Semantic Web Technologies for Social Translucence and Privacy Mirrors on the Web

Mathieu d'Aquin and Keerthi Thomas

Knowledge Media Institute, The Open University, Milton Keynes, UK
{mathieu.daquin, keerthi.thomas}@open.ac.uk

Abstract. While the open, collaborative and distributed nature of the Web makes it an effective social platform, it also makes it difficult to implement appropriate privacy management features. In this paper, we discuss the notions of social translucence and privacy mirrors, originally defined to be used as guiding principles for the design of (ubiquitous) systems implementing some form of social processes, and how such principles apply to privacy management for online interactions. The focus of the paper is therefore on the technological challenges that are raised in trying to implement these principles (visibility, awareness, accountability) in an open Web context. We show through several examples of such implementations how the issues of data integration, data processing and the need for inference that are raised by such approaches provide an interesting use of semantic technologies, for individual monitoring and sense-making of personal information on the Web.

1 Introduction

As argued for example in [20], privacy is a complex notion to define, depending on the context and the purpose of providing a definition. [19] includes a survey of the different views on privacy, alongside different perspectives. Besides the very naive ones (e.g., "privacy is withdrawing information from other"), one of the most natural definitions of privacy is the one described by [11] as "privacy is not simply an absence of information about us in the minds of others, rather it is the control we have over information about ourselves." Similarly, Altman [2] (as cited by [17]), defines privacy as "the selective control of access to the self".

As some have shown [1], any control of personal information is determined by the user's perception of privacy - what information users regard as private, from whom, and in what context, and depending on the privacy risks how the users might make trade-offs. Thus, user's control of personal information directly relates to feedback and awareness being necessary for privacy management, and systems based on the exchange of personal information should be "socially translucent system", i.e., systems that "support coherent behavior by making participants and their activities visible to one another" [10]. There is no arguing that social, professional or commercial interactions on the Web rely extensively on the exchange of private, personal information. However, on the

Web, the circulation of such information is happening in an un-restrained, fragmented and distributed environment. We therefore believe that one of the core, and immensely difficult challenge for privacy is to enable the three principles of social translucence–visibility, awareness and accountability [10]–with respect to the open and unbounded exchange of personal information on the Web [6].

In this paper, we argue that, to tackle the technological aspects[1] of such a challenge–achieving social translucence for what concern Web-mediated communications–Semantic Web technologies are not only useful but they are needed. We show through preliminary models of systems related to what is described in [18] as "privacy mirrors", how the data integration, analytics and sense-making capabilities that Semantic Web technologies can deliver in the complex and distributed context of Web (personal) information can lead to increasing individuals' control over their own personal information, through raising their awareness of the privacy implications of their actions on the Web.

2 Social Translucence and Privacy Mirrors

The notion of social translucence (as defined in [10]) is broader than its application to privacy management. It concerns the design of systems with a social process component, on aspects where coherent behaviours from the user(s) are important. It is argued that, to achieve such coherent behaviours, it is necessary for the system to make such behaviours visible and understandable to the users. In other words, while it might not require to be fully transparent on every aspects of its operations, the system should let the user see the relevant information for them to make sense of their own activities, so to adequately tune their use of the system, and of the social interactions it enables. [10] therefore characterises translucent systems according to three main principles: visibility, awareness and accountability.

There are obvious connections between the idea of social translucence and privacy. As an objection to it, one might argue for example that social translucent systems need to be designed carefully, so that visibility does not come to contradict privacy. It is indeed the main argument for describing such systems as translucent: they need to make enough information available to the user to enable their informed use, but not to become so transparent as to enable unintended use of the information made visible against the user.

A more constructive connection however between social translucent systems and privacy is one where the principles of visibility, awareness and accountability are used to enable a coherent and informed behaviour from the users with respect to the distribution and propagation of their personal information. This idea is especially well illustrated in "Privacy Mirrors", a framework for designing

[1] While we are aware that it is only a restricted view, and that other aspects related for example to policies and the law are crucial, in this paper, we explicitly focus only on the technological and tool-support aspects of Web privacy. The interested reader might refer to [21], which makes the connection between the work presented here and some non-technological considerations.

socio-technical (ubiquitous) systems so that they integrate the necessary tools of "awareness and control to understand and shape the behaviour of the system" [18]. It is based on five basic characteristics: history, feedback, awareness, accountability and change.

