

# Informed Consent on the Semantic Web - Issues for Interaction and Interface Designers

Paul Shabajee<sup>1</sup>

<sup>1</sup> Institute for Learning and Research Technology, University of Bristol, UK  
and HP Labs, Bristol, e-mail: paul.shabajee@bristol.ac.uk

**Abstract.** This discussion paper highlights differences between semantic web-based and more traditional web-based applications that raise significant privacy and other ethical issues for users when they provide personal data (about themselves or others) into a semantic web environment. The paper introduces the technologies, highlights some key differences and reviews the concepts behind informed consent. It then argues that a critical challenge for interaction and interface designers of semantic web (and related) applications is to ensure that users are enabled to give their fully informed consent before they provide such data. Finally it explores specific aspects of the challenges and potential approaches to meeting them.

**Keywords:** semantic web, interaction design, user interface design, informed consent, legal issues, ethics

## 1. Background and Introduction

This paper has been motivated by the fact that the issues related to privacy in the context of semantic web-based (and similar) applications, have been a recurring theme in many pieces of work that the author has been involved in over the last five years, including work on projects as part of the Semantic Web Advanced Development Europe<sup>1</sup> and follow on projects, and more recently in the Iugo<sup>2</sup> a research project. Iugo is investigating the issues surrounding the design and development of a system to integrate information (from across the web) related to Research and Development (R&D) events, including information about the events, speakers, delegates, presentations, projects, blogs, electronic chat logs, etc.

Put simply, the issues seem to arise because publishing data in semantic web and similar environments<sup>3</sup> is qualitatively and quantitatively different from more traditional modes of web publication. These differences (see ‘What’s So Different With the Semantic Web?’ below) compound what are already problematic issues

---

<sup>1</sup> <http://www.w3.org/2001/sw/Europe/>

<sup>2</sup> [http://www.jisc.ac.uk/index.cfm?name=vre\\_iugo](http://www.jisc.ac.uk/index.cfm?name=vre_iugo)

<sup>3</sup> See - Welcome to the Semantic Web and ‘Web 2.0’, below

related to privacy in web-based environments – in particular the associated, potentially negative consequences for users.

This paper aims to build on existing work in the combined areas of privacy issues and HCI (Human Computer Interaction) by exploring the additional issues raised by semantic web type technologies. Consequently, it aims to be part of the process of stimulating debate and discussion of these issues, hopefully leading to further concreted research and development activity. In order for such debate to be fruitful the author believes that members of different communities need to come together, and so the intended audience is drawn from HCI and semantic web communities and also more widely, those with interests in issues of *informed consent* and security, privacy and education. Clearly not all members of those communities will be familiar with all the issues or technologies.

To these ends the paper is structured as follows, 1) it briefly characterizes semantic web and similar environments, 2) explores the existing issues of privacy and HCI on the web, and the additional issues raised by functionalities of the semantic web, 3) highlights the apparently limited awareness of these issues amongst developers and researchers as well as users, and 4) introduces the concepts surrounding *informed consent* and specifically how they may be part of the potential solution to these issues. Finally it explores potential issues for debate and potential for future research and development activities.

## 2. Welcome to the Semantic Web and ‘Web 2.0’

The semantic web can be characterized as a machine readable<sup>4</sup> ‘web of data’ [1]. The central principles on which it is based are very simple and well detailed elsewhere<sup>5</sup>, however as many reading this paper may be unfamiliar with these principles, a brief description will be given here.

In broad terms the goal of the semantic web technologies is to enable the automatic integration of information from across the web and web-like environments (e.g. corporate intranets) and even personal information environments (e.g. linking information between applications on a single desktop computer). The underlying goal is the ability to provide significant advantages to human users, (e.g. greater and more effective and efficient access to information, the development of more customized services, etc...). Some fundamental aspects that will be useful to understand in the context of this paper are:

- 1) Effective integration of information requires that things (real world and digital objects, and abstract concepts e.g. people, publications, academic subjects) and the relationships between them (e.g. authorship, that a paper is

---

<sup>4</sup> ‘Machine readable’ in this context means that a computer program can automatically access and process the data in way pre-defined by a computer programmer.

