Undefined 0 (0) 1
IOS Press

A Taskonomy for the Semantic Web

Editor(s): Krzysztof Janowicz, Pennsylvania State University, USA and Pascal Hitzler, Wright State University, USA
Solicited review(s): Manfred Hauswirth, DERI, National University of Ireland, Galway, Ireland and Mark Gahegan, University of Auckland, New Zealand
Open review(s): Frank van Harmelen, Donovan Artz, Pascal Hitzler, Osma Suominen

Tom Heath
Talis Systems Ltd.
Knights Court, Solihull Parkway,
Birmingham Business Park, B37 7YB,
United Kingdom
Email: tom.heath@talis.com

Abstract. The modalities of search and browse dominate current thinking about interaction with the Web. Given the Web's origins as a global hypertext system, it is understandable that these document-centric interaction patterns prevail. However, these modalities alone are inadequate as a conceptual model of interaction with the global Linked Data space that is the Semantic Web. Realising the full potential of the Semantic Web requires a fundamental reconsideration of Web interaction patterns in the light of Linked Data, and this renewed conceptualisation must drive the research agenda related to user interaction and the Semantic Web. This paper argues that a fundamental understanding of user goals and tasks is the appropriate perspective from which to approach the research and development of Semantic Web applications. However, the *Web* in Semantic Web should not detract from the potential for cross-platform data interoperability enabled by the Semantic Web technology stack. In this context we propose a *taskonomy* of data- and object-centric user tasks derived from an analytical abstraction of existing research, not simply in the fields of Web search and browse but also email and Instant Messaging, that can help shape the direction of research and application development in the Semantic Web field.

Keywords: Semantic Web, Linked Data, Email, Instant Messaging, User Interaction, Taskonomy, Activity-Centred Design.

1. The Ubiquitous Document Metaphor

"The correct approach to the support of behavior is activity-based classification." [15]

Look closely enough and it becomes apparent that document-centric metaphors are fundamental to the concept, design and realisation of our most widely used computing systems. Computers have desktops, files live in folders, we add pages to our Web sites. The terminology of email (*mailbox*, *postmaster*, *blind carbon copy*, *attachments*) reflects a communication platform conceived in the era of memos and postal systems. The Web emerged from a desire to share information between scientists [2], and owes much to the influence of the document-centric fields of hypertext and information retrieval:

"Computers give us two practical techniques for the man-knowledge interface. One is hypertext, in which links between pieces of text (or other media) mimic human association of ideas. The other is text retrieval, which allows associations to be deduced from the content of text. In the first case, the reader's operation is typically to click with a mouse (or type a reference number) – in the second case, it is to supply some words representing that which he desires. The W3 ideal world allows both operations, and provides access for any browsing platform." [1]

1.1. Classifying Web Search

This original model of the Web has defined our view of how it is used ever since, with searching and browsing remaining the prevalent lens through which we view human interaction with the Web [16], even occurring in more data-centric analyses such as [19].

The dominance of the document-metaphor manifests itself not only in the computing applications we develop, but in the research conducted to try and understand how people use the Web in practice. While work such as [6] has attempted to understand the range of activities conducted on the Web, e.g. banking, job-hunting, or finding travel information, numerous others have attempted to identify and classify various forms of Web search:

- Guha, McCool and Miller [7] distinguish between *navigational* searches, where "the user is using the search engine as a navigation tool to navigate to a particular intended document" and *research* searches, where the user is "trying to locate a number of documents which together will give him/her the information s/he is trying to find" (pp. 702).
- Broder [4] identifies three types of Web search: *navigational* and *informational* searches, that map closely onto the *navigational* and *research* searches of Guha et al. [7], and *transactional* searches where the user intends "to reach a site where further interaction will happen" (pp. 6), such as a shopping site or a site where images or music can be downloaded.
- Related work by Rose and Levinson [17] yielded top-level categories with many similarities to those of Broder [4], in addition to a number of more specific sub-categories (e.g. *download*, *entertainment*, *interact*, and *obtain*).
- Morrison, Pirolli, and Card [13] describe a taxonomy of Web activities with three variables: the *purpose* of a search, the *method* used, and the *content* of the information being searched for.