Both social translucence and privacy mirrors have been defined in the context of the design of systems implementing some form of social processes. However, as already argued in [6], improving control and privacy management over online interactions would require to achieve similar principles and characteristics in a pre-existing, complex and open system (the Web) making it much more challenging. Focusing on the first basic principles of privacy mirrors, providing history, feedback and awareness on a user's social activities on the Web raises technological issues that are not present in closed, designed (even if ubiquitous) systems: the need to collect information from various sources, to integrate it in a coherent, manageable view, to analyse it and make sense of it in an open context (i.e., not relying on domain and system specific assumptions), etc. These are the challenges that we believe need to be tackled to achieve a more translucent approach to Web interactions, and where Semantic Web technologies can play a major role in providing a more informed way of privacy management for Web users.

In the next sections, we discuss some of these aspects through examples of using Semantic Web technologies to improve the user's visibility and awareness of their own Web activities.

3 Data Integration for Web Interactions

Making information visible to users about their own activity in a system that implements social processes first requires for the system to collect such information in a way which is manipulable and transferable to the user. There are however specific challenges that appear when trying to achieve this on online activity data. Activity data typically sit in the logs of websites and Web applications, and are exploited by online organisations, in an aggregated form, to provide overviews of the interactions between the organisation's online presence and its users (most commonly in the form of website analytics). In [5] we looked at the technological challenges that are faced when trying to invert this perspective on activity data: provide individual users with an overview of their interactions with the online organisation. Such challenges can be summarised as:

Fragmentation and heterogeneity: Activity data is typically held in log files that have different formats, and might not be easily integratable from one system (website, application) to another.

User identification: Recognising and identifying a user within the data is typically a problem faced by any activity data analysis. However, when taking a user-centric perspective, a user needs to be identified over several systems and the consequences of inaccurately recognising a user can be more critical.

Data analysis: Activity data is generally available through raw, uninterpreted logs from which meaningful information is hard to obtain.

Scale: Tracking user activities through logs can generate immense amounts of data. Typical systems cope with such scale through aggregating data based on clusters of users. Here, we need to keep the whole set of data for each individual user available to provide meaningful analysis of their interaction with the organisation in a user-centric way.

In [5], we showed how we investigated and handled these challenges through relying on semantic technologies, especially RDF for the low level integration and management of data, ontologies for the aggregation of heterogeneous data and their interpretation, and lightweight ontological reasoning to support customisable analysis of user-centric activity data. This involved in particular building tools to:

1. Convert and integrate the data from their log format into RDF, following the schema provided by given ontologies.
2. Create ontology level definitions of types of resources and activities to enrich the initial data, to then apply ontology reasoning as a way to classify traces of activities according to these ontological definitions.
3. Realise additional ad-hoc processing of the data, to improve interpretability. This included in particular deriving the location of the user from their IP (using a dedicated online API), deriving human readable labels for the user agent (e.g., "Chrome 8") from the complex user agent strings included in the logs, deriving general date/time information from the timestamp included in the logs (e.g. day of the week), etc.

As discussed in the next section, even within the limited scope of information systems within one organisation, such semantic data integration of a user's interactions with online systems made it possible to create user-centric views over such interactions, making such activities visible to them and leading to a greater awareness of their own behaviour.

Taking a broader view on activity data, in [3] we experimented with the user-centric collection, aggregation and processing of information regarding an individual's interaction with Web systems. As shown in Figure 1, the idea there was to intercept all the Web traffic coming out of a user's computer, and integrate the relevant information into an RDF representation, based on a simple generic ontology of Web interactions. A number of studies were made on the basis of this tool, collecting large scale information for individual users (100 million triples for a single user over 2.5 months). In [9], we showed how, based on such data, we could reconstruct a complete user profile, aggregating and mapping personal information regarding dozens of attributes from the exchanges of such information with thousands of websites. This demonstrated in particular how the basic principles of Semantic Web technologies (linking and sense making out of heterogeneous information) can support the construction of a user-centric and user-interpretable view of an originally fragmented set of interactions and exchanges.