<sup>5</sup> A good starting point is the W3C Semantic Web page - <http://www.w3.org/2001/sw/>

about a particular subject) can be uniquely identified. The semantic web provides simple mechanisms to do this. It is important to realize that a single ‘thing’ may have many unique identifiers (e.g. a person – passport number, work ID number, etc.).

- 3) In order for machines to be able to automatically process/integrate some kinds of information (e.g. different aspects of people) it is necessary to have a way to represent a simplified (machine readable) model of key aspects of those things (e.g. that both presenters and delegates at a conference are both types of people or that a particular person may have multiple e-mail addresses, etc. These models are called *ontologies* [1].
- 4) For all this to work on a web-scale you need to define clear technical standards for representing (in machine readable format) the various types of information and ontologies that will be used across the web. These are the basic languages: RDF (Resource Description Framework), RDFS (RDF Schema), and OWL (Web Ontology Language – correct acronym) [1].
- 5) In order to really make this work you need other components, such as ways of representing ‘rules’ (e.g. conditional statements like “if x has value y then w has value z”) that cannot be written down as a statement; ways of robustly tracking the provenance and authenticity of particular statements, ontologies, etc.; and more broadly standard mechanisms for assessing and asserting trust. As yet these and a few other required elements have not been standardized<sup>6</sup> (see also discussion later on the Policy Aware Web).
- 6) The semantic web is still evolving. Active standards and working groups include, e.g. semantic web services, rules interchange and RDF Data Access.

In summary the *semantic web* is simply data linked with tightly defined metadata (data about data) about how one piece of data is related to others, combined with an ontology(ies) to help relate the data more richly. These explicit (machine readable) statements of relationships between identifiable things enables computer programs to use simple mechanisms to use those relationships to enhance functionality of information systems – that gets you a long way, but other requirements exist (such as representing ‘trust’) to complete the vision, and these are under development.

## 2.1 Web 2.0

Web 2.0 [2] is a phrase that, while it has no firm definition is widely used to describe a class of web-based applications that share some common goals with the semantic web e.g. making data more sharable/interoperable, joining data from multiple sources. To that extent they are ‘semantic web like’. Broadly common characteristics of ‘Web 2.0’ applications include [3]:

- User generated and/or user influenced content
- Applications that use the web (versus the desktop) as a platform, in innovative ways
- Similar visual design and shared functional languages

---

<sup>6</sup> Rule Interchange Format Working Group Homepage - <http://www.w3.org/2005/rules/wg>

- Leveraging of popular trends, including blogging, social tagging, wikis, and peer-to-peer sharing
- Inclusion of emerging web technologies like RSS, AJAX, APIs (and accompanying mashups), Ruby on Rails and others
- Open source or sharable/editable frameworks in the form of user-oriented "create your own" APIs

So called 'mashups'<sup>7</sup> are of particular relevance here because they provide examples of information and service integration that mirrors some of the goals of semantic web integration, but generally using more diverse approaches. They are generally web-based applications that draw content from multiple sources to create a new site or service. They generally use APIs (application programming interface) provided by the original data source sites – although some use other methods to obtain data, such as scraping data from plain HTML web sites. Widely cited examples of characteristic Web 2.0 sites include Flickr, Google Maps, del.icio.us, last.fm and sites that draw content from multiple sites such as Technorati and Yuan.CC Maps.

### 3. Privacy, the Web and Semantic Web

#### 3.1 Personal Information Integration

*A friend's daughter mentioned that she had a new boyfriend. She mentioned his name and other passing details. My friend (after debating with himself about the ethics involved) used Google to see if he could find out a little about the boyfriend. He could – in a very little time he had found the boyfriend's professional homepage which linked to his (personal) flickr photo sharing pages. Other Google results led to other incidental links that collectively painted a fairly clear (if incomplete and potentially inaccurate) picture (literally!) of the man. Interestingly it seems the daughter had not (significantly) 'Googled' her boyfriend and so the father knew some information about her boyfriend that she did not. It was also clear to my friend that his daughter might rather he had not found the flickr pages.*

This kind of personal anecdote<sup>8</sup> is common and illustrates the fact that on the web 'personal' information that was previously simply unfindable without very significant effort (professional qualifications, photographs, statements of opinion, etc.) can be virtually effortlessly findable using existing web search engines. It also highlights specifically that personal information that *we publish about ourselves* and *over many years* are often the raw material for such searches – these might include personal web-pages, mail-list contributions, photos, blog entries, logs of IRC (Internet Relay Chat) activities, feedback to others' blogs or websites holding information about events we

<sup>7</sup> [http://en.wikipedia.org/wiki/Mashup\\_%28web\\_application\\_hybrid%29](http://en.wikipedia.org/wiki/Mashup_%28web_application_hybrid%29)

<sup>8</sup> In all cases where anecdotes are presented slight changes have been made where necessary to protect the identities of these individuals.

have attended, etc. In researching this paper, the author came across numerous personal anecdotes of unintended exposure of personal data by individuals.

