Supporting User Tasks and Context: Challenges for Semantic Web Research

Tom Heath, Martin Dzbor, Enrico Motta

Knowledge Media Institute, The Open University,
Walton Hall, Milton Keynes, MK7 6AA, United Kingdom
{t.heath, m.dzbor, e.motta}@open.ac.uk

Abstract. Whilst the tasks users perform online are often complex and wide-ranging, the tools currently available may not adequately support them. Attempts to classify user behaviors online have tended to focus on the medium of the web, where searching and browsing are seen as the primary modes of interaction. This paper introduces a comprehensive user-oriented classification of online tasks that emphasizes the user's goals without assuming the use of particular internet tools or technologies. Taking greater account of a user's context is also discussed as an essential component in better supporting performance of tasks online. Finally we consider how Semantic Web technologies can support the development of task-focused context-aware tools.

1 Introduction

The internet provides a platform for users to perform many varied tasks, such as finding information, exploring new ideas and communicating with others. In many circumstances this platform is immensely powerful and user tasks are well supported. For example, someone wanting to find large numbers of documents on a particular subject is likely to have success with regular search engines.

However, not all tasks that users perform (or may wish to perform) online are well supported by current tools and technologies. Consider the following scenarios:

1.1 Scenarios of Internet Usage

Locating a Book. Juan wants to buy a present for his cousin, and is looking for a book that Alice had read and recommended to him. He thinks the book is called "The Sergeant's Guitar", but he can't remember the author. Searching his favourite online bookshop for this title returns no results. Juan has to contact Alice, who tells him the book is actually called "Captain Corelli's Mandolin". With this clarification Juan is able to locate the book in the online bookshop and orders it for his cousin.

Arranging a Trip. Matt is arranging travel from his office in Liverpool to a conference being held in Slovenia. Using a travel web site he looks for flights from local airports to the Slovenian capital Ljubljana. Whilst some flights are available they are infrequent and expensive. Knowing that Adam has been to Slovenia before, Matt consults him for advice before making further plans. Rather than flying to Ljubljana, Adam recommends booking a cheap flight to Klagenfurt in Austria with a budget airline; from there frequent trains run across the border to Slovenia. Whilst the total journey time will be slightly longer, the tickets will be substantially cheaper than if he were to fly directly to Ljubljana.

On the conference web site Matt reads that there is a train station near the conference venue. He follows a link to the web site of the rail company, checks the online timetable, and finds that trains run directly to this station from Austria. Revisiting the conference web site he checks the list of recommended hotels and visits each of their web sites, but finds they're all full for the duration of the conference. He remembers that Adam recommended staying with a local family, and that a Tourist Office could arrange this. He locates the appropriate tourist office web site through a search engine and sends them an email explaining his requirements.

2 Problem Analysis

In the scenarios above, the users expend considerable time and attention in completing their tasks. Whilst the outcomes are generally successful, Juan and Matt encounter a number of obstacles along the way. Some of these obstacles pertain to the specific tools available to them, whilst others reflect wider issues of the technologies and architectures of the internet in its current form.

Query Precision. When Juan is unable to remember the exact title of the book he is looking for, the search engine on the bookshop web site isn't able to accommodate his imprecise query; it takes his query literally and returns no results, even though the terms he has entered bear a strong semantic relation to the real title of the book. As far as the search engine is concerned *captain* has no relation to *sergeant*, as the engine has no representation of the semantic links between terms, just of their linguistic syntax. Furthermore, it certainly isn't aware that Juan knows Alice, and that the "Captain Corelli's Mandolin" he is looking for is the same book that she reviewed favourably on her web log.

Manual Coordination. Planning his journey to the conference requires Matt to make separate arrangements with many different parties, each of which is largely unaware of his overall goal. The travel web site Matt originally uses can only provide information about flight routes he specifically requests. It is incapable of reasoning about alternative means of reaching the same destination, or of using knowledge held by Matt's social network to help complete the task. Similarly the airline is unaware of his final destination and so cannot automatically provide information about train connections from the airport. The tourist office may be aware that he'll be attending a conference if he mentioned it in his email, but they are unlikely to know that the

conference starts early every day so his hosts will need to provide breakfast before 7am. Ensuring all of these conditions are met falls to Matt. Information about the task is not shared or reused, meaning he must explicitly state his requirements to each party and manually assemble information from the various sources if his task is to be completed. All other conference delegates must do the same.