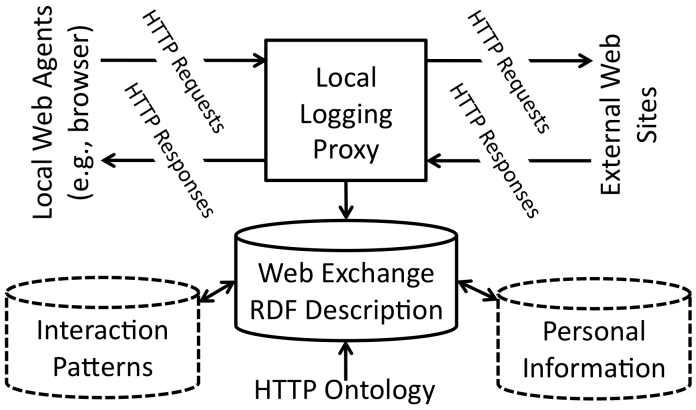

Fig. 1. A local proxy system for collecting information about a user's Web interactions.

A common objection to this approach (creating an aggregated view of personal data for the benefit of the user) is that it might itself create a privacy threat. We qualify such an objection as naive, considering that it is merely making available to the user a process that is otherwise in use by large online organisations outside of the user's reach and control. Of course, we do not mean to say that it does not create additional risks. It is indeed crucial that such risks are appropriately tackled through the secure, local storage and access to the collected data in any implementation of the approach. However, since we focus here on the role of semantic technologies, we see these aspects and their implications on computing resource requirements outside the scope of the paper.

4 Personal Analytics: Data Analytics on Personal Information

Of course, building a user-centric view over activity data is not sufficient to make this information visible to the user, and to make the user aware of their activities. However, an integrated view created through semantic technologies allows one to provide access and visualisation mechanisms that can help the user navigating and making sense of these data. The interesting aspects here is that such data might include many different dimensions. Through the integrated representation of the user's Web activity data built using mechanisms such as the ones described in the previous section, we can obtain representations that adequately represent all these dimensions, and that interlink them.

For example, using the data collected through the local Web proxy in [3], a number of simple displays can be shown to the user that summarise their own activities, from the time of the day when they are most active to their common search keywords or the geographical location of the websites they interact

with. However, one can also exploit the interlinked nature of these dimensions to investigate in more details some activities that might appear suspicious. For example, one might notice from a map-based visualisation that they interacted with websites located in Nigeria, and query the system to obtain more information about the times of these interactions, or the type of information exchanged with these websites.