Issues not only exist for the individuals who are *subjects* of searches but also individuals making the searches – in the previous example the father had a ethical choice to make as to whether to conduct the searches. Others’ such as prospective employers may also have legal issues or organization guidelines to consider with respect to equal opportunities legislation, etc.

Perhaps less obviously, incidental information about us and our activities may be captured (often unknown to ourselves) in ways that may immediately or at a later date become public on the Web. An interesting illustration comes from Merkitys<sup>9</sup> a piece of software run on mobile camera phones that captures incidental information as the picture an allows users to post it to Flickr (see “Web 2.0 above) along with tags representing the additional captured data this includes location information via GPS (Global Positioning System) device, GSM Cell info (i.e. mobile phone mast identifier) and the Bluetooth (unique) identification numbers of other devices *near* the camera at the time. These pieces of data are automatically used to tag the images that are then uploaded to Flickr. If you were incidentally near that device when it took a photo and someone knows your devices’ Bluetooth IDs, they can then search Flickr for that tag and see photos that (at least assuming no fraud) were taken when your device(s) (and so inferring you) were present.

Less publicly (data held internally to organizations), data derived from what Johnson [4] calls the increasing ‘*instrumentation of human action*’ can be combined from different sources and used to draw up very comprehensive pictures of individuals and their activities e.g. shopping transactions (via store loyalty cards), car registration number recognition systems, internet browsing and searching histories via your ISP(s), employers records, and even government-held records (as services become more ‘joined up’) – for example, in Scotland the Government is developing a system that will “... *provide the Citizens of the UK with a single convenient interface with multiple Government bodies.*”<sup>10</sup>. High profile examples such as AOL’s recent inadvertent publication of search log data demonstrates that even the ‘secure’ environments pose risks<sup>11</sup>.

These are very well known and studied issues from legal, technical and ethical perspectives (see for example [4] and [5]) and there are a number of initiatives that are currently exploring the issues and seeking ways to mitigate the web-based and other examples. The most comprehensive example is probably the Policy Aware Web [5], [6], which aims to provide a robust mechanism for a machine readable manner of describing, recording and monitoring how information can be, and is, accessed, used and re-used in the context of legal and other policy frameworks. Others include the

---

<sup>9</sup> <http://meaning.3xi.org/>

<sup>10</sup> <http://www2006.org/speakers/kinney/>

<sup>11</sup> <http://news.bbc.co.uk/1/hi/technology/5255732.stm>

Platform for Privacy Preferences (P3P) Project<sup>12</sup> and, in the context of copyright, the Creative Commons<sup>13</sup> initiative. Although it should be born in mind that technical solutions alone cannot offer a comprehensive panacea, not least because standards must be taken up and used before they can have an impact – the standards themselves must be *usable* by developers and designers for take up to be effective.

In these contexts, there are numerous existing legal frameworks (e.g. copyright, moral and database rights, data protection/privacy, etc.) and additionally various ethical frameworks related to the use and re-use of kinds of information (e.g. for IT professionals the ACM code of ethics [10] contains many elements that refer directly to privacy and associated issues of dignity of users e.g. “3.5 *Articulate and support policies that protect the dignity of users and others affected by a computing system: ...*”). For an overview of many of the key (pre-semantic web) issues, see [4].

However regardless of legal or ethical frameworks, in a web-based and highly ‘instrumented’ (see above) environment there are numerous places where individuals expose their own information inadvertently. Two more anecdotes collected during the writing of this paper further illustrate the point.