In both these cases the user makes the best use of the tools available to them on the internet, even though these tools might not be well adapted to the true task the user is trying to perform. Furthermore, the same tools take little account of the user's context, such as the knowledge and previous experiences of those around them, when often this provides crucial assistance in performing a task.

3 Conceptualising User Tasks Online

To assess how well existing tools support users in completing tasks online, and how they might be better supported, it is important to understand the types of tasks people perform on the internet. The majority of literature in this area focuses specifically on the medium of the web rather than the internet as a whole, an issue discussed in greater detail below.

3.1 Web Activities as Forms of Searching and Browsing

Previous research has sought to identify and classify user behaviours on the web, mainly by identifying specific modes of searching or browsing. At the most basic level Guha, McCool and Miller [1] distinguish between *navigational* and *research* searches. In a navigational search "the user is using the search engine as a navigation tool to navigate to a particular intended document", whereas a research search is characterised by the user "trying to locate a number of documents which together will give him/her the information s/he is trying to find" (pp 702).

Broder [2] describes a taxonomy consisting of three types of web search: *navigational*, *informational*, and *transactional*. The navigational and informational types map closely onto the navigational and research searches proposed by Guha et al [1], with transactional searches consisting of queries 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. However, the range of possible transactions a user may wish to perform, and the underlying reason for wishing to perform them is not explored.

Related work by Rose and Levinson [3] yielded top-level categories with many similarities to those of Broder [2], but also a number of more detailed sub-categories such as *download*, *entertainment*, *interact*, and *obtain*. Despite a number of examples being given to illustrate these sub-categories, the distinctions between them are often based on technical aspects of how the target object will be used, rather than the 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). The authors give 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 it appears the user is trying to locate something that they can then make use of irrespective of how this is done.

One common factor in these studies is a search- or browse-centric perspective on web use. These "two predominant interface modes" [4] (pp 177-178) are often taken as the window through which to study user actions on the web. However, such a perspective may prevent a real understanding of the user's goals in being online. In the scenarios described above, the users have very clear tasks they wish to perform. To what extent can the classifications outlined here accommodate these tasks?

In the *locating a book* scenario, Juan is trying to locate a book that he knows exists and he uses the search engine on the bookshop web site to try and do so. This could be seen as analogous to the navigational searches proposed by Guha et al [1] where the user tries to locate a known document, or by Broder [2] where the user is searching for the web site of a known organisation or individual. In this case the target is a book, but the principle of trying to locate a known item is the same and this task seems fairly well accounted for by the classifications described above. However in the case of Broder [2], consideration is not given to the reason why the user wishes to locate a particular web site. Presumably visiting the site is not an end in itself, but part of the strategy for performing another task such as finding a phone number or arranging car rental.

The focus on classifying search behaviours means none of the schemes discussed so far can account for the task Matt carries out in the *arranging a trip* scenario. Whilst the *resource-interact* goal of Rose and Levinson [3] and the *transactional* queries of Broder [2] 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 arranging something as an overarching reason for being online, or even carrying out a search. Whilst no queries such as "arrange trip to conference" (the task Matt is performing) are reported, this likely reflects that users are aware of the limitations of search engines and therefore do not enter such queries, rather than a lack of desire to perform such tasks.

3.2 Distinguishing Between Needs and Strategies

Drawing on work in domains such as organisation science Choo, Detlor and Turnbull [5] highlight a distinction between a user's *information needs* and the *information seeking strategies* they employ to meet these needs. A similar distinction could also be made between the task a user intends to carry out online, and the strategies they use to complete this task.

