Usability and the Semantic Web

Closing Keynote at the
Third European Semantic Web Conference
Budva, Montenegro
June 14th, 2006

Anthony Jameson
DFKI, German Research Center for Artificial Intelligence
and
International University in Germany

Note on additional text

- Text that appears in this smaller font – or in marginal notes – was not shown during the presentation at ESWC 2006
- This text has been added so as to enable readers to understand the slides independently of the live presentation

Contents

Introduction
Summary

Introduction

Many innovative technologies aim to enhance a user's interaction with a system is some respects; but they typically raise usability challenges which, if not dealt with, may outweigh the intended benefits

Research on semantic web technologies has so far focused mostly on the technology, but the past 3 years have seen an increase in interest in interaction design and evaluation

The main part of this talk discusses three of the key usability challenges, approaches that have been taken to them, and issues that remain open

Three general challenges

1. Reducing effort and complexity in querying and search

In the ideal case, users could obtain information via semantic methods by straightforwardly characterizing their information need in terms of elements of the ontology(ies) used in the system

In most cases, the ontologies (and other information sources) are too large, complex, and otherwise unsuited for end–user inspection

Designers of query interfaces for the semantic web have been creative and often successful in devising ways of allowing users to benefit from the existence of an ontology without confronting them with its complexity

A goal related to that of reducing effort is the goal of ensuring adequate *expected benefit*, which can be relatively difficult with semantically based interfaces; two strategies are discussed briefly

2. Conveying adequate mental models

The same design solutions that reduce effort and complexity in querying can also make the resulting behavior of the system difficult to understand and predict, as is illustrated here by a discussion of an intermediate SmartWeb prototype

Research and experience in human–computer interaction on *mental models* yields a number of results and ideas about when and why it is important for a user to have at least some vague understanding of how a system works internally and about ways of conveying an appropriate mental model

3. Providing adequate motivation for content provision

Many semantic web application scenarios presuppose that some users will invest effort in providing or enhancing content (e.g., by annotating web pages)

Theory and experience from several fields have yielded a number of ideas about the conditions under which users may be motivated to do such work

The roles of users in semantic web research and development

In more mature fields that involve novel forms of human–computer interaction, it is often hard to publish a paper concerning a new interactive system unless it includes some empirical evidence that the novel aspects of the system are well accepted by users

By contrast, empirical research with users is found only sporadically (though increasingly) in the semantic web field

Several apparent reasons are discussed

Finally, some general hints about how to involve users effectively in research on semantic web technologies are given, with emphasis on the diversity of roles that users can play and the proven effectiveness of interdisciplinary teams in designing useful and usable systems

Two Perspectives

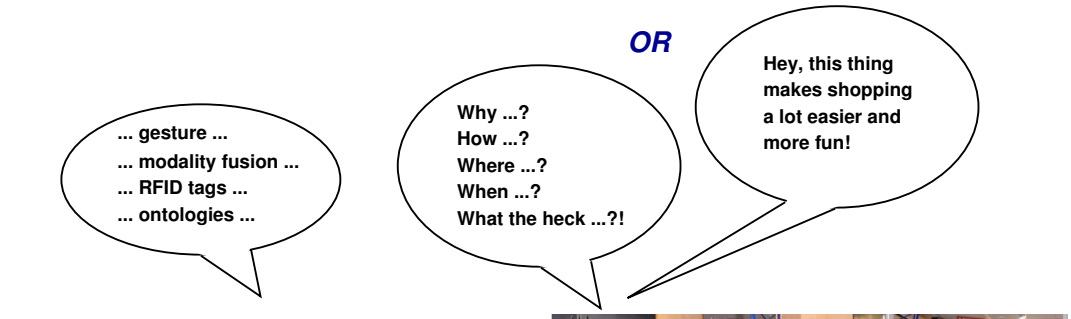

Usability of Novel Technologies

Innovative Technology	Usability Challenges
Spoken dialog systems	Dealing with speech recognition errors Knowing what to say
Systems that adapt to their users	Understanding and controlling system behavior Putting up with additional demands on attention
Systems with semantic web technology	?

Hasn't This Been Done Before?