1. *On buying a new digital camera, a very technically competent individual entered the owner data into the camera (name and phone number) in the hope that if it were lost it could be returned. This person also had a Flickr account and uploaded images taken with the new camera. Being a supportive member of the community they used the Flickr setting to allow others to see the exif data (metadata embedded in the photographs). What they did not realize was that the owner details were exposed as part of that information. They had therefore unwittingly exposed that data and also exposed the link between their identity and their Flickr id, which they may not have wished to do.*
2. *Another individual purchased a new Internet domain name for their personal blog. While they used a pseudonym in the Blog they did not mind others knowing who they were. They blogged the fact that they had a brand new computer system and expressed opinions about the features and associated gadgets. What they had forgotten was that when registering the domain name they had given their home address. And so the publicly available record contained their personal details. This made them feel very uncomfortable, having said to anyone who cared to investigate, that £1000s of new equipment was at that address.*

### **3.2 What’s So Different About The Semantic Web?**

If all this is possible now and the issues are already significant, what difference does it make if you add in semantic web type environments? Johnson [4 p117] lists five ways in which computers and IT (pre-semantic web) has changed what was

---

<sup>12</sup> <http://www.w3.org/P3P/>

<sup>13</sup> <http://creativecommons.org/>

traditionally called ‘record keeping’ in a privacy context:

1. it has made a new scale of information gathering possible
2. it has made new kinds of information possible (e.g. transaction generated data)
3. it has made a new scale of information distribution and exchange possible
4. the effect of erroneous information can be amplified
5. information about ones life may endure much longer than ever before

This gives us a useful starting point from which to look at the additional functionality that a semantic web based infrastructure adds to this:

1. *Universal*<sup>14</sup> *open standards* for representing data, metadata and model (ontology) representation. All data will be [or queried as though it is. See Ease of Integration below] in semantic web formats RDF/RDFS/OWL. Traditionally a fundamental barrier to integration was the vast number of (generally bespoke) formats.
2. *Identifiably* – fundamental to the semantic web is the use of a *standard* framework for the representation of unique identifiers. This means that all data held in semantic web environments will use a common (or at least interoperable) form(s) or representation for identifiers. Issues arise because increasingly single identifiers are used across multiple systems and because in semantic web applications distinct identifiers can be linked together very simply. In previous environments a key barrier to integration was the massively diverse range of identifiers and their formats and difficulty in mapping between them.
3. *Ease of Integration* – the semantic web was designed from the ground up to make data and metadata integration as easy as possible<sup>15</sup> in a real world, web scale environment. Integrating data is very much simpler (in principle) than existing relational database or other (e.g. XML based) systems. This ease of integration is at present taking a significant step forward with the development of ways of making existing relational databases (via so called SPARQL<sup>16</sup> end points) and other data sources e.g. desktop applications (e-mail, calendars, etc. – e.g. Aperture Framework<sup>17</sup>) and corporate information systems such as LDAP (Lightweight Directory Access Protocol) data, available to semantic web applications.
4. *An integrated means of simple inference over all data* – it has always been possible to make logical inferences over (appropriate, i.e. logically consistent) digital data. However in the past, because data in different systems used different representation standards, inferences over information in multiple data sources was highly problematic. In the case of semantic web data the use of standards and the mathematical basis of the languages make such inferences (assuming necessary consistency in the data and models) relatively trivial.
5. *Persistence of data* – the inherent reusability (share-ability) of semantic web data and the potentials for the creation of new data generated from inference, mean that data is more likely to become distributed, and so persist (even when/if the original

---

<sup>14</sup> Universal as in, if data is to be available to as semantic web data it must use the appropriate standards – as early ‘web’ data by definition used HTML.

<sup>15</sup> Given constraints of time available to produce the standards.

<sup>16</sup> SPARQL - query language for the semantic web (<http://www.w3.org/TR/rdf-sparql-query/>)

<sup>17</sup> <http://aperture.sourceforge.net/>

source has been removed) in a semantic web environment than on the current Web; that is, unless there are very systematic controls over reuse – see Policy Aware Web above.

The effects of these abilities and features can be immediately significant. For example, clearly anyone can make a statement on the semantic web, just as they can in other publishing environments (news papers, etc.). However there is a *very significant* difference. The semantic web version of a statement can be *immediately* integrated and can be represented in *results from queries from across the whole web*. For example, I might have multiple online personas to allow me to separate my personal from my professional activities. Now what if someone finds out about the multiple personas and publishes the fact in a semantic web environment – easily done at present as part of a FOAF<sup>18</sup> (Friend Of A Friend) file? Any search made about one persona would then include results related to the other. Further, just as in other types of publication, if the information is true, there may be no legal recourse to have the statement removed – especially given differences in legislation between jurisdictions. The traditional ‘anonymity’ of the Web becomes insecure. This example also highlights the need for trust and provenance tracking mechanisms.