Morrison, Pirolli and Card [6] describe a taxonomy of web activites with three variables: the *purpose* of a search, the *method* used, and the *content* of the information being searched for. Whilst these variables appear neatly defined, the classification of some activities suggests the variables may not be mutually exclusive in the form proposed by the authors. For example, some methods are seen to be triggered by a particular goal (*find*, *collect*) whereas others are not (*explore*, *monitor*). In this case it may be that explore and monitor actually represent goals in their own right, and should be classed under *purpose*.

Sellen, Murphy and Shaw [7] describe a classification that identifies six activities carried out on the web (*finding*, *information gathering*, *browsing*, *transacting*, *communicating*, *housekeeping*), based on a study of web use by twenty-four knowledge workers. This classification is not limited to describing variations of searching or browsing, and does attempt to capture the user's needs or goals in using the web. However, by focusing purely on web-based tasks (excluding communication by email, for example), the classification does make assumptions about the strategies being used in performing tasks online.

3.3 Summary

The literature outlined above demonstrates that there are many ways to conceptualise the activities people perform on the web. But to what extent do these classifications represent a valid account of users' goals when online? In general, the classifications address just a small selection of the tasks users may wish to perform online, they characterise component parts of much larger tasks which are not identified or accounted for, or draw distinctions between tasks where these may not actually exist. By taking a search-centric view of web usage some classifications also make assumptions about the strategies a user might employ. Even some schemes that attempt to distinguish needs from strategies remain driven by the principle of an *information* need and *information* seeking strategy, rather than a *task* need and a strategy for performing it.

These factors suggest that a fuller understanding of the range and nature of tasks performed online is necessary. In contrast to current classifications, any broader conceptualisation must adequately account for the scenarios given at the start of this paper, and must not assume the use of specific technology such as search engines or web browsers. In fact, rather than focusing solely on the web as the medium, the only assumption made should be of the user performing tasks using an internet connected device. Distinguishing the web from the rest of the internet in the case of task performance would be to confuse the task need with the strategy employed.

4 A User-Oriented Classification of Online Tasks

Drawing on the schemes described above and the discussion of their limitations, the following classification is proposed as a model of tasks users perform online.

Table 1. a user-oriented classification of online tasks

Task	Definition	Example
Locating	Looking for an object or chunk of information which is known or expected to exist; it may or may not have been seen before by the user.	Locating an article from a journal, an image for a school project, 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.	Exploring a philosophical theory to understand its central tenets; getting background information about an organisation before a job interview.
Monitoring	Checking known sources that are expected to change, with the express intention of detecting the occurrence and nature of changes.	Monitoring news web sites during an election; checking email accounts for new messages; watching discussion fora for new ideas or information.
Grazing	Moving speculatively between sources with no specific goal in mind, but an expectation that items of interest may be encountered.	Following links that spark your interest on someone's web log, just to see what you find.
Sharing	Making an object or chunk of information available to others.	Sharing holiday photos through an online photo album; uploading a journal article to your personal web site.
Notifying	Informing others of an event in time or a change of state.	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.	Writing on your web site that you like a certain film or artist, or that you own a certain book.
Discussing	Exchanging knowledge and opinions with others, on a specific topic.	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.	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	Arranging travel and accommodation for an international conference.

	certain time.	
Transacting	Transferring money or credit between two locations; may or may not have some consequence in the offline world.	Paying a bill, or transferring money between accounts.

Relating this classification to the work of others, the *informational* goal of Rose and Levinson [3] maps clearly to the task of *exploring* described in Table 1, whilst the *resource* goal relates closely to the *locating* task introduced above. However, the *navigational* goal of Rose and Levinson [3] has no equivalent here as it is concerned merely with getting to a specific web site the user has in mind; it doesn't address the task the user intends to perform when they reach the site in question. The same criticism applies to the taxonomy of web searches developed by Broder [2], where in both the *navigational* and *transactional* types the user is attempting to reach web sites where they can perform their task. Considering the taxonomy of Morrison et al [6] raises an issue mentioned previously, that *explore* and *monitor* as they characterise it may actually represent tasks not methods. If this is the case then they correspond well to the tasks of *exploring* and *monitoring* introduced here in Table 1.

