

The Usability Imperative Inherent in the Semantic Web

Duane Degler and Renee Lewis

A tremendous amount of hope – and hype – has been attached to Tim Berners-Lee’s concept of the Semantic Web, where machine-readable ‘meaning’ enriches the promise of the web. Creating a positive, successful, trust-worthy experience for users is crucial to its success. What does that mean? What is imperative for it to become the ‘next generation’ web? Most importantly, why must the usability community play a leading role to shape the Semantic Web in a positive, user-centered way?

What is the Semantic Web?

What is the Goal?

Before going into the more recent, well known descriptions and visions of the Semantic Web, we feel it might be useful to start with a somewhat older – but no less exciting and relevant – vision:

“Science has provided the swiftest communication between individuals; it has provided a record of ideas and has enabled man to manipulate and to make extracts from that record so that knowledge evolves and endures throughout the life of a race rather than that of an individual...

Our ineptitude in getting at the record is largely caused by the artificiality of systems of indexing. When data of any sort are placed in storage, they are filed alphabetically or numerically, and information is found (when it is) by tracing it down from subclass to subclass. It can be in only one place, unless duplicates are used; one has to have rules as to which path will locate it, and the rules are cumbersome...

The human mind does not work that way. It operates by association. With one item in its grasp, it snaps instantly to the next **that is suggested by the association of thoughts**, in accordance with **some intricate web of trails** carried by the cells of the brain... Yet the speed of action, the intricacy of trails, the detail of mental pictures, is awe-inspiring beyond all else in nature.”

Extracts from “As We May Think” (Vannevar Bush 1945 – emphasis added)

We have the computing power. We have the interconnectedness of the web. Are we any closer to the achievement of “enabling man to manipulate and make extracts... so that knowledge evolves and endures”? Is the Semantic Web part of the path that takes us there?

‘Semantic’ is typically defined as relating to the meaning of things, particularly words and symbols. The W3C (World Wide Web Consortium), which is leading the effort to establish the Semantic Web, has a definition on its Activity Statement page on the subject: “*The Semantic Web provides a common framework that allows data to be shared and reused across application, enterprise, and community boundaries.*” (W3C 2001). There doesn’t appear to be anything particularly transformational in that statement, and certainly no reference to the idea of ‘meaning’ in the simpler goal of sharing data – however, the statement is immediately followed by a quote from one of the seminal articles:

“The Semantic Web is an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation.”

From “The Semantic Web,” Scientific American (Berners-Lee, et al, 2001)

So, What is the Vision?

What particularly have Tim Berners-Lee and others been considering as part of a vision of the Web? In “Weaving the Web” the vision is described as:

- In the first part, the Web becomes a much more powerful means for collaboration between people...
- In the second part...collaborations **extend to computers**...
 - Machines become capable of analyzing **all the data** on the Web – the content, links, and transactions between people and computers...
 - When it arrives, the day-to-day mechanisms of trade, bureaucracy, and our daily lives will be handled by **machines talking to machines**, leaving humans to provide the inspiration and intuition...
 - The intelligent "agents" people have touted for ages will finally materialize...
 - This machine-understandable Web will come about through the implementation of a series of technical advancements and **social agreements** that are now beginning...

Summarized from “Weaving the Web” (Berners-Lee 1999 – emphasis added)

How far should we interpret the vision as going? Could we go all the way to:

- Ubiquitous access to ‘knowledge’...
- Through serendipitous discovery and collaboration...
- Via self-promoting content...
- Using omnipotent data...
- So that computers can converse intelligently...
- All on my behalf?

(Whoa!)

It raises an interesting question - does the ‘average user’ want that kind of relationship with a computer or the web? Will it help overcome our routine frustrations, interaction barriers, and the information overload we experience today? What sort of technical advances and “social agreements” are required for such a thing to become a reality?

Around the same time that Tim Berners-Lee was writing the above vision, he also had a simpler, more limited perspective (1998 [2]): “The concept of machine-understandable documents does not imply some magical artificial intelligence which allows machines to comprehend human mumblings. It only indicates a machine's ability to solve a **well-defined problem** by performing **well-defined operations** on existing **well-defined data**.”

Not all the visions come from Tim Berner-Lee or the W3C. Tom Passin, in his upcoming book “The Explorers Guide to the Semantic Web” (Passin, in press 2004), identifies a number of interpretations of the Semantic Web. The interpretations range from a focus on machine-readable data, improved searching, better annotation, to automated infrastructure for machine processing of web content, and even acting as a servant to humanity!