All of the examples above demonstrate that even without semantic web functionality, the current web environment is full of applications and situations where information about individuals, spread out across the Web, can be manually integrated by people with effective IT based information skills. *People are necessary* because the information sources are generally ‘siloes’, i.e. self contained and unlinked to related sources. In a semantic web-based environment the problems are amplified because it is capable of simply short circuiting these silos, allowing the dynamic integration (e.g. you can ask a query across all) of data across all the silos.

Importantly, without an infrastructure like that beginning to be explored by the Policy Aware Web initiative [5] and others investigating *access control* and *trust*, the integration of such data seems largely outside the scope of exiting data-protection legislation. The data is held in different systems (let us assume each with its own valid data protection policies), however the end user can integrate the data (about an individual) via a query simply through their web browser-based applications.

It seems unclear what the status of this kind of situation has in legal terms, because at no point anywhere other than browser window is the information so integrated, stored or cached. In this hypothetical situation, there has seemingly been no breach of data protection laws prior to the query being made and yet, in the browser, the end user sees information that (if held by a single organization) would clearly be in breach of such legislation.

---

<sup>18</sup> <http://www.foaf-project.org/>

## 4 Informed Consent?

The fundamental question of informed consent can be naively summed up simply as, ‘Do you *fully* understand what you are about to do or commit to of your own *free will*?’ However, the individual terms needs some unpacking, as we shall see below.

Informed consent is a centrally important aspect of medical practice and research, as well as social science and other forms of research [7]. It is fundamentally associated with issues of basic rights – for example to have control of what happens to ones body, or to not take part in something like a research study if one does not agree with its purposes – and where actions involve some risk of harm (physical, psychological, financial, professional, etc.).

In previous sections the author has attempted to demonstrate that entering data into a semantic web-based system has potential entailments. Those entailments might include serious negative consequences to the person in exposing themselves, or those they know, to personal risks of various kinds, ranging from mild embarrassment to serious damage to their professional reputation or inadvertently exposing themselves to risk of burglary or stalking. Consequences may also be more broad, exposing information that may be joined to other data about them or other people/things, in ways they had not intended, possibly a long time in the future. Additionally, information may be re-used in ways that they did not intend or for purposes that they may actively disagree with (e.g. for moral, religious, political or commercial reasons).

Informed consent is thus a principle that seems appropriate here. In fact as van de Geest et al [7] point out it is already the case that the EU Data Protection Directives of (1995 and 2002) requires that a data subject gives informed consent to their information being processed, with “consent” meaning, “... *any freely given specific and informed indication of his wishes by which the data subject signifies his agreement to personal data relating to him being processed.*”<sup>19</sup>. There are many other national equivalents around the world. However as they indicate, “*The 2002 directive even suggests that consent can be dealt with sufficiently by ticking off a checkbox on a web site.*” They argue that such an approach is inappropriate and insufficient in the context of users providing information for personal profiles in online environments.

Friedman, et al [8], as part of the Value Sensitive Design Research Lab<sup>20</sup>, explore the issues of informed consent and interface design directly in their Informed Consent Online project<sup>21</sup>. More broadly the concept behind Value Centered Design is “... *an approach to the design of technology that accounts for human values in a principled and comprehensive manner throughout the design process.*” They present a model of informed consent for information systems, composed of six components. The first two related to users being ‘informed’ and the other four to ‘consent’.