Several of the activities identified by Sellen et al [7] have direct equivalents in this classification. For example, their activity of *finding* maps directly to that of the *locating* task presented here, whilst both classifications define *transacting* in similar terms. The *information gathering* activity captures aspects of both the *exploring* and *evaluating* tasks introduced in this paper. Similarly their concept of *browsing* encompasses elements of *monitoring* and *grazing*, without distinguishing the two as this classification does. Whilst the similarities between tasks such as *locating* and *finding* provide a degree of validation for this classification, these examples also highlight the greater granularity of the tasks introduced in this paper.

The classification presented here addresses a wider range of user tasks than those described in Section 3 above. One reason for this greater coverage is that it explicitly includes tasks such as *notifying* and *sharing* that assume an audience or recipient other than the user. Secondly, this classification doesn't make assumptions about the technology being used in performing the task, only that the user is online by way of some form of internet connected device. For example, *notifying* might take place via email, and *discussing* could take the form of an instant messaging conversation. This serves to not limit the classification to a specific domain such as searching using a conventional web search engine, or a specific internet medium such as the web.

4.1 Linked Tasks

During any one online session, a user may perform a number of tasks that, whilst distinct, are in some way related; these could be thought of as linked tasks. For example, you may have heard that a concert is on in the city where you live. You would like to go to the concert, and so use a listings web site to find out that it starts at 8pm. Thinking that your friends might like to go as well, you then email them to let them know about the concert, mentioning the start time. In this case the first task is

clearly an example of *locating*, as you set out to find a certain piece of information, whilst the second task constitutes *notifying*. Here the two tasks bear a thematic relationship but remain tasks in their own right, each addressing a particular goal. Similarly, *monitoring* a news web site may reveal a story of interest that results in the user *grazing* related sites with the expectation of finding other relevant items. Shopping online can be seen as a further example of linked tasks. The act of paying for goods or services can be classified as *transacting*, and this may be preceded by *locating* a specific item to purchase or *evaluating* a number of different options.

4.2 The Role of User Contexts

As the scenarios introduced earlier demonstrate, users rarely perform tasks in isolation. Taking the *arranging a trip* scenario as an example; without the knowledge gained from those around him Matt would likely have booked the more expensive flight to Ljubljana. He may also have begun a long and detailed search for alternative hotels within reach of the conference venue when he found that all official hotels were full. Similarly Alice's knowledge is crucial in helping Juan *locate* a book in the first scenario, both in recommending the book initially based on her own previous experience of reading it, and in clarifying the title.

In fact, a number of aspects of a user's context can be identified that may have significant roles to play in shaping the nature of the task and the way in which it's performed. These might include factors such as a user's *social networks*, their *previous experiences*, *preferences* they hold, their *current location*, services or third parties they *trust*, or the *resources* they have available for performing the task.

Crucially these context factors are likely to manifest themselves differently depending on the task being performed. For example, in tasks such as *notifying* or *sharing*, members of a user's *social network* may be seen as the audience for the task or the beneficiaries of its outcome, rather than sources of assistance as in the scenarios above; *discussing* on the other hand might involve contribution from all individuals, presumably for mutual benefit. Taking the factor of *trust* as an example, the extent to which a user *trusts* a third party web site may be of great significance if they are carrying out a *transacting* task such as paying for goods or transferring money. However, in contrast, if they are *exploring* a controversial topic and simply want to survey a broad range of opinions it may not matter whether they trust the sources they find or not.

5 Tool Support for Online Tasks

5.1 Conventional Internet Tools

If the classification presented in Table 1 represents the tasks people perform online, how are these tasks supported by current tools available on the web, and the internet as a whole? Some existing tools address the needs of these tasks fairly well. For example, software that reads news feeds from multiple web sites and aggregates the

results on a user's desktop are a successful and widely used means of *monitoring* many sources at once. Unfortunately a similar level of uptake has not been seen with tools that monitor multiple email accounts, perhaps due to a lack of standardised ways of accessing web-based email accounts, and users often have to perform this task manually.