1.2. Distortions in the Search-centric Lens

As comprehensive models of Web search, these classifications have a number of limitations. For example, in the work of Broder [4], the range of possible *transactions* a user may wish to perform, and the underlying reasons for wishing to perform them, are not explored. Similarly, consideration is not given to why a user may wish to *navigate* to a particular Web site or document. Presumably this destination does not represent an end in itself, but part of the strategy for performing another task, such as finding a phone number or arranging car rental.

In addition, while Rose and Levinson [17] give a number of examples to illustrate their sub-categories, the distinctions between them are often based on technical aspects of how the target object will be used, rather than the fundamental nature of the task the user is performing. For example, the target of the *download* goal is "a resource that must be on my computer or other device to be useful" (pp. 15), and the authors cite the example of a piece of software. However, the same definition could equally apply to the adult movie example used to illustrate the *entertainment* sub-category. In both cases the key feature is the attempt to *locate* something specific; drawing arbitrary category distinctions between these serves only to obscure the commonality in the underlying goal of the user.

At first glance the variables proposed by Morrison, Pirolli and Card [13] appear neatly defined. However, the classification of some activities suggests the variables may not be mutually exclusive in the form presented by the authors. For example, some methods are seen to be triggered by a particular goal (*find*, *collect*) whereas others (*explore*, *monitor*) are not. On the contrary, there is a strong argument that *explore* and *monitor* represent goals in their own right, and should be classed under *purpose*.

The focus of these studies on classifying search behaviors may be valuable in informing the ongoing development of Web search engines. However, by taking a search-centric perspective on Web usage these classifications may often obscure the true goal of the user in being online and perpetuate the ubiquity of the document metaphor in attempts to understand how the Web can support people in achieving their goals. The search-specific focus of these studies means none can account for more complex tasks performed on the Web. While the *resource-interact* goal of Rose and Levinson [17] and the *transactional* queries of Broder [4] suggest an intention to carry out further interaction beyond the search (perhaps indicating a greater overall goal), the search itself is still seen as the user's primary task. No mention is given of, for example, *arranging a holiday* as an overarching reason for being online, or even for carrying out a search. While analysis of search query logs is unlikely to show many queries such as "arrange holiday", this likely reflects an awareness among users of the narrow scope of search engines rather than a lack of desire the use the Web for this purpose.

2. That Which Is Not Search

A number of studies have investigated a broader range of tasks beyond simply Web search:

- Sellen, Murphy, and Shaw [18] describe six types of activity carried out on the Web, based on a study of Web use by twenty-four knowledge workers: *finding*, *information gathering*, *browsing*, *transacting*, *communicating*, *housekeeping*).
- Kellar, Watters, and Shepherd [11] report a study into how people use "web browser navigation

mechanisms", in which participants were asked to classify their Web usage according to the following tasks: *fact finding*, *information gathering*, *just browsing*, *transactions*, and *other*. This classification was informed by previous studies, such as [18], but refined based on the findings of pilot studies with users.

– Kellar [10] refines the top level classification of Kellar et al. [11], grouping *fact finding*, *information gathering*, and *browsing* under an *information seeking goal*; *transactions* are joined by *communications* under *information exchange*; lastly a new top-level goal of *information maintenance* is added, containing a single *maintenance* task.

The classifications of [18,11,10] are not limited to describing variations of Web search and do attempt to capture the user's needs or goals in using the Web, with some success. For example, the concept of *transacting* is a first-class citizen of all three classifications, without the degree of indirection present in the notion of a *transactional search*. In contrast, it is not clear whether the notion of *browsing* represents a goal in its own right, or simply an activity in support of some further (unspecified) goal. As already noted, this confusion of purpose and method is a consistent theme in attempts to understand how people use the Web and how the Web supports individuals in meeting their goals.

2.1. Fragmented Platforms for Communication

Perhaps the most significant limitation in all the work reviewed to this point is the focus purely on Web-based tasks. Sellen, Murphy and Shaw [18] define their *communicating* task as "Using the Web in order to participate in chatrooms or discussion groups" (pp. 229), but exclude email activities from the data. Similarly, Kellar [10] introduces a *communications* task, but uses email as an illustrative example.