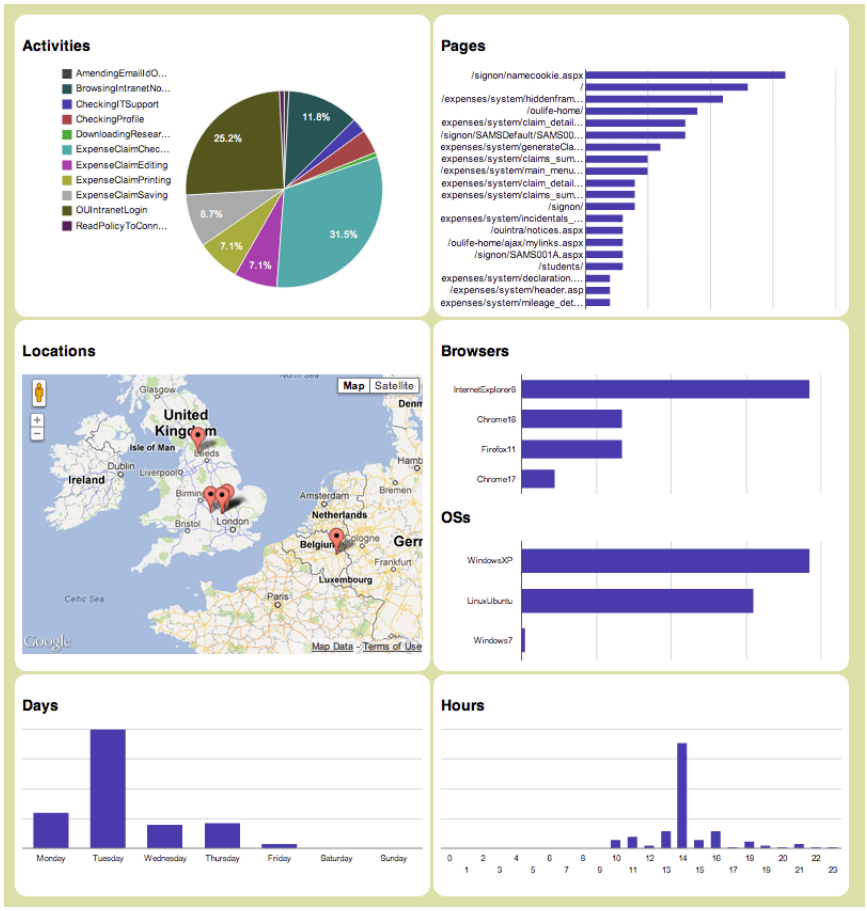

Fig. 2. Personal analytics dashboard.

To show the benefits and possible use of such approaches, which we call "personal analytics", we described in [7] a study where we developed a personal analytics dashboard using information from the logs of various systems in an organisation (The Open University), as described above and in [5]. The idea of this study was to see what would be the uses and concerns members of this organisa-

tion might have when given meaningful access to a record of their own activities. The dashboard, a view of which is visible in Figure 2, was built from querying a triple store containing the integrated information about the activities of the user on all considered systems, after this information had been processed to extract interesting dimensions (time, location, type of activity). The interface was also built to allow both usages of obtaining overviews of the activity (through charts showing information about the overall user activity) and of querying the data for specific information (by enabling the users to filter the displayed information on all dimensions, based on values in one of them–e.g., clicking on the time chart for "Sunday" to see only activities carried out on a Sunday). The study was then conducted by interviewing participants at the time they were being confronted to their own online activities through the dashboard. While the goal was not to focus on the responses specifically related to privacy, we could certainly identify different views, matching the common fundamentalist/pragmatist/unconcerned divide [12]. More interestingly however, this study allowed us to identify possible uses one might make of having a view over their own activity data:

Self reflection: Some participants of the study saw value in simply being able to reflect on what they do, on their own usage of resources and especially on how it reflected the way they worked.

Improving the use of resources: Very much related to the previous use case, the most commonly mentioned potential purpose of having access to one's own activity data is to improve the use of online resources, and to make it more efficient. This includes cases where the analysis of past activities would lead to a change in behaviour with respect to interactions with online resources.

Tracing anomalies: In a way somehow more specific than the one above, a commonly mentioned use of activity data is to trace back and find evidence or information related to some kind of anomalies in the users activities. Participants mention scenarios such as being able to check that an activity was properly carried out, or to analyse a situation that might have led to a privacy related issue.

Ensuring transparency: This use case is somehow different from the three others, as it does not really relate directly to the use of online resources or even activity data, but generally to the relationship with the organisation collecting these data. Indeed, while several participants did not see the collection of activity data as being in anyway worrying, it is often mentioned that it is simply "good to know what they know".

While, out of the four usages mentioned here, not all explicitly relate to privacy, they show that there is value in making activity data visible and understandable to users, and that even simple data integration and analytics techniques can help achieving that. It appears clearly that, while only a preliminary attempt, such personal analytics tools have the potential to make the Web a more translucent system, increasing users' awareness of their own activities and their ability to shape their use of online systems appropriately.

5 Reasoning to infer hidden privacy connections

In the previous sections, we showed how semantic technologies can be useful in creating the features of translucent systems for the Web, through data integration, manipulation and analytics. However, having access to data does not necessarily imply being able to make sense of it. Here, we argue that an even more critical aspect to a "Web privacy mirror" is that it should provide a way for the user to compare what can be understood from their activities as observed by the system to what they intended them to be. This requires the ability to interpret information integrated in the way described above at a higher conceptual level than what can be done with just the raw data.

This point was argued in [4] and described previously in this paper, where data obtained from the local web proxy was used to confront the user so as to measure their trust of online organisations and the sensitivity they attribute to their own personal data. Information was first computed on the basis of the raw data related to which online systems received which piece of personal information (e.g., `facebook.com` received the user's phone number). Based on this information, two measures were calculated:

- The observed trust the user gave to an online system, based on the idea that for a system to receive critical information implied that they were trusted.
- The observed criticality of a piece of personal data, based on the idea that if a piece of data is shared with many untrusted online systems, it must have low criticality.