There is yet a third common view of the Semantic Web, which often appears in conference presentations. This is a technologist's-eye view as represented to people interested in understanding new computing approaches, XML, the Semantic Web, and the wealth of new international standards being developed:

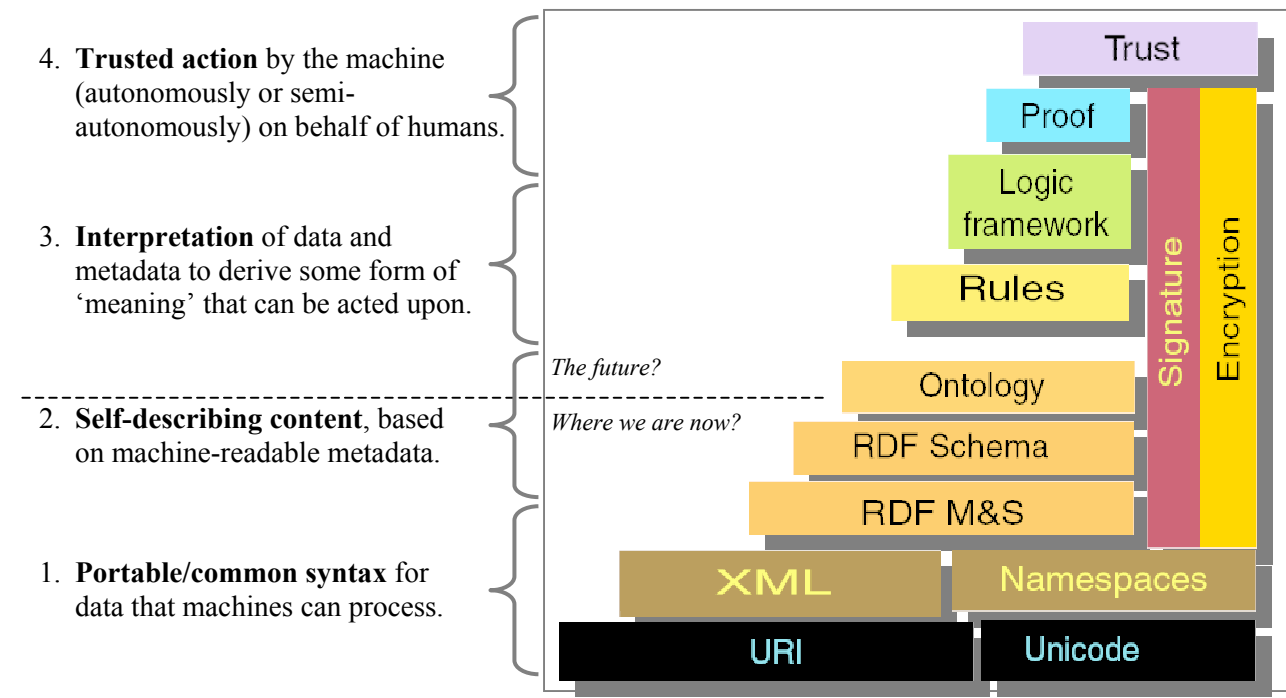


A fictitious example of an all-too-common style of presentation slide.

Well, maybe slides full of XML aren't making things that much clearer, after all.

Suffice to say that there are many visions and views. Each one provides a glimpse into the types of problems that people are trying to solve. Many of these problem areas are commonly worked on within the usability community. Later in this paper we will explore some scenarios to better understand the usability implications that are inherent in them.

The W3C uses something known as the "Semantic Web Layer Cake" (right) to illustrate what it believes to be the core elements. This provides a starting point for a less technical inventory of the parts (left).



(“Layer Cake” image, Berners-Lee, 2003, slide 15)

For now, we will take as given that the generic equivalents of these parts, in some form, can make up the Semantic Web. It is not the purpose of this paper to explore the completeness of the ‘Layer Cake’ as a model of the required parts of the Semantic Web. Indeed, there have been a number of versions of the ‘Layer Cake’ over the years, as well as other forms of models proposed by other groups – it is not unusual in complex software environments to be making up the ingredients shopping list, baking it, decorating it, and serving it at the same time!

Let’s look briefly at the parts, because what is there needs to be influenced by people who are focused on usability, and what is not there holds implications for the usability of the Semantic Web, and the role that usability specialists need to play in its future.

Part 1. A Portable/Common Syntax for Data that Machines can Process

Working on the simple principle that “if the computer can’t read it, you can’t use it,” the foundation for the Semantic Web is XML. At its heart, XML is nothing more than:

- Identification of a few character symbols that computers know to look out for when reading data, such as:
 - Pointy brackets: *< there will be something in here you can interpret >*
 - The equal sign and quote marks: *< some common data aspect = “some specific value” >*
 - (Usdin, 2003)
- Definitions of how to assemble the stuff inside the pointy brackets, mainly:
 - Elements: *< RecognizableThing >*, for example *<Heading>* or *<Name>* or *<DateOfBirth>*
 - Attributes: *< Element SpecificCriteria = “some value” >*, for example *< Heading Level=“2”>* or *<Address Type=“home”>* or *<PurchaseOrder NumberFormat=“LLL9999”>*

The benefits of using XML are:

- It doesn’t necessarily distinguish between data and ‘content’ – so it’s all data to the computer
- It’s machine independent and more future-proof – it is understood equally well by software running on Windows, MacOS, Unix, Linux, etc.
- It’s human-readable – really useful if you want to take a look to see what something is, if a specific software application isn’t available
- It allows you to describe more than just how something looks, which is a step forward into the world of data from HTML

The use of XML should give both designers and developers a lot more control over how a computer behaves when it uses data. It also puts more demands on designers to think about what they want the computer to do, and how to represent those data needs to the user to avoid ‘garbage in, garbage out’.