As intuition would suggest, and these findings corroborate, the Web is regularly used for two-way communication of the sort conducted through email, chatrooms and discussion groups. *Communicating* accounted for just four percent of observed activities by Sellen et al.[18], however it is likely that the inclusion of email in the analysis, in addition to increased use of Webmail services in recent years and the advent of Web-based *microblogging* services, would result in a significantly higher percentage if the study was repeated.

An examination of analogous research into how people use email reveals, unsurprisingly, significant usage of this *platform* to conduct asynchronous communication. For example, in a study of 20 office workers, of varying roles, Whittaker and Sidner [20] focus on three main email functions: *task management*, *personal archiving*, and *asynchronous communication*. Of particular note in this case is that Whittaker and Sidner found evidence of email being used for a significantly wider range of tasks than purely asynchronous communication, for which is was originally conceived. They refer to this process as *email overload*.

Additional evidence for the overloading of email as a platform comes from an investigation of members of a large research laboratory [12], in which email was identified as supporting the following work functions: *information management*, *time management*, and *task management*. It was noted that those participants for whom email served an information management function may have job roles that involve staying abreast of developments by tracking information, a task that bears a noteworthy similarity to the *monitoring* activities identified in the work of Kellar [10] and Morrison, Pirolli and Card [13].

Ducheneaut and Belotti [5] identify additional ways in which people use email to perform tasks. For example:

– All but one participant reported regularly using email to exchange files with others
– Eighty percent of respondents reported using email to arrange meetings
– Seventy-two percent of participants used email to make decisions.

In the case of the latter point, how this decision-making is achieved in practice is not discussed by Ducheneaut and Belotti. Presumably email is used as a medium for discussion from which a decision can be reached.

While mainstream use of Instant Messaging (IM) is a more recent phenomenon than adoption of email and the Web, research in this area reveals noteworthy overlap with the tasks identified in the literature discussed above related to email and the Web. While use of IM for *communicating* is to be expected, the literature demonstrates that this platform also plays a role in tasks such as *finding* and *arranging*.

Following a study of how IM is used in the workplace, Nardi, Whittaker and Bradner [14] describe a number of informal communication tasks supported by this technology:

– *quick questions and clarifications*

- *coordination and scheduling*
- *organising impromptu social meetings*
- *keeping in touch with friends and family*

Whilst the latter three tasks seem rather distinct, further examination shows that they share a common theme of the participants making arrangements. Such arranging may vary from purely work-related, such as scheduling a meeting, to workplace social arrangements such as meeting colleagues for lunch, to coordinating social activities with friends and family outside of work.

Isaacs, Walendowski, Whittaker, Schiano and Kamm [9] refine the work of Nardi et al. [14] by showing how the prevalence of particular tasks in IM usage varies across different types of users. Frequent IM partners, or those seen as heavy users, used the medium predominantly for working together. Lighter users or infrequent partners generally used IM to carry out scheduling tasks.

At a general level, the findings of Isaacs et al. [9] are supportive of Nardi et al. [14]. Evidence was found for a number of similar functions:

- *simple questions and information* bears a strong likeness to Nardi et al.'s *quick questions and clarifications*, whilst placing more emphasis on the simplicity of the exchange rather than the duration over which it occurred.
- directly equivalent *scheduling* and *coordination* tasks are present in both classifications.
- the *social talk* task of Isaacs et al. is comparable to aspects of the *keeping in touch with friends and family task*

The evolution of email into a platform for a wider range of tasks than those for which it was originally intended is consistent with the development of the Web into a general purpose platform for a variety of tasks that go beyond its original role as an information distribution platform. While evident in literature and practice, it should be noted that the *function creep* affecting email and the Web does not imply that these platforms are well adapted to the tasks for which they are being used. On the contrary, they may represent the best of several poor options.

With the literature also indicating that Instant Messaging supports a broad range of heterogeneous tasks, any comprehensive attempt to understand the tasks and goals of Web users must not examine this platform in isolation, but instead take a holistic view of how all Internet platforms (e.g. Web, email, IM) are used. If the goal of understanding user tasks and goals is to drive improvements in available applications and services, this understanding must be shaped not by existing applications and services that embody potentially counterproductive assumptions and metaphors, but by the notion of *activity-centred design* [15] and a fundamental examination of the underlying goals of Internet users. To do otherwise would be to confuse purpose with the method employed.