In many circumstances traditional search engines are an effective means of locating objects or information, although the *locating a book* scenario illustrates the type of situation where this is not the case. Furthermore, searches are largely limited to textual content due to the complexity of indexing other media such as images or music.

A number of question answering engines such as Ask Jeeves[1] are available that may be able to help *evaluate* if a certain piece of information is true, although the user may not be sure whether to trust the source of the answer. Furthermore, many comparison web sites exist that are able to evaluate the cheapest place to buy a product, or the fastest route between two points, but they are only able to use information explicitly represented in their databases, rather than reasoning about alternatives that may meet the user's criteria. This is highlighted in the second scenario, where the travel web site Matt uses is only able to provide information about routes he specifies, rather than reason about alternative ways of reaching the same destination.

As these examples and the problem analysis given above demonstrate, there is a need for tools and technologies that better support the user in performing tasks online.

5.2 Applications of Semantic Web Technology

A number of tools are discussed below that go some way to addressing these shortcomings, and move towards greater support of the kinds of tasks identified in Table 1. Whilst their features may be described in different terms by their authors, these tools can all be seen to support aspects of the *exploring* task introduced above. To varying degrees they all draw on Semantic Web technologies or principles to support their additional functionality. The Semantic Web vision [8] [9] proposes an extension of the current web that takes it from a collection of interlinked documents for human consumption to a space where information is sufficiently structured, and the rules that define this structure sufficiently explicit, as to allow machines to understand and reason with it. Fundamental to this vision are the basic building blocks of knowledge represented using the Resource Description Framework (RDF), and rules for logical inference in the form of ontologies.

Guha et al [1] describe a system known as TAP, which seeks to support what they term *research searches*. By using Semantic Web data describing concepts and their relationships to others entities, the system is able to provide search results tailored to the concept being searched for. This principle is illustrated with the example of a search for the musician Yo-Yo Ma that returns "his current concert schedule, his music albums, his image, etc." (pp. 702). If however the search term denoted a researcher rather than a musician, the system might return information about the

[1] http://www.ask.com/

person's publications or their research interests. In terms of the *exploring* task, this approach may help the user by providing links to background information not easily assembled using conventional web search engines.

Also supporting users in exploring concepts or entities is the browsing tool Magpie [10]. In contrast to TAP this tool assumes that the user has been able to reach a document that contains some concepts or entities of interest. A user-selected ontological layer over the original document then allows the invocation of semantic services related to those concepts. This serves the purpose of providing related information that may not be explicitly mentioned on the page being viewed.

Another tool that builds on the browser metaphor and applies it to the Semantic Web is Haystack [11]. Here the user is able to browse arbitrary collections of RDF metadata through a point and click interface, with links being made between semantically related items. Crucially this tool is able to gather information on a particular topic from multiple sources and assemble it in one place, in contrast to conventional models of web browsing where the user may have to visit several different pages or sites to gather related pieces of information.

Whilst implemented differently (on the web rather than on the desktop), the application CS AKTive Space [12] provides a similar ability to explore relations between concepts or entities, although in this case the system is limited to the domain of computer science research in the UK.

One feature these tools have in common is the ability to present the user with new pieces of information, or make new connections between concepts or entities that might not otherwise have been apparent; this ability is a key feature of the Semantic Web. To this end they make a significant step towards supporting users in the task of *exploring* concepts or entities to gain additional knowledge or understanding.

However, many of the other tasks identified earlier are not so well supported. For example, resolving the issues highlighted by the scenarios presented in Section 1 requires tools adapted to *locating* and *arranging* that go beyond the traditional search engines and travel web sites currently available. Semantic Web technologies such as those outlined in [9], may provide the technological basis for building such tools, by enabling the creation of large, distributed, and dynamic knowledge bases, and the means to reason across them.