These measures are dependent on each other and need to be calculated recursively.

Calculating such measures had the objective of confronting the user, who might have their own idea of their behaviour, with information about what can be inferred from their actual, observed behaviour. This was achieved through an interface showing the ranking of online systems based on trust (right of Figure 3) and of pieces of personal information based on criticality (left of Figure 3). The user then had the ability to move each element around, and dynamically see the effect on the calculation of the two measures on other elements. For example, if the user was to manually increase the criticality of a given piece of information, they would see the websites that have received this piece of information move up in the trust ranking.

Besides the complex technological implementation to achieve this result, the interesting aspect of this system is that it represents an implementation of the five characteristics of privacy mirrors applied to the environment of the open Web: it provides *history* of the interactions between the user and online systems. It uses this history to give *feedback* to the user regarding what can be inferred or concluded from observing these interactions leading to a greater *awareness* of the consequences of the user's own behaviour. This in turn increases *accountability* by making it possible to trace the origin of certain inferences and it also effects *change* by providing information on the user's behaviour in relation to the perceived or calculated privacy state.

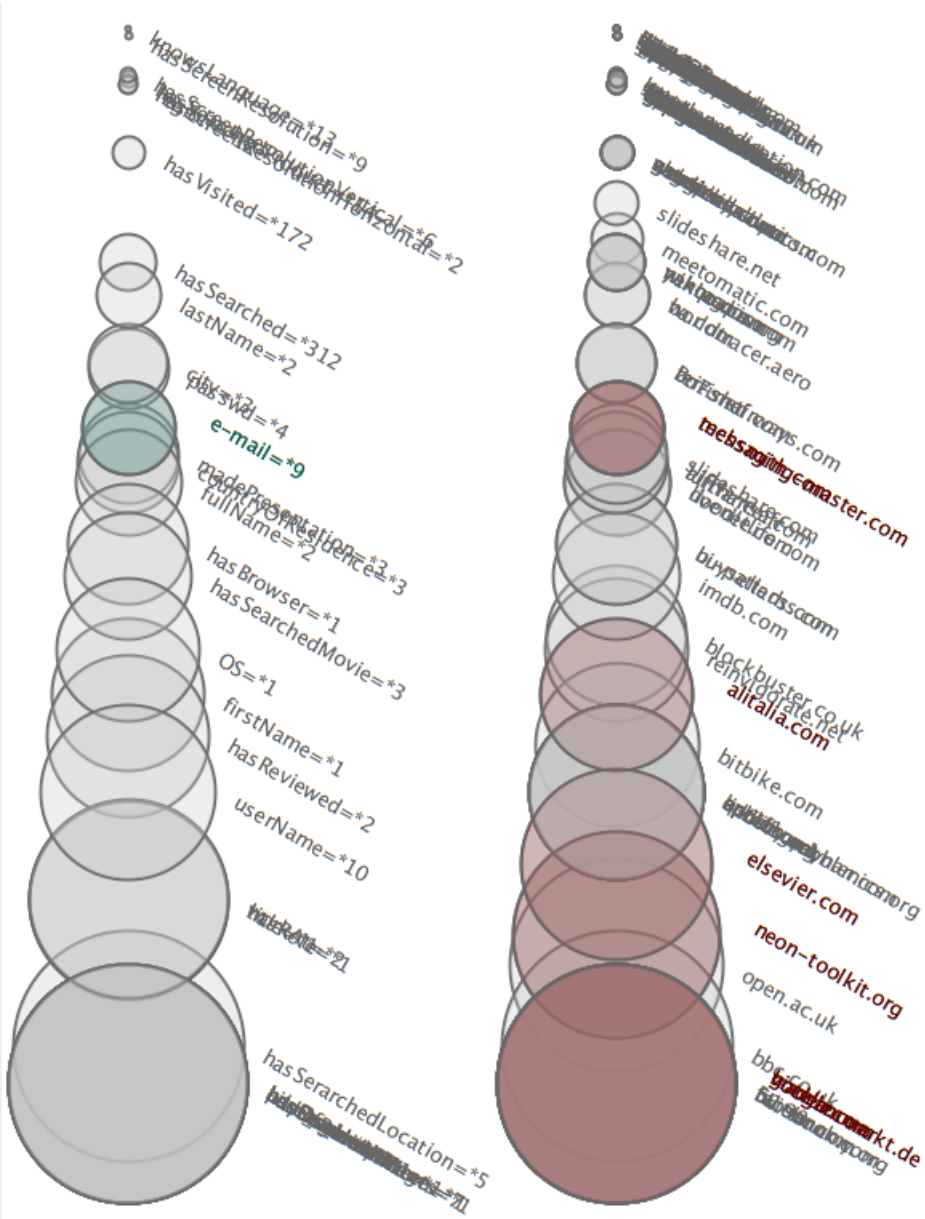

Fig. 3. Observed measures of trust in websites (right) and criticality of pieces of personal information (left) for a given Web user.

6 Related work

As [12] has shown, users can be classified according to three main categories (unconcerned, pragmatics, and fundamentalist) with respect to privacy. However, such studies only consider what the users declare to be their behaviour online. In [13], the authors show that, when looking at the actual behaviour of users in concrete scenarios, there is little being done even by the most privacy and technology aware of users i.e., the fundamentalists, which effectively supports the protection of personal information online. Our hypothesis here, supported by the notions of social translucence and privacy mirrors, is that this is partly due to a lack of visibility and awareness of Web users of their own activities, which is in turn partly due to the lack of appropriate technological support for such awareness and visibility on the Web.

Indeed, the inherent complexity and fragmentation of the flow of personal information on the Web makes it difficult for an individual Web user to monitor and make sense of their own information without appropriate technological support. In contrast with such a complexity, some of tools currently available [15, 14] only provide coarse-grained control options, are not able to support the process of assessing the implications of the potential aggregation of information published in a distributed fashion, and do not consider the complexity and variety of channels through which online activities happen. Typically, users would simply use popular Web search engines to check websites where their name appears [16]. Services such as Garlik[2] have emerged recently, providing more advanced approaches to monitor the Web for personal information. Other commercial services such as Trackur[3] and Visible Technologies[4] are able to monitor social media sites for references to a person.