3. A Web of Actions

To summarise the arguments so far, it is apparent from the literature that Internet platforms such as email, IM and the Web are widely used to support tasks for which they were not originally intended. While this is acceptable (of course), and perhaps an inevitable indicator and consequence of their success, it is not necessarily optimal from a user perspective, as applications developed for one purpose may not be well adapted to others. Is your email client optimised for task and project management? In addition, many of the applications we use to interact with email and the Web, and their underlying conceptual models, reflect a document-centric perspective that is not adequate for a world of Linked Data [3].

Linked Data, and the Semantic Web that has arisen from the large-scale publication of Linked Data, is about *things* and the *connections* between things. Linked Data is about the ability to publish descriptions of any aspect of any thing. It's about giving *identifiers* to those things, and maybe even interacting directly with those things, rather than just with documents that describe them.

A Web that is document-centric only enables users to interact directly with documents; it does not allow users to interact with or perform actions on the things described by those documents. The Linked Data paradigm changes that, by encouraging data publishers to assign HTTP URIs to any object or concept they wish to refer to. The existence of these identifiers paves the way for applications that support direct interaction with things identified by URIs, or at least interaction that is less incumbered by the layer of indirection inherent in document-centricity.

If Linked Data is about *things* rather than just *documents*, what happens to the old metaphors that underpin so many of our computing systems? Are desktops and filing systems appropriate metaphors for organising and accessing things that are not documents,

or is a broader perspective required? Does consideration not of what can be done with documents, but what can be done with *things in general* provide a different perspective – a set of thing-centric and action-centric metaphors to shape and inspire the Linked Data applications we build? What Linked Data brings to the table are the means to identify the things we want to act upon, and also to describe the kinds of actions that it is possible to perform on things of certain types.

In a world where people are identified by URIs, should a person *A* who wants to share a photo with person *B* have to choose between multiple platforms (e.g. email, IM, photo-sharing Web site, social networking Web site) in order to share the photo? Does it make conceptual sense to create a new email message and then attach the photo to that *document*, or simply to post the photo to the URI of the recipient and allow her to decide how it is handled on arrival? In the latter case, the recipient benefits because she gets to *choose* how and where the photo is received and stored, while the sender benefits precisely because he does not have to.

Similarly, is the fragmented nature of current communication channels optimal, whereby a person *C* wanting to notify person *D* of something must choose between multiple channels through which to achieve this goal, many of which may be suboptimal for the recipient at a particular time? Should it not be up to the recipient to choose the notification method, with the sender simply posting notifications to a canonical identifier for that recipient, perhaps accompanied by some indicator of the perceived urgency of the notification?

4. A Taskonomy for the Semantic Web

With these questions in mind, the following list presents a *taskonomy* of user activities and goals online. This taskonomy was developed by distilling the tasks and activities identified in previous literature, and removing those that represented *means* or *methods* rather than *ends* or *purposes*, or only reflected artefacts of existing Internet platforms.

- **Locating**: Looking for an object or chunk of information which is known or expected to exist.
  Example: Locating an article from a journal, an image for a school project, a colleague's phone number, or information about a book a friend recommended.
- **Exploring**: Gathering information about a specific concept or entity to gain understanding or background knowledge of that concept or entity.
  Example: Exploring a philosophical theory to understand its central tenets; getting background information about an organization before a job interview.
- **Grazing**: Moving speculatively between sources with no specific goal in mind, but an expectation that items of interest may be encountered.
  Example: Following links that spark your interest on someone's blog.
- **Monitoring**: Regularly or repetitively checking known sources that are expected to change, with the express intention of detecting the occurrence and nature of changes.
  Example: Monitoring news Web sites during an election; checking email accounts for new messages; watching discussion fora for new ideas or information.
- **Sharing**: Making an object or chunk of information available to others.
  Example: Sharing holiday photos with a colleague; uploading a journal article to your personal Web site.
- **Notifying**: Informing others of an event in time or a change of state.
  Example: Emailing a group of friends to tell them you will be going to a concert at the weekend.
- **Asserting**: Making statements of fact or opinion available, with no discursive expectation.
  Example: Writing a review of a film, or stating on your Web site that you own a certain book.
- **Discussing**: Exchanging knowledge and opinions with others, on a specific topic.
  Example: Posting a comment on a discussion forum stating that you disagree with a previous post, explaining why, and then receiving responses from others.
- **Evaluating**: Determining whether a particular piece of information is true, or assessing a number of alternative options in order to choose between them.
  Example: Choosing which film to see at the weekend, based on what's showing, where, and at what time.
- **Arranging**: Coordinating with third parties to ensure that something will take place or will be possible at a certain time.
  Example: Arranging travel and accommodation for an international conference.