Because the Semantic Web is part of the Web there is also, within its foundations, the concept of addresses (Uniform Resource Identifiers, or URI). These identifying addresses can go beyond the more well know URL (Uniform Resource Locator) used on the Internet, and can conceivably address anything that is stored in a known place. In practice, for this information to be generally usable, the URI needs to identify an address that people can find on the Web. The degree to which this ‘Web find-ability’ becomes a core aspect may influence how the Semantic Web develops – there are a lot of things that may have semantic descriptions and a provenance which are not generally findable on the Web. What then?

Part 2. Self-Describing Content, Based on Machine-Readable Metadata

Without metadata, we have no Semantic Web. Because of this simple fact, a huge amount of discussion and development has focused on this aspect of the Semantic Web. We can say from personal experience and from looking at the experiences of many others – *this is hard, and it's very important*.

What do we mean by ‘self-describing’? We mean that every item of data and piece of content should describe as much about itself as it is able. Only then will it truly be able to be ‘known’ by a computer.

Descriptions might include:

- What something ‘is’ – at this moment in time
- Where it came from, who was involved along the way
- In what contexts it was created, where it is relevant/appropriate, and how it is likely to behave in those contexts
- What relationships or actions it expects to participate in, and when it is not allowed to participate
- The nature of those relationships it's participating in, given the current context (this is the heart of self-description!)

Descriptions are part of what help us – and computers – know when one thing is related to, or relevant to, another thing. Humans do this easily and naturally. Every time we describe something (for example, a maple tree), we use a short series of words that mean the same thing (maple, deciduous tree, Canadian symbol), or describe aspects of the subject (used for syrup, the one in my backyard, the leaf is shaped like this). Other people usually are quick to understand what we mean.

Anyone who uses a search tool knows that computers do not perform this ‘disambiguation’ well. It is common to be flooded with ‘hits’ because lots of documents contain similar words. It is also common to get documents that aren't relevant, even though they use certain words (if you search for “maple” you will get documents about trees, syrup, Canada, furniture, and hockey, as well as documents that contain the phrase “the Poplar is not like a maple”).

There is a scenario from a 2001 conference that illustrates hopes for self-description and meaning:

Peter investigated a point that bothered him: Bill used the term “Service” in an unusual way. He wrote: “Acme computing will run the trust rating service for semanticweb.org” (a sentence from Bill). His assistant found no problem so he hit: “service”, the assistant displayed “service in {database} equivalentTo: infrastructure.” Peter asked for “metainfo”, which raised “Updated today by negotiating with Bill’s assistant.” Peter again asked for “Arguments!”: “Service in {database} conflicts with service in {web}”. “Explain!” “In operating system and database, the term services covers features like fault-tolerance, cache, security, that we are used to put in the infrastructure. More evidence?” Peter was glad he had not to search the whole web for an explanation of this. The two assistants detected the issue and negotiated silently a solution to this problem.

Part of a Report from an EU workshop on the Semantic Web (Euzenat 2001)

From a usability perspective, some interesting questions arise:

- How is Peter’s question represented to the assistant? What sort of a “problem” is the assistant looking for that it didn’t find, and did Peter explicitly identify possible problems initially?
- Is the term “service” sufficiently self-described, and more importantly the *context* in which the term was used in conversation sufficiently clear, for an informative definition to be presented?

- Is every word in the sentence Peter typed able to be interrogated? Will the assistant interrogate “trust rating service” differently from “service”? How does Peter select one or the other?
- When might the negotiation with Bill’s assistant(s) take place? When Peter poses the question? Or is every possible question/answer negotiated in advance when two agents “detect” an issue? Is Peter informed of the detection? How does Peter review the appropriate question/answer sets?
- How does Peter remain in control of the assistant’s activities on his behalf, and how are the myriad possible problems, negotiations, and solutions filtered and presented to him?


There are a number of parts to self-description, but they fall into two general categories:

- Knowledge representation syntax (the formalism for writing specific statements about data and relationships in a way that can be read by machine)
- Knowledge representation models/structures (the formal constructs/sets of relationships between descriptive items –often described as taxonomies and ontologies)

There are a number of forms that are available to support the encoding of metadata in a more machine-readable way. The approach chosen by the W3C is the Resource Description Framework (RDF), the RDF Schema (RDF/S), and the Web Ontology Language (OWL). Other possible forms include Topic Maps (in its XML formalism, XTM), the Knowledge Interchange Format (KIF), and others. It is important to note that these are not at the same levels of abstraction, nor do they focus on the same purpose, so it is inappropriate to compare them equally. A lot of work is going on in the technical community to look at translation and connections between them.

The underlying principles of these various approaches are very similar – we move beyond simple lists of keywords (like the <meta> tag in an HTML document) and represent relationships. These relationships help us describe things in the world in a way that helps the user when in a conversation with a computer.