Many discussions have taken place recently regarding the addition of a 'do-not-track' option in Web browsers. However, to understand the extent and potential impact of tracking mechanisms in relation to their own online behavior requires appropriate information for Web users. Tools such as Gosthery[5] have appeared, which display to users the tracking mechanisms used on visited websites, but without the ability to monitor the connections to such trackers across time, or to link to relevant actions (e.g., the use of the do-not-track option, or blocking the connection to the tracking service). Similarly, a number of basic tools have emerged recently that provide Web users with personal analytics features, based for example on their Facebook activities[6] or on their browser's history (using dedicated browser extensions).

---

[2] http://garlik.com
[3] http://www.trackur.com
[4] http://www.visibletechnologies.com
[5] http://www.ghostery.com/
[6] see the personal analytics feature of Wolfram Alpha - http://www.wolframalpha.com/facebook/

7 Conclusion

In this paper, we discussed the notions of social translucence and of privacy mirrors, and how achieving such approaches to privacy in an open Web context requires technological advances that are typically found nowadays in Semantic Web systems. We discussed these aspects by illustrating several systems we developed in the past and that implemented, at least partially, this idea of semantic translucent Web interaction. Indeed, achieving social translucence on the Web, especially to implement privacy mirrors, creates interesting technological challenges that fit the Semantic Web agenda. The ability to integrate, analyse, manipulate and reason over meaningful, processable data coming from distributed heterogeneous systems is here required, and to make privacy management feasible on the Web, will need to become part of the normal activities of everyday Web users. While we focused here in some specific examples, we believe this to be a general statement that needs to be considered in many different contexts. For example, we are currently studying the privacy settings of online systems such as Facebook, not from a purely technical point of view, but from the perspective that they currently obscure, more than they make visible, the consequences of using such systems on theirs users' privacy state. We believe that, through the appropriate semantic modeling of these mechanisms and through giving users the possibility to employ such semantic model to build their own view of their own activity, we can increase their ability to control the dissemination of their personal information, and therefore their privacy.