– **Transacting**: Transferring money or credit between two parties.
  Example: Paying a bill.

With this taskonomy as a reference point, what forms of applications should we develop to exploit the unique capabilities of the Semantic Web?

Starting with the classic modalities of *search* and *browse*, we should not be developing Semantic Web search engines and browsers that crudely apply existing (document-centric) interaction styles to Linked Data [8]. Instead we should build services that allow us to *locate* specific information as efficiently as possible by incrementally supplying as much concrete information as we can and as is needed in order to narrow the search space sufficiently. This does not mean simply tweaking *Query by Example* interfaces to work over RDF data, but enabling query terms to be combined with background contextual information about the user in order to refine the result set.

Interfaces for *exploring* should not just be document browsers, but applications that integrate and summarise information about a specific thing of interest, based on its type, and adapt the interface accordingly.

Monitoring applications need to be able to adapt their interfaces based on the rate of change of different information sources, and the relative significance of these changes. Some information sources (e.g. Twitter) will change frequently with little consequence, while others (e.g. natural disaster warning systems) will change rarely but carry great significance. Effective monitoring applications that do not fragment user attention across multiple channels will need to account for each combination of information significance and rate of change, and adjust their behaviour accordingly.

Sharing, *notifying*, *asserting*, and *discussing* are currently supported by applications that frequently tie data to the application or system in or through which it was created. For example, notification emails stay in email systems, discussion forum posts stay unconnect to related posts made in different fora, files are uploaded to specific systems and then shared with others, rather than vice versa. A shift is needed such that applications emphasise the posting of notifications, discussion points, assertions, files, etc. into the Web at large, from where they can be retrieved as required by authorised parties, rather than simply into an application-specific silo.

Applications wishing to support *evaluating* and *arranging* may stand to gain the most in the near term from Linked Data and the Semantic Web. It is not hard to envisage how price comparison Web sites could be enhanced through Linked Data, such that products can be evaluated not simply based on price, but on local availability, delivery times, product reliability, guarantee terms, and environmental impact. This capability is feasible at present, but very costly due to the complexity of integrating data from numerous sources, each with proprietary interfaces.

Similarly, many domain specific *arranging* applications exist, such as flight comparison and booking Web sites. Where these applications fall short is in their rigidity; integration of arbitrary data relevant to a trip, but not specifically flight related, comes at a significant cost and not all types of information will warrant the investment despite potential value to a long-tail of users.

How *transacting* applications will truly benefit from the Semantic Web is not immediately clear. In one respect *transacting*, *sharing*, and *notifying* have much in common: in all cases the Semantic Web infrastructure allows a recipient to be uniquely identified, independently of any specific application or service. The result of this could be a democratisation of online payment services for end users, supported by common protocols and standards for payment interoperability.

Grazing, as defined above, is an activity with no specific, explicit user goal. In contrast it is likely to serve as a *displacement activity* that allows the user to defer performance of another (likely more important) task. Key factors in *grazing* would appear to be novelty, serendipity, and human interest. With much of the promise of the Semantic Web centred on increased precision, it is unclear what form a *grazing* application for Linked Data may take. Further research may increase our understanding of *grazing* and reveal whether this is a meaningful match for Semantic Web technologies.

5. Conclusions

As the adoption of Linked Data and Semantic Web principles and technologies continues, informal questions are increasingly being asked about the kinds of applications that could and should be developed to make best use of these technologies. Meaningful answers to these questions can only be achieved through principled analysis that attempts to understand the areas in which Linked Data and the Semantic Web can make a unique contribution relative to conventional technologies.