---

<sup>19</sup> [http://www.cdt.org/privacy/eudirective/EU\\_Directive\\_.html](http://www.cdt.org/privacy/eudirective/EU_Directive_.html)

<sup>20</sup> <http://www.ischool.washington.edu/vsd/>

<sup>21</sup> <http://www.ischool.washington.edu/vsd/projects/informedconsent.html>

1. **Disclosure** – refers to providing accurate information about the benefits and harms that might reasonably be expected from the action under consideration. Questions might include, what information is captured?, who has access to it?, what will it be used for? Van de Geest et al [7] take this further and include the need to disclose information about available alternatives beyond simply not providing the information or taking part – common and necessary elements in a medical context.
2. **Comprehension** – refers to the individual’s accurate interpretation of what is being disclosed. Here we must decide what constitutes adequate comprehension.
3. **Voluntariness** – that is individuals could reasonably resist participation should they wish to. This includes not only direct coercion but also restriction of substantive choice or manipulation of information, such as providing misleading information or conducting psychological manipulation (e.g. “everyone else is ...”).
4. **Competence** – refers to possessing the mental, emotional and physical capabilities needed to give informed consent – e.g. a person with Alzheimers’ or a child may not be capable of making the required reasoned judgments to make the decision.
5. **Agreement** – refers to the user having a reasonably clear opportunity to accept or decline to participate. This component may appear to be the most straightforward, however this is deceptive. For example, users may forget (many months later) that their conversations in a chat room are being recorded. Also, as in other contexts, many types of agreement are implicit – for example, walking onto a football field, an individual has implicitly agreed to participate in the normal activities of the game. Deciding where the boundary between implicit and explicit agreement exists in an electronic environment can be problematic.

Additionally, in medical and other research studies, participants have the right to withdraw at any point from a treatment or study without needing to give any reason. This in particular is problematic in the current stage of semantic web development, because in the present environment much semantic web data is harvested and re-used on multiple sites with no indication of how it should or should not be used – RSS<sup>22</sup> is a prime example of such a format. In such cases *withdrawing* data from publication is simply impossible, especially if that information was used to make inferences and the new inferred information re-published.

6. **Minimal Distraction** – this criterion arose from empirical investigations following their earlier work. It refers to meeting the preceding criteria without unduly diverting the individual from the task at hand. This was added to the previous 5 criteria because users “*became numbed to the informed consent process and disengaged from the process in its entirety.*” They acknowledge that this is problematic but crucial if informed consent is to be realized in practice.

---

<sup>22</sup> [http://en.wikipedia.org/wiki/RSS\\_\(file\\_format\)](http://en.wikipedia.org/wiki/RSS_(file_format))

As Friedman et al [8] state, in *face to face* interactions a dialog takes place with physical cues helping to ensure adequate understanding. More broadly, as Weitzner [9] points out, in real world environments, “*we rely on various feedback loops to establish what is acceptable versus unacceptable behavior.*” This is the point that makes this process fundamentally related to interaction and interface design in semantic web based environments.

The current practice of gaining informed consent in a web environment is somewhat different from what might be seen as ideal in the context of this model. In most cases websites provide a ‘privacy statement’ of some form, during the sign up process, making these issues clear. It may also provide a ‘terms of use’ statement that details how a user of the site should behave, their responsibilities, legal conditions and other issues related to the use of the service.

However, in general, acceptance of the terms of use and privacy statement is largely a matter of the user ticking a box or pressing a button that states that the user has ‘reviewed’ and/or ‘agrees’ to it/them. These *terms* are often hyper-linked from the sign up page or in some cases in a scrollable frame within the page. This was the case for 18 of 20 web-based services reviewed while writing this paper, ranging from photo-sharing, social book-marking, blogging, online video conferencing, social/business networking, personal information sharing ‘spaces’, online e-mail and personalized web portal applications. Where this was not the case (in two of the 20) they did not appear to provide a privacy statement or terms of use statement for review at the time of signup. If multiple services were provided by one company, the sign up systems were common (in form) to all of those services.

While this approach to providing ‘disclosure’ may meet the *letter of law* it seems very inadequate as a means of obtaining *genuinely* informed consent, and would certainly not be seen as sufficient in either medical or social science research contexts. Indeed, users tend not to read agreement policies during online registration processes in any case [8].

The difference between medical/formal research contexts and providing data to websites *may* be one of proportionality, i.e. the necessary effort required to ensure consent is adequately given can be based on a risk/benefit analysis. The potential risks of medical intervention may be high and the impact on individuals who take part in sensitive social science research could be devastating, or if trust in the professional academic research community (to keep information confidential) were degraded, the impact would be disastrous to that community. However in the case of a single website holding data used only on that website the risks *may* be seen as minimal.

However as argued above, the semantic web and similar technologies bring with them more significant implications and risks – especially while there is no mechanism such as that proposed by the Policy Aware Web initiative to mitigate the situation.