What might that look like? As an example, it is easy to see and interpret the relationships in the following illustration, using conceptual graph representation:



(Sowa 2001)

One of the more popular – and simple – uses of RDF to model relationships on the web is something known as ‘Friend of a Friend’ (FOAF – see Dumbill 2002, Brickley 2003). FOAF allows people to describe their relationships with other people in a way that can be easily read by a computer. It has the potential to become a machine-readable equivalent of the ‘six degrees of separation’ concept.

One of the interesting points about FOAF is that it tells other people about ‘me’ in a completely open, machine-readable way. In a very simplistic example, Duane can express that he’s a friend of Renee, without any obligation on Renee’s part to say anything about her perception of a friendship with Duane. All there is at this point is a subjective statement from Duane. There is a little bit of awareness gained (you can explore if there *may* be a connection there), but there is little reliability or verifiability in the statement. However, you know a little more about Duane – he is a little more fully described (of course, if Duane also declares he’s a friend of Abraham Lincoln, you might need to check whether he’s referring to the former President or to a living individual).

Once Renee does make such a statement, she has begun to describe herself. She has also allowed for the creation of a third piece of information – a ‘friendship’ between two parties – which may be either implied or explicitly stated in data terms. The information may take on a different meaning or reliability if twenty people who know Duane and Renee express that they are friends.


If we go beyond the simple expressions of friendship, we might discover that Duane and Renee are both related to a company. It has a client who describes a successful project outcome. Both Duane and Renee are described as having lead roles in that project. Do those statements, when woven together in the context of wanting to understand their working practices, add up to a positive reference? (We hope so!)

Why does this matter? Any person (or even software agent) using the Semantic Web to explore these relationships must have a clear indication of the nature of the data used to represent these relationships, and the reliability of that data. There also needs to be a sense of when to stop – do we need or want to investigate all twenty friends, or every member of the project team and all their descriptions, in order to come to a conclusion? At some point, the web of self-describing relationships has to align with a trust model that supports how you ‘believe’ in specific expressions.

Self-description? So what?

It helps for something to describe itself so we can be sure we know what it is, and just as importantly what it is *not*. Beyond that, the key thing to remember from a usability perspective is that **having descriptions about concepts and relationships allows us to begin describing context**. When we can describe and interpret context, we can begin to adapt the information and the interface to the specific needs of the user.

How little descriptive information do we need to make the Semantic Web work? We may never know. What happens if we don’t have everything we need to understand the conditions and context of the situation? Will the Semantic Web **only** work when ‘everything’ is stored and available? We don’t know yet. We need to evaluate when ‘enough’ is appropriate for the situations users face.



Part 3. Interpretation of Data and Metadata to Derive ‘Meaning’

Now that we’ve established the importance of terms and their relationships, it’s time for the bad news. How does a computer know the difference between what the terms/concepts mean, and **what we mean** when we use them? Again, context is the key!

This is where we move from the formative present to the idealized future of the Semantic Web. It needs methods of expressing logical arguments and exposing them to scrutiny. It needs rules – internal rules for what it does with data and how it lets data interact with each other, and external rules that express how and when it communicates with users and performs actions. These things need to be secure from tampering and verifiable. If computers are to act on our behalf, then we have to be confident that they can read and interpret information, no matter how subtle, in a way that is consistent with how we would read and interpret that same information ourselves.

This is a very dynamic, immediate thing to ask of our information resources and processing systems. If we think back to that horrible day of September 11th, 2001, we can begin to consider the issue. At 8am, we use the Internet to explore a subject, transact business, or answer a specific question. In doing so, we use the term “threat.” By 10am that morning, when we turn to the Internet and use the word “threat” we mean something almost completely different. How does the computer know that? How does all that self-describing data – that was in fact described by humans before that date and time – know that?

Is it enough to rely on proven authority and provenance? Tim Berners-Lee (1998 [2]) said: “I have no expectation of any inference capability in the SW [Semantic Web] core design... Any real SW system will work not by believing anything it reads on the web but by checking the source of any information. (I wish people would learn to do this on the Web as it is!). So in fact, a rule will allow a system to infer things only from statements of a particular form signed by particular keys.”

Does that not risk rigidifying ‘meaning’ to previously believed statements from the past? How do we resolve competing claims of provenance, in order to ‘disambiguate’ meaning? Where do we discover the location of the authoritative source? Finally, as a practical consideration, must we be connected all the time in order to maintain our ‘feelers’ out to the web of meaning?

Part 4. Trusted Action by the Machine on Behalf of Humans

The Significant Role of Agents

Agents are needed to act on our behalf, traveling and interpreting the semantic information space, doing things that have been explicitly requested by us (such as make a payment for a selected item) and implicitly identified (such as ask permission for the right to use some item of information that is being sought to complete a task).

Tim Berners-Lee talked about Agents early on in his discussions of the Semantic Web:

My agent needs to know about me

- What may people know about me?
- What do I need to know about them?
- What am I prepared to pay for?
- What will I allow myself to do?

Human interface challenge

From “World Wide Web – Hopes for the Future” slide 17 (Berners-Lee 1998 [1])

One simple point says it all: “**Human interface challenge.**” Many of us have difficulty finding the time and the instructions for configuring personal preferences on a web portal or adjusting the preference settings on our word processor. It’s hard to imagine how we will be able to make decisions about every possible implication and permission that will be required to set an agent loose on our behalf.

