptions in intelligent environments. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems (CHI EA '04)*. 1155–1158.
27. Sameer Patil, Roberto Hoyle, Roman Schlegel, Apu Kapadia, and Adam J. Lee. 2015. Interrupt Now or Inform Later? Comparing Immediate and Delayed Privacy Feedback. In *Proc. ACM Conference on Human Factors in Computing Systems (CHI '15)*. 1415–1418.
28. Veljko Pejovic and Mirco Musolesi. 2014. InterruptMe: designing intelligent prompting mechanisms for pervasive applications. In *Proc. ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. 897–908.
29. Veljko Pejovic, Mirco Musolesi, and Abhinav Mehrotra. 2015. Investigating The Role of Task Engagement in Mobile Interruptibility. In *Proc. Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. 1100–1105.
30. Hillol Sarker, Moushumi Sharmin, Amin Ahsan Ali, Md. Mahbubur Rahman, Rummana Bari, Syed Monowar Hossain, and Santosh Kumar. 2014. Assessing the Availability of Users to Engage in Just-In-Time Intervention in the Natural Environment. In *Proc. International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)*. 909–920.
31. Florian Schaub, Bastian Könings, Michael Weber, and Frank Kargl. 2012. Towards Context Adaptive Privacy Decisions in Ubiquitous Computing. In *Proc. International Conference on Pervasive Computing and Communications Workshops*. 407–410.
32. Gunnvald B. Svendsen, Jan-Are K. Johnsen, Live Almås-Sørensen, and Joar Vittersø. 2013. Personality and technology acceptance: the influence of personality factors on the core constructs of the Technology Acceptance Model. *Behav. Inf. Technol.* 32, 4, 323–334.
33. Eran Toch. 2011. Super-Ego: A Framework for Privacy-Sensitive Bounded Context-Awareness. In *Proc. International Workshop on Context-Awareness for Self-Managing Systems (CASEMANS '11)*. 24–32.
34. Alan Westin. 1967. *Privacy and Freedom*, The Bodley Head Ltd