## 5 Why is This Relevant to Interaction and Interface Design?

The previous section demonstrated that informed consent is multifaceted. Ensuring that a person is meaningfully giving their *informed consent* must be a process, a *dynamic and conditional conversation*. When seeking consent, doctors and researchers (who follow appropriate guidelines) must ensure that all of the criteria are met. This must of its nature be a *two way process*. In an automated electronic environment these *conversations* must take place via interactions between users and computer systems. Given the limitations of human/computer interaction compared to human/human interaction, developing interactions and interfaces that ensure that users understand the implications of doing so, seems an urgent challenge.

It is urgent not least because it seems to be an essential component in the development and uptake of the technologies; as Weitzner et al [5] points out, “*people will not share information freely in an environment that is threatening or antithetical to basic social needs such as privacy, security, the free flow of information, and ability to exercise their intellectual property rights as they choose.*” Further, trust (in technologies and organizations) is often seen as a prerequisite for the adoption of online services of whatever form [7].

Looking to the future, if initiatives such as Policy Aware Web (or similar) are to work, the design of user interactions and interfaces will be critical if they are to meet their key goal of providing transparency, that is “*...both people and machines needs to be able to discover, interpret and form common understandings of the social rules under which any given resource seeks to operate. Can it be shared, copied, commented upon, made public, sold, etc?*”

An interface, as well as providing means to practical ends, also conveys to a user a set of *embedded values* that underlie design [4], [8]. At least part of the building of necessary trust in users is that the interfaces that they use and the interactions they have communicate clear values that reassure them that their concern for their best interests has been central to the design of the system.

Working to understand the design requirements of user interactions and interfaces that genuinely help ensure that users are giving their informed consent is thus a critical issue for the whole semantic web community, with an obvious and fundamentally important role to be played by those working in user interaction and user interface design. However a more multidisciplinary approach is required, with active collaboration between members of a range of communities – HCI, semantic web researchers and developers, privacy and security research communities, educationalists, lawyers, governments, commercial companies, and most importantly a full range of end users themselves.

## 6 Moving forward – Exploring Solutions

The privacy related issues that are arising with the semantic web (and indeed wider web) are co-evolving. Any set of solutions, be they technical, design, cultural, educational, legal, etc. will co-evolve in that wider context. This co-evolution seems to offer significant challenges because of the new functionality and affordances the technologies provides. Below are some areas where the author has identified aspects of the problem space that would benefit from further investigation and/or offer potential solutions to specific problematic areas; there are no doubt many others.

1. **Better understanding the issues** – while many issues are being well articulated by current researchers such as Friedman, van de Geest and Weitzner, their colleagues and many others, there is much more that is yet to be explored, not least because as yet there are few large scale semantic web applications in wide spread use. However there are many instances from the use of other existing technologies that highlight issues needing attention. Understanding these issues will require collaborative working between members of a range of communities – see previous section. The relationships between technical solutions and interfaces that allow the most effective implementations seem like an important area of investigation.
2. **Design Methodologies** – the development of new approaches and increasing awareness of existing approaches to the processes of interaction and interface design can take account of these issues – for example, the ‘Value Sensitive Design’ approach [8] which uses a tripartite methodology made up of a) conceptual investigations, b) empirical investigations, and c) technical investigations.
3. **Development and evaluation of specific approaches** to interaction and specific interface designs. These will include development of existing approaches such as ‘peripheral awareness’ based interfaces developed by Friedman et al [7] (section 24.3.5), and exploring ways to mitigate specific instances, such as that suggested by a user of Flickr (in relation to the exposing of personal exif metadata, in the anecdote above) to show users an example of the metadata including potential problems, at the time that the option is selected. Moving towards more generalisable understandings can assist in future design, e.g. in nearly all of the studies cited above, the issue of users having inaccurate or incomplete mental models of the workings of electronic environments is highlighted.
4. **Education of users** – many of the issues detailed above arise because of lack of knowledge and understanding on the part of users. Issues for educational processes are wide ranging – covering technical understandings, social rules and expectations, identification of potential issues or threats, etc. These may be through specific, long term educational processes (e.g. embedding the issues in the curricular at all educational levels, as part of ‘information literacy’) or via publicity or educational campaigns (as has been done in the case of threats such as virus protection and phishing by online service providers) or within the interfaces themselves, through offering contextual tours to users or by embedding cues in the interface to draw the user’s attention to key learning points.