Prove It! Otherwise, How Do I Trust It?

We’ve heard it on the playground, in the office, and on the street: “Oh yeah? Prove it!” Tim Berners-Lee has even proposed the idea of an “Oh yeah?” button, which sounds like a great idea (1998 [1], slide 14). What would that look like? How would I decide what to point it at? A whole page? A sentence? Where the cursor is currently resting (that’s difficult on a read-only web page)? What sort of response would I get back, and would I need to understand the constructs of the logical argument, or just trust the source of the confirmation of proof?



Oh Yeah?

“The first uses, such as access control on web sites, involve validation of a previously prepared proof, not a requirement to answer an arbitrary question, find the path the [sic] construct a valid proof.”

Tim Berners-Lee (1998 [2])

It will be interesting to see what proposals arise for the way in which proof can be articulated in machine-readable and human-readable terms. It is important to reinforce the need for users to be able to articulate their expectations of what ‘proof’ means to them in different contexts, and how they expect any previously prepared proof to be wrapped in the context in which it was prepared, to see how that matches their own context. According to Tim Berners-Lee (2003, slide 24), the ‘Web of Trust’ has these attributes:

- All statements on the Web occur in some context
- Applications need this context in order to evaluate the trustworthiness of the statements
- The machinery of the SW does not assert that all statements found on the Web are “true”
- Trustworthiness is evaluated by each application

We would add: **trustworthiness is judged by users of an application, in the context of their goals.**

There has been plenty of research conducted on the attributes of interfaces and interaction that encourage trust (for example, D’Hertefelt 2000, Sisson 2000, Cheskin 2000), particularly in the arena of e-commerce. The lessons may be applicable in our future Semantic Web design discussions.

What Might the Impact Be?

We’re Overwhelmed – Will the Semantic Web Help or Add Confusion?

The Web has grown at a phenomenal rate. The amazing thing is that it has grown completely organically, using only the very simple conventions of HTML and a pervasive, increasingly inexpensive infrastructure. Some of the reasons for its success include:

- Anyone can do it, and it is now inexpensive and easy to do
- Anyone can see millions of examples of other people who have done things, and borrow the things they like best (spreading what’s good very quickly)
- There are few, or no, dependencies – I can put something up at any time, in any way, without having to rely on other pages or sites being fully developed before my work can be seen or used
- It is very forgiving – the worst thing that happens is that you don’t see something, and have to back up, and in many cases the viewer technology (the browser) accepts small inconsistencies and errors in the technical coding and still manages to display the result

The Web has always been about **the publisher**. Focusing on the quality of the user's experience is something that has had to be fought for over many years.

However, the *Semantic Web* is all about **ME, the user!** Data is described in a way that has meaning **to me**, in context of **my situation**, with agents acting as surrogates **on my behalf**. If it isn't usable, it might as well not exist!

Will the Semantic Web focus on **me**? According to stories about what it could be like, apparently so:

At the doctor's office, Lucy instructed her Semantic Web agent through her handheld Web browser. The agent promptly retrieved information about Mom's *prescribed treatment* from the doctor's agent, looked up several lists of *providers*, and checked for the ones *in-plan* for Mom's insurance within a *20-mile radius* of her home and with a *rating of excellent or very good* on trusted rating services. It then began trying to find a match between available *appointment times* (supplied by the agents of individual providers through their Web sites) and Pete's and Lucy's busy schedules. (The emphasized keywords indicate terms whose semantics, or meaning, were defined for the agent through the Semantic Web.) In a few minutes the agent presented them with a plan...

From "The Semantic Web," Scientific American (Berners-Lee, et al, 2001)

Certainly that scenario appears a lot simpler than the current experience of searching the web and then calling the insurance company's 1-800 number. But what are the assumptions in those simple words – assumptions that create challenges for the design of applications and information resources so they operate effectively in a 'semantic' world?

- Lucy has to have personalized, secure access to the web from the doctor's office.
- Her 'handheld web browser' (a PDA-sized screen?) needs to understand or be told that she is at the doctor, and that the issue is with her mother, rather than herself.
- She has to have full permission to instruct her software agents to perform on her mother's behalf – was all that permission-setting arranged in advance, or is that part of the transaction she is carrying out at the time? Does her mother need her own computer access at the same time?
- The doctor's agent, and thus computer record-keeping system, has to be able to clearly explain semantically the specific situation and required actions/prescriptions, and possibly additional instructions. It has to be instructed to do so on request. It also has to be sure that the right portion of a potentially complex diagnosis and treatment plan is what the practitioner receives, and not necessarily everything that might not be relevant to that person's work.
- The trust that Lucy's mother will place in that practitioner comes from a rating system that may need to explain, in real time on the handheld, where and how that rating system is derived, and possibly provide to her mother more background information about that practitioner, in an understandable and personally relevant form.
- Four independent schedules are coordinated in the background, and options presented – how are personal preferences, desires, and conveniences negotiated? How do Lucy and Pete describe to the computer the way that they will allocate the chauffeuring chore? How are the schedule negotiations and rules represented to the parties involved?
- Finally, the schedule options are presented. How many options are presented? Every option that fits in the calendar (imagine the clutter!)? Or only one or two? If they aren't appropriate, how many 'clicks' will Lucy and Peter have to perform to go through appropriate alternatives? How does the computer signpost the options (like search results? "2 of 47 appointments shown")?