However, the claim that Semantic Web technologies are needed does not imply that they are sufficient. It appears clearly that privacy implications cannot be purely represented through straightforward data modeling and be considered only within an ontological reasoning framework. While some aspects do fit well the ontological perspective promoted by the Semantic Web (e.g., in defining basic implications on the inclusion of users in certain categories in Facebook), what is required here are higher level inferences regarding the information these users have access to, which could be achieved for example by integrating notions of epistemic logic into the considered reasoning framework (see [8] for a description of some preliminary work in this direction). Of course, another critical challenge is that we are suggesting here to provide everyday web users with the ability to employ complex and sophisticated Web, data management, semantic analysis and reasoning technologies on their own data, collected from their own Web activities, and for their own purpose. Critical work is required on the usability and human-interaction aspects of Semantic Web technologies in order to make such a vision feasible.

References

1. Anne Adams. Users' perception of privacy in multimedia communication. In *CHI '99 extended abstracts on Human factors in computing systems*, pages 53–54, Pittsburgh, Pennsylvania, 1999. ACM.
2. I. Altman. Privacy: A conceptual analysis. *Environment and Behavior*, 8(1), 1976.

3. M. d'Aquin, S. Elahi, and E. Motta. Personal monitoring of web information exchange: Towards web lifelogging. *Web Science*, 2010.
4. M. d'Aquin, S. Elahi, and E. Motta. Semantic monitoring of personal web activity to support the management of trust and privacy. In *Trust and Privacy on the Social and Semantic Web (SPOT) workshop at ESWC*, 2010.
5. M. d'Aquin, S. Elahi, and E. Motta. Semantic technologies to support the user-centric analysis of activity data. In *Social Data on the Web (SDoW) workshop at ISWC*, 2011.
6. M. d'Aquin, M. Rowe, and E. Motta. Self-tracking on the web: Why and how. In *W3C workshop on Web Tracking and User Privacy*, 2011.
7. M. d'Aquin and K. Thomas. Consumer activity data: Usages and challenges. Knowledge Media Institute, Tech. Report kmi-12-03, 2012.
8. M. d'Aquin and K. Thomas. Modeling and reasoning upon facebook privacy settings. In *Demo session at the International Semantic Web Conference, ISWC*, 2013.
9. S. Elahi, M. d'Aquin, and E. Motta. Who wants a piece of me? reconstructing a user profile from personal web activity logs. *Linking User Profiles and Applications in the Social Semantic Web (LUPAS) workshop at ESWC*, 2010.
10. T. Erickson and W. A. Kellogg. Social translucence: an approach to designing systems that support social processes. *ACM transactions on computer-human interaction (TOCHI)*, 7(1):59–83, 2000.
11. C. Fried. Privacy [a moral analysis]. *Yale Law Journal*, 77:47593, 1968.
12. Harris Interactive. The harris poll 17: most people are privacy pragmatists who, while concerned about privacy, will sometimes trade it off for other benefits. Online: http://www.harrisinteractive.com/vault/Harris-Interactive-Poll-Research-Most-People-Are-Privacy-Pragmatists-Who-While-Conc-2003-03.pdf, 2003.
13. C. Jensen, C. Potts, and C. Jensen. Privacy practices of internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies*, 63(1-2), 2005.
14. Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. A "nutrition label" for privacy. In *Proc. of 5th Symposium on Usable Privacy and Security*, pages 1–12, Mountain View, California, 2009. ACM.
15. KnowPrivacy. Policy Coding Methodology. http://knowprivacy.org/policies_metho-dology.html. [Accessed: 6th March 2013].
16. M. Madden and A. Smith. Reputation management and social media. Report from the PewResearchCenter, 2010.
17. S. T. Margulis. On the status and contribution of westins and altmans theories of privacy. *Journal of Social Issues*, 59(2), 2003.
18. E. D. Nguyen, D. H. ; Mynatt. Understanding and shaping socio-technical ubiquitous computing systems. GVU Technical Report;GIT-GVU-02-16, 2002.
19. H. Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2010.
20. D. J. Solove. *Understanding Privacy*. Harvard University Press, 2010.
21. K. Thomas and M. d'Aquin. On the privacy implications of releasing consumer-activity data. Knowledge Media Institute, Tech. Report kmi-13-02, 2013.