Workshops

- WWW 2004
- ESWC 2005
- ISWC 2005 (and 2006)

Other individual papers

- [Too many to mention]

My goal today

- Highlight *a few key themes*
- Discuss them with reference to concrete examples from real systems

Tasks and Challenges

Type of task	Issues discussed today
Searching / querying	How can we minimize complexity for the end user?
	How can we ensure the minimally necessary understanding of the system's processing?
Adding information to ontologies	How can we induce users to do the necessary work?
	How can we involve users in the design process?

The Usual Focus

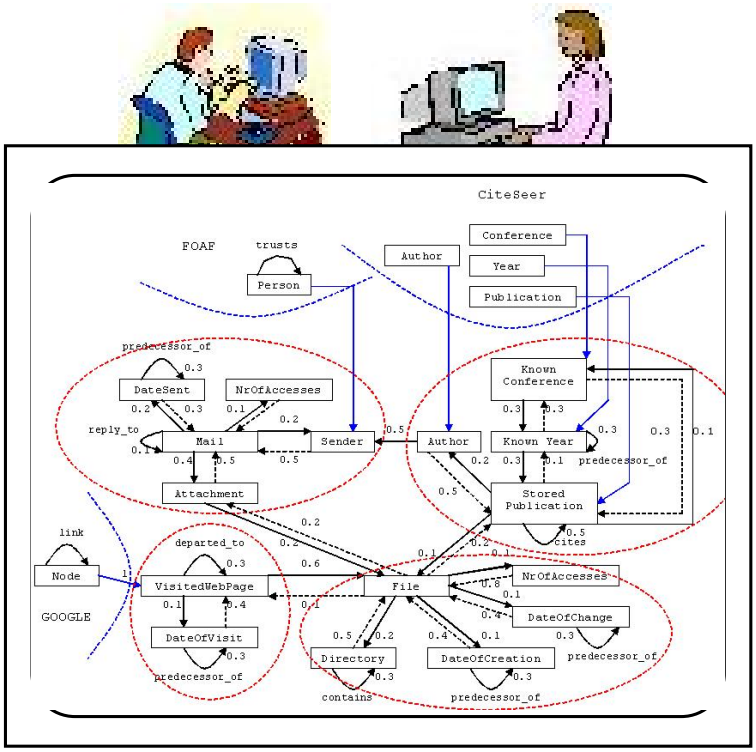

The Focus in This Talk

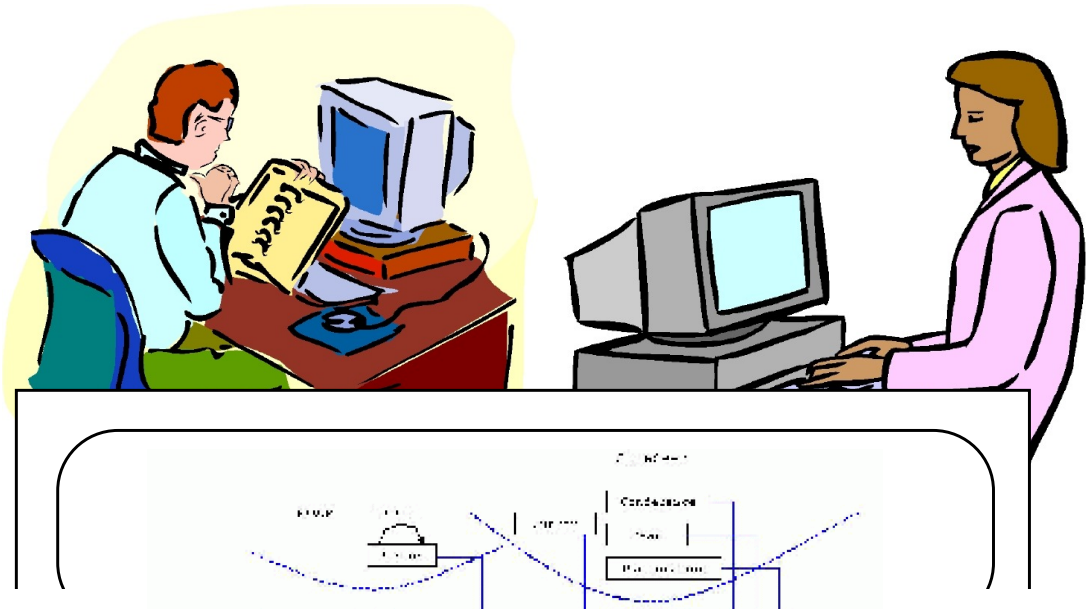

Reducing Effort and Complexity in Querying
A Query in seMouse

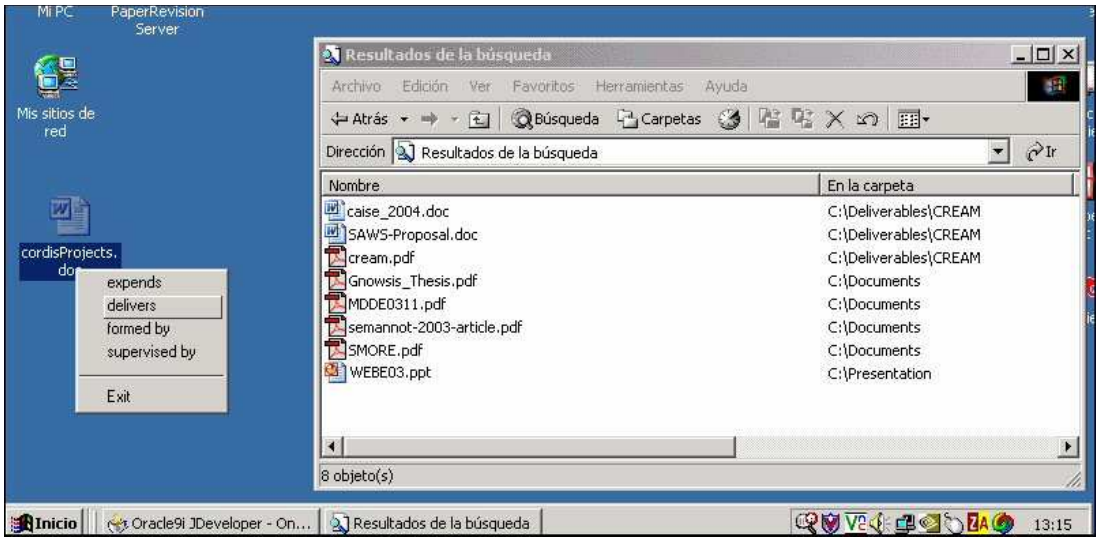

- Querying with seMouse involves straightforward selection from the elements in the desktop ontology
- It is unclear how usable this approach would be if the ontology were orders of magnitude larger

13

Iturrioz, J., Anzuola, S. F., & Diaz, O. (200 urning the mouse into a semantic device: The seMouse experienc Y. Sure and J. Domingue (Eds.), *Proceedings of the Third European Semantic Web Conference*. Berlin: Springer.

Halo 2: Background

Project Halo is a multistaged effort of Vulcan, Inc. whose ultimate goal is a large, widely available knowledge base that can answer scientific questions from various fields, in particular questions that require general reasoning and computational capabilities

Since it is impractical to have knowledge engineers involved at all times in the formulation of such a large body of knowledge, tools are being developed and evaluated in Phase 2 of the project that enable domain experts to formulate knowledge independently after a limited amount of training

The slides and comments about Halo in this talk refer to examples from one of the two teams that are pursuing somewhat different approaches in parallel: the team led by ontoprise, which also includes groups from the Open University, iSoco, DFKI, Carnegie–Mellon University, and Georgia Tech

More detailed reports on the results from this team, as well as the other team – led by SRI and including groups from Boeing, the University of Texas, and Carnegie–Mellon University – will be presented once the evaluations currently in progress have been completed

14

For a report on Phase 1 of the Halo project, see Friedland, N., Allen, P., Matthews, G., Witbrock, M., Baxter, D., Curtis, J., Shepard, B., Miraglia, P., Angele, J., Staab, S., Moench, E., Oppe , H., Wenke, D., Israel, D., Chaudhri, V., Porter, B., Barker, K., Fan, J aw, S., Yeh, P., Tecuci, D., & Clark, P. (2004). Project Halo: Towards a digital Aristotle. *AI Magazine, 25*(4), 29–48.

Halo 2: Knowledge Querying