There are a number of complex tasks inherent in this scenario, and many different people, software applications and information resources that have to interoperate and communicate. It is easy to imagine how confusing and frustrating the user experience could become without thoughtful, user-centered design. Certainly this task wouldn't be simple currently using the phone, e-mail, and the Web. What happens when the Semantic Web agents are operating autonomously or semi-autonomously on our behalf? How can we be sure that the Semantic Web is:

- So easy that anyone can describe themselves and their information semantically
- Able to clearly show what a 'good' (complete, understandable, semantically rich, trustable, not overwhelming) experience is, when much of the activity is happening in the background using semantic applications and agents
- Able to grow organically and with few dependencies, while also moving toward the 'web of meaning' idea
- Forgiving of differences in language and meaning, while also being clear and respectful of semantic 'shades of gray'

Context is King!

How do users easily and effectively describe what 'relevant' means to them at any point in time? We can often have issues with context even when dealing with people face-to-face. How well will software agents operating on a user's behalf know all the rich aspects of context and user goals at a particular point, in order to represent that situation to other agents and people, for them to then carry out their own actions on the user's behalf? There will be limitations to consider, and the impact of those limitations will make the difference between usable and not usable.

While the main *technical* focus of the Semantic Web is machine interpretation of self-describing information and services, the ultimate purpose of that interpretation is to make the web **easier and more useful for people**. To fulfill the claim of increasing relevance and value, there needs to be a way for every user to describe what they consider to be relevant and valuable to them – at any point in time, no matter how much that may change from one minute to the next. This is what we call the 'challenge of context' (Degler and Battle, 2000). We have all been disappointed by the unfulfilled promise of software applications and utilities that claim to know what we are working on or what our preferences are, only to deliver information or a service that clearly illustrates the application's rudimentary and inadequate knowledge of us, our circumstances, and our needs. **What are the risks that we end up not with meaning, but with a flood of meaningless complexity?**

The following is another scenario that helps us explore context and relevance:

During her stay at Honolulu, Clara ran into several interesting people... her assistant popped up a note with a link to a vCard that reads: "This guy's profile seems to match the position advertisement that Bill put on our intranet. Can I notify Bill's assistant?" Clara hit the "explain!" button. "I used his company directory for finding his DAML enhanced vita: he's got the required skills as a statistician who led the data mining group of the database department at Montana U. for the requirement of a researcher who worked on machine learning.", Clara hit then the "evidence!" button. The assistant started displaying "I checked his affiliation with university of Montana, he is cited several times in their web pages: reasonably trusted; I checked his publication records from publishers' DAML sources and asked Bill's assistant a rating of the journals: highly trusted. More details?". Clara had enough and left her assistant inform Bill's.

Part of a Report from an EU workshop on the Semantic Web (Euzenat 2001)

While the story is simple (and compelling – who wouldn't want candidate pre-qualification to be this easy?), there are considerations. What context would the semantic-enabled toolkit need to consider?

- Clara is away from the office. How current is her 'watch list' of job positions from the company intranet? If the company is large, how many openings are on it, and were they all investigated for all possible candidates? What level of detail fits her current context, and the company's?
- Why does the job candidate search stop with people she met? Might part of her role on that trip be to meet qualified people who are there that she doesn't know? Or was the trip mainly personal?
- If she met this person, would it be appropriate for her to add her insights or conversation details to the notification to Bill? She has the greatest situational awareness, having met the person.
- Are we assuming that the Montana U information resources are current, and specific to the context of the job opening? Is there sufficient agreement between the company semantics and Montana U's semantics about the meanings of "statistician" or "machine learning"?
- Based on what Clara knows at that time, how many questions might she have? How many different semantically labeled buttons might there be on her interface besides "explain!" and "evidence!"?

The Vital Role for Usability

The Semantic Web represents a fundamental change in how we interact with computers – what role and responsibility do we have to the users who will be affected by that change? What techniques do we use to help people – particularly developers and decision-makers – understand that change? Some of the issues that need to be addressed are being articulated (for example, Shum 2001, Shirkey 2003). We believe that there are some critical roles for usability practitioners in the arena of the Semantic Web:

- To provide approaches that help applications effectively model the user's context, profile, experience, tasks, goals, and needs
- To make the underlying interpretations and activities of semantic applications/environments apparent to the user (not hidden), while not intruding on the user (facilitating actions as background processes)
- To provide interaction techniques that keep the user in control of the environment
- To provide analysis, design, and testing techniques that effectively represent the needs of users
- To enhance the web's ability to support computer-mediated conversations between people in a way that provides increasing semantic richness, and promotes semantic representation of other modes of communication (e.g. non-verbal) between people in collaborative environments

Some Initial Imperatives – a Proposal

Imperative 1:

Put the user's goals first and foremost, not as an afterthought to technical implementations

We are in the early stages of a profoundly transforming technology – potentially on a par with the telegraph, telephone, and e-mail – in its capacity to move a functionally limited tool (the computer), using a primarily broadcast medium (the Web), toward a conversation, where the computer may begin to genuinely facilitate human interaction.