5. **Education of designers and developers** – education needs to span not only users, but designers and developers. The author, while investigating these issues, has spoken to many developers and researchers involved in the design and running of web sites and research projects. He has found *numerous* examples, many of them high profile and successful projects, where those involved were simply unaware of issues of direct relevance to their work, such as obtaining or even considering legal rights and related issues (e.g. copyright, moral and database rights when collecting data from external websites, recording participants without explicit permission, or discussing how the recordings are to be used.) or ethical issues, such as collecting research data via API's of Web 2.0 type applications, or screen-scraping data from Web pages and yet not contacting the *research subjects* (individuals whom the data is about or who created/own the data) in any way to seek their permission. Another aspect of education in the context of designers and developers is awareness raising with regard to standards and technologies available to them as part of a larger tool kit of solutions.

## 7 Conclusion

The semantic web provides very significant potential benefits to end users and developers in a number of key ways; ease of information integration, more efficient searching, more effective re-use, etc. (see above). However, with that functionality there are new and possibly substantial risks for end users when they provide personal information into an open or semi-open semantic web environment – risks that for the most part users cannot be expected to be aware of.

While there are initiatives under active development to help mitigate many of these risks, at present such systems are not in place. Even in the longer term many *risk creating* behaviors on the part of users are accidental or unknowing. Indeed even when, and if, such mitigating technologies are in place, they must have effective models of user interaction and user interfaces to actually produce their desired effect.

Given that this is the case, there is an associated ethical (and in many cases legal) obligation on the part of service providers to try to ensure that users are aware of these risks so they can make truly informed judgments before providing personal information. Meeting these challenges will require multidisciplinary collaboration and exploration of a wide range of issues and potential solutions.

## References

1. W3C (2006) W3C Semantic Web Activity. Available online: <http://www.w3.org/2001/sw/>
2. O'Reilly, Tim.: What Is Web 2.0: Design Patterns and Business Models for the Next Generation of Software (2005). Available online: <http://www.oreilynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>

3. Seomoz.org.: What is Web 2.0? (2006). Available online: <http://web2.0awards.org/web20-zeitgeist.php#whatis>
4. Johnson, Deborah G.: Computer Ethics, 3rd ed. Prentice Hall. (2001)
5. Weitzner, Daniel J., and Hendler, Jim. and Berners-Lee, Tim. and Connolly, Dan. Creating a Policy-Aware Web: Discretionary, Rule-based Access for the World Wide Web (2005), in Ferrari, Elena and Thuraisingham, Bhavani M. (2005), Web And Information Security, IRM Press, London. Earlier version Available online: <http://www.mindswap.org/~hendler/2004/PAW-final.pdf>
6. Weitzner, Daniel J., Abelson, Harold., Berners-Lee, Tim., Hanson, Chris., Hendler, James., Kagal, Lalana., McGuinness, Deborah L., Sussman, Gerald Jay and Waterman, K Krasnow Transparent Accountable Data Mining: New Strategies for Privacy Protection, Computer Science and Artificial Intelligence Laboratory Technical Report, MIT-CSAIL-TR-2006-007, January 27, 2006. Available online: <http://www.w3.org/2006/01/tami-privacy-strategies-aaai.pdf>
7. Van der Geest, Thea., Pieterse, Willem Pieterse and de Vries Peter.: Informed Consent to Address Trust, Control, and Privacy Concerns in User Profiling, Workshop on Privacy-Enhanced Personalization, 10<sup>th</sup> International Conference on User Modeling, Edinburgh (2005). Available online: <http://www.isr.uci.edu/pep05/papers/InformedConsent.PDF>
8. Friedman, Batya., Lin, Peyina and Miller, Jessica K. (2005) Informed consent by design. In L. Cranor and S. Garfinkel (Eds.), Security and Usability (chapter 24). O'Reilly.
9. Weitzner, Daniel J.: Testimony of Daniel J. Weitzner before the United States Senate Committee on Commerce, Science, and Transportation (2000). Available online: <http://commerce.senate.gov/hearings/0525wei.pdf>
10. ACM.: ACM Code of Ethics and Professional Conduct (2005). Available online: <http://www.acm.org/constitution/code.html>