The fundamental shift from *document-centricity* to *thing-centricity* brought about by the Linked Data paradigm creates opportunities for new forms of *activity-centred applications* and also challenges the research and development community to reassess the established metaphors that underpin computing applications and services. The ubiquity of documents in human culture suggests that alternative metaphors may not be easily identified, however the taskonomy presented in this paper can form the basis for discussion and innovation in the research community that can begin to address these challenges.

References

[1] T. Berners-Lee, R. Cailliau, J.-F. Groff, B. Pollermann, World-Wide Web: The Information Universe, *Electronic Networking: Research, Applications, and Policy*, **2**(1) (1992), 52–58.

[2] T. Berners-Lee, R. Cailliau, A. Luotonen, H.F. Nielsen, A. Secret, The World-Wide Web, *Communications of the ACM* **37**(8) (1994), 76–82.

[3] C. Bizer, T. Heath and T. Berners-Lee, Linked Data - The Story So Far, *International Journal on Semantic Web and Information Systems* **5**(3) (2009), 1–22.

[4] A. Broder, A Taxonomy of Web Search. *ACM SIGIR Forum* **36**(2) (2002), 3–10.

[5] N. Ducheneaut, V. Bellotti, E-mail as Habitat: An Exploration of Embedded Personal Information Management, *Interactions* **8**(5) (2001), 30–38.

[6] C. Emmanouilides, K. Hammond, Internet Usage: Predictors of Active Users and Frequency of Use, *Journal of Interactive Marketing* **14**(2) (2000), 17–32.

[7] R. Guha, R. McCool, E. Miller, *Semantic Search* In Proc. 12th International Conference on World Wide Web (2003).

[8] T. Heath, How Will We Interact with the Web of Data?, *IEEE Internet Computing* **12**(5) (2008), 88–91.

[9] E. Isaacs, A. Walendowski, S. Whittaker, D.J. Schiano, C. Kamm, *The Character, Functions, and Styles of Instant Messaging in the Workplace*. In Proc. ACM Conference on Computer Supported Cooperative Work (2002).

[10] M. Kellar, *An Examination of User Behaviour During Web Information Tasks*, In Proc. CHI'06 Extended Abstracts on Human Factors in Computing Systems (2006).

[11] M. Kellar, C. Watters, M. Shepherd, *The Impact of Task on the Usage of Web Browser Navigation Mechanisms*, In Proc. Graphics Interface (2006).

[12] W.E. Mackay, Diversity in the Use of Electronic Mail: A Preliminary Inquiry, *ACM Transactions on Office Information Systems* **6**(4) (1988), 380–397.

[13] J.B. Morrison, P. Pirolli, S.K. Card, *A Taxonomic Analysis of What World Wide Web Activities Significantly Impact People's Decisions and Actions*, In Proc. SIGCHI Conference on Human Factors in Computing Systems (2001).

[14] B. Nardi, S. Whittaker, E. Bradner, *Interaction and Outeraction: Instant Messaging in Action*, In Proc. ACM Conference on Computer Supported Cooperative Work (2000).

[15] D.A. Norman, Logic versus usage: the case for activity-centered design, *Interactions* **13**(6) (2006), 45,63.

[16] C. Olston, E.H. Chi, ScentTrails: Integrating Browsing and Searching on the Web, *ACM Transactions on Computer-Human Interaction* **10**(3) (2003), 177–197.

[17] D.E. Rose, D. Levinson, *Understanding User Goals in Web Search*, In Proc. 13th International Conference on World Wide Web (2004).

[18] A.J. Sellen, R. Murphy, K.L. Shaw, *How Knowledge Workers Use the Web*, In Proc. SIGCHI Conference on Human Factors in Computing Systems (2002).

[19] F. van Harmelen, A. ten Teije, H. Wache, *Knowledge Engineering rediscovered: Towards Reasoning Patterns for the Semantic Web*, In Proc. 5th International Conference on Knowledge Capture (2009).

[20] S. Whittaker, C. Sidner, *Email overload: exploring personal information management of email*, In Proc. SIGCHI Conference on Human Factors in Computing Systems (1996).