Imperative 2:

The Semantic Web must be usable and valuable – to me, the user!

One of the core concepts of the Semantic Web is that it enables **computers and people to work in cooperation** (Berners-Lee, et al, 2001). It is our contention that to do this successfully, the Semantic Web must be **usable**. This requires new methods and new techniques – as well as exploring how existing techniques can work if we use them effectively, in the right places with the right priorities.

Imperative 3:

Keep it simple

Greater richness of information and interaction runs the risk of greater complexity presented to the user – if web applications ‘know more’ about the situation and the data available, how do we focus on what’s important.

Imperative 4:

Don't rigidify Context – it must accurately represent the user's vantage point

Context is king, and yet context is subjective (subjectivity has not been the strong-suit of computing to date). There is a need for new methods that expose context, and allow a user's contextual description to be articulated. We need a balance between what can and cannot be automated. We also need a ‘macro’ view of contextual patterns – future work in usability design patterns could make an important contribution to this imperative.

Imperative 5:

Users have to apply judgment and give permission to agents – this must be clear, transparent, and directly controlled

The common theme is that computers are going to work more actively with more data on our behalf – which implies the need for more decisions and judgments about what the computer is allowed to do and not do. Whose decisions and judgments? Users! This requires information, context, understanding, and clear representations in the interface. The range and sources of information that support decisions will be immense, and yet must be simple and clear when they are presented for inspection.

Imperative 6:

Respect and show all limitations (i.e. what the computer doesn't know or can't show) for proper interpretation by the user

The computer and the Semantic Web are not – and will never be – all knowing and all-inclusive. Interpretations and actions will always be based on limited descriptions (metadata) and data. While at any point in time the amount of data is likely to be ‘good enough’ for the purpose at hand, the judgment as to whether what is available is good enough rests with the user, on whose behalf actions are taken.

Imperative 7:

The user must be able to easily and quickly control changes in meaning

Somehow, even though everything has the potential to be connected to everything else in a complex, global web of meaning – each and every one of us has to have facilities available that allow us to change interpretations and relationships that create ‘meaning’ in terms of computer processing of data. The human controls required to get the best of this technology will have to be worked out very thoughtfully.

Do We Have the Tools to Deal With This?

Is the current suite of methods and tools available to the analyst, usability specialist, or software developer sufficient to design for the Semantic Web? We expect that there will be many changes to consider in the design tools that we use. Among the things we might consider are:

- How do we model users and scenarios in a complex, individualized, personalized environment?
- How do we represent context in the design process? During development and testing?
- How do we prototype an environment that is deeply data-driven, without building it all first?
- How do we test, if the process is so deeply situational and the application behavior is so much more autonomous?
- How do we gather feedback on peoples' actual experiences of use?
- How do we describe design, since many of the 'parts' of an application may be individual, autonomous components or agents that are not coded by the same people, at the same time?

We can't know in advance everything that we will need. We have to try using our existing tools and refine them as we discover what new methods we need. We have to constantly talk about our experiences – successes and failures – and look for new ways of working.

Questions to Explore as We Get Involved in “the Dialogue”

We know the usability community must be pro-active in helping to shape the Semantic Web. What are the things we need to think about? Below are some initial questions we feel need to be explored:

- What aspects of the Semantic Web have an interface? Even if there is no interface, what are the interactions?
- How do you remain in control, when so much is going on 'behind the curtain'?
- What are the fundamental implications of agents, component technologies and web services - building blocks for web-based applications and interfaces? The more that applications are built from shared, distributed components, how much harder is it to shape the user experience?
- How do users and technologies work together to build the semantic connections? How are they expanded and enhanced when users collaborate with each other electronically?
- How are user personas, preferences, experiences and interests modeled? Once these connections are built then how are they changed as the user's expectations and needs change?
- How can we create better-managed semantic environments to help us avoid drowning in the ocean of metadata?
- How does a user's interaction with an application help filter out the myriad of possibilities, hone in on what is relevant, and make sure that the interpretations being made in the background are suitable to a particular situation that day, and not a situation faced last week or last month?
- How much of this can be done without requiring many hours of my time informing my machines/appliances of my desires, circumstances and vocabulary preferences?
- What interactions are required so that semantics can support online conversations and interactions? Can semantic technologies enable person-to-person and group interaction, and make the idea of online community more practical? What facilities are required to make that easier?
- How do we keep the language from becoming stale - or, put another way, how do we *maintain* the Semantic Web?