In the *locating a book* scenario, this might enable Juan to specify that the book he is looking for is called something like "The Sergeant's Guitar". A system that could make use of background knowledge about the semantic links between 'sergeant' and 'captain', and between 'guitar and mandolin', might be able to identify "Captain Corelli's Mandolin" as one possible match within the online bookshop, rather than returning no results. Similarly when *arranging a trip*, a Semantic Web application could take Matt's destination as an input, reason about ways of reaching that destination and propose a number of travel itineraries, leaving Matt to choose the one that best meets his needs. In both these cases, tools that draw on aspects of Juan and Matt's contexts, particularly knowledge held by those around them, would be beneficial in completing the tasks.

Of the tools discussed above, perhaps the only one to take any account of user context is Magpie. The user's selection of an ontological layer could be seen to reflect some aspect of their context, in that subscription to a shared conceptualisation likely reflects their perspective on a domain to some extent. However, this representation of

context is implicit and does not approach the richness of the factors proposed in this paper.

6 Conclusions: From a Semantic Web to a Task-Focused Context-Aware Internet

In conclusion, unless tools are developed that are adapted to the task the user wishes to perform, and that take into account the contexts in which the user exists, the kind of obstacles highlighted in the scenarios above are likely to remain. Task-focused and context-aware tools could provide a more effective means for users to perform tasks online than current web tools, and Semantic Web technologies may provide the platform for developing them, if the following challenges can be met. Firstly, can a user's contextual data be captured and made available on the Semantic Web in a meaningful and reusable form, and how might this be achieved? Secondly, can tools be developed that are able to reason about the contextual information needed to assist in the performance of a particular task?

Such tools should be extended to cover the full spectrum of tasks users perform online, and also operate across a wider range of internet platforms such as email and instant messaging. Not to do so would draw a distinction between the web and the wider internet based on technical grounds such as the particular protocol being used. Distinctions of this sort hold little meaning for the average user, who is concerned primarily with performing a task irrespective of how tools are implemented. To this end it may be more appropriate to envision a task-focused context-aware internet, where all online activities can benefit from the use of semantic technologies.

Acknowledgements

This research was partially supported by the Advanced Knowledge Technologies (AKT) project. AKT is an Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University.

References

1. Guha, R., McCool, R., Miller, E.: Semantic Search. In: Proc. Twelfth International Conference on World Wide Web (WWW2003) (2003) 700-709
2. Broder, A.: A Taxonomy of Web Search. ACM SIGIR Forum 36 (2002) 3-10
3. Rose, D. E., Levinson, D.: Understanding User Goals in Web Search. In: Proc. 13th International Conference on World Wide Web (WWW2004) (2004) 13-19
4. Olston, C., Chi, E. H.: ScentTrails: Integrating Browsing and Searching on the Web. ACM Transactions on Computer-Human Interaction (TOCHI) 10 (2003) 177-197

5. Choo, C. W., Detlor, B., Turnbull, D.: Information Seeking on the Web: An Integrated Model of Browsing and Searching. In: Proc. 62nd Annual Meeting of the American Society for Information Science (1999) 3-16
6. Morrison, J. B., Pirolli, P., Card, S. K.: A Taxonomic Analysis of What World Wide Web Activities Significantly Impact People's Decisions and Actions. In: Proc. Conference on Human Factors in Computing Systems (CHI'01) (2001) 163-164
7. Sellen, A. J., Murphy, R., Shaw, K. L.: How Knowledge Workers Use the Web. In: Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI2002) (2002) 227-234
8. Berners-Lee, T.: Semantic Web Road Map. World Wide Web Consortium (W3C) (1998)
9. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. Scientific American 284 (2001) 34-43
10. Dzbor, M., Domingue, J., Motta, E.: Magpie - Towards a Semantic Web Browser. In: Proc. 2nd International Semantic Web Conference (2003)
11. Quan, D., Karger, D. R.: How to Make a Semantic Web Browser. In: Proc. 13th International Conference on World Wide Web (WWW2004) (2004) 255-265
12. Schraefel, M. C., Shadbolt, N. R., Gibbins, N., Harris, S., Glaser, H.: CS AKTive Space: Representing Computer Science in the Semantic Web. In: Proc. 13th International Conference on World Wide Web (2004) 384-392