References

- Berners-Lee, Tim. World Wide Web - hopes for the future. Keynote speech at the Quebec Internet Forum, Montreal, Canada. March 1998 [1]. Online: <http://www.w3.org/Talks/1998/0227-WebFuture/overview.htm>.
- Berners-Lee, Tim. What the Semantic Web can Represent. W3C. September 1998 [2]. Online: <http://www.w3.org/DesignIssues/RDFnot.html>.
- Berners-Lee, Tim, with Mark Fischetti. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web by Its Inventor*. San Francisco: Harper, 1999.
- Berners-Lee, T., Hendler, J., Lassila, O. The Semantic Web. *Scientific American*, May 2001. Online: <http://www.scientificamerican.com/2001/0501issue/0501berners-lee.html>.
- Berners-Lee, Tim. The Semantic Web and Research Challenges. W3C. 2003. Online: <http://www.w3.org/2003/Talks/01-sweb-tbl/>.
- Brickley, Dan. Identifying things in FOAF. RDFWeb and the FOAF (Friend of a Friend) project. July 2003. Online: <http://rdfweb.org/mt/foaflog/archives/000039.html>.
- Bush, Vannevar. As We May Think. *The Atlantic Monthly*. July 1945. Online: <http://www.theatlantic.com/unbound/flashbks/computer/bushf.htm>.
- Cheskin Research Study on eCommerce Trust. Cheskin Research. July 2000. As found online (PDF): <http://www.debmcdonald.com/trust.htm>.
- Degler, D. and Battle, L. Knowledge Management in Pursuit of Performance: the Challenge of Context. *Performance Improvement Journal*, 39(6), July 2000, 25-31. Online: <http://www.ipgems.com/writing/kmcontext.htm>
- D'Hertefelt, Sim. Trust and the perception of security. *The Interaction Architect*. January 2000. Online: <http://www.interactionarchitect.com/research/report20000103shd.htm>.
- Dumbill, Edd. XML Watch: Finding friends with XML and RDF. IBM DeveloperWorks. June 2002. Online: <http://www-106.ibm.com/developerworks/xml/library/x-foaf.html>
- Euzenat, Jérôme (Ed.). Report from the NSF-EU Workshop "Research Challenges and Perspectives of the Semantic Web", Sophia- Antipolis, France. October 2001. Online: <http://www.ercim.org/EU-NSF/semweb.html>.
- Passin, Thomas. *Explorer's Guide to the Semantic Web*. Greenwich, CT: Manning Publications, in press (2004). Online reference: <http://www.manning.com/passin>.
- Shirkey, Clay. The Semantic Web, Syllogism, and Worldview. First published on the "Networks, Economics, and Culture" mailing list. November 2003. Online: http://www.shirky.com/writings/semantic_syllogism.html
- Shum, Simon Buckingham. Human-Centered Aspects of the Semantic Web. Presented at NSF-EU Workshop "Research Challenges and Perspectives of the Semantic Web", Sophia-Antipolis. October 2001. Online: http://www.ercim.org/EU-NSF/semweb/slides/buckingham-shum/SBS-NSF-EU-SemWeb_files/frame.htm.
- Sisson, Derek. Trust & Trustworthiness. *Philosophe.com*. February 2000. Online: <http://www.philosophe.com/commerce/trust.html>.
- Sowa, John (Ed.). Conceptual Graphs Draft Working Document. ISO/JTC1/SC 32/WG2. April 2001. Online: <http://users.bestweb.net/~sowa/cg/cgstand.htm>.
- Usdin, B.T. It's the Markup, Stupid! Presented at Extreme Markup Languages 2003. Montreal, Quebec, Canada. August 2003. Online: <http://www.idealliance.org/papers/extreme03/html/2003/Usdin01/EML2003Usdin01-toc.html>.
- W3C Semantic Web Activity Area. Established 2001. Online: <http://www.w3c.org/2001/sw/>.

About the Authors

Duane Degler, IPGems

Email: ddegler@ipgems.com

Duane Degler has consulted on a broad range of organizational performance issues, with experience in user-centered design, information/knowledge strategy, process analysis, international system implementation, training, and multimedia. He is currently involved in the development of the Pensare Group, as well as supporting clients such as the Social Security Administration. He has worked with government agencies, legal firms, insurance companies, major auto manufacturers, and several health organizations. His work and writings are breaking new ground in design and content/knowledge management related to the adoption of semantic technologies. Prior to that, his interface/system design work has received three US awards. He has managed pioneering multimedia projects in the '80s and spent much of the '90s in the UK involved in knowledge management research and consulting, contributing to his multi-disciplinary approach to interaction design. Duane holds a BS in Broadcast Communications and an MS in Organizational Communications.

Renee Lewis, Pensare Group

Email: renee.lewis@pensaregroup.com

Renee Lewis has over 18 years of experience developing and delivering object and Internet-based solutions for government and commercial clients. She is currently President and Founder of the Pensare Group, LLC, a consulting company specializing in technology adoption and communication strategies. Today, she supports clients such as the Social Security Administration with a contextual retrieval engine, and three other advanced research companies focused on bayes net, analytic search engines, and collaboration portals. She has held leadership positions in technology companies focused on market areas such as drug safety, e-Learning, health care, and telecommunications. She also spent 10 years as a consultant for Booz, Allen and Hamilton, focused on implementing distributed database architectures for both commercial and government clients. She holds a BS from the Pennsylvania State University and an MS from George Washington University in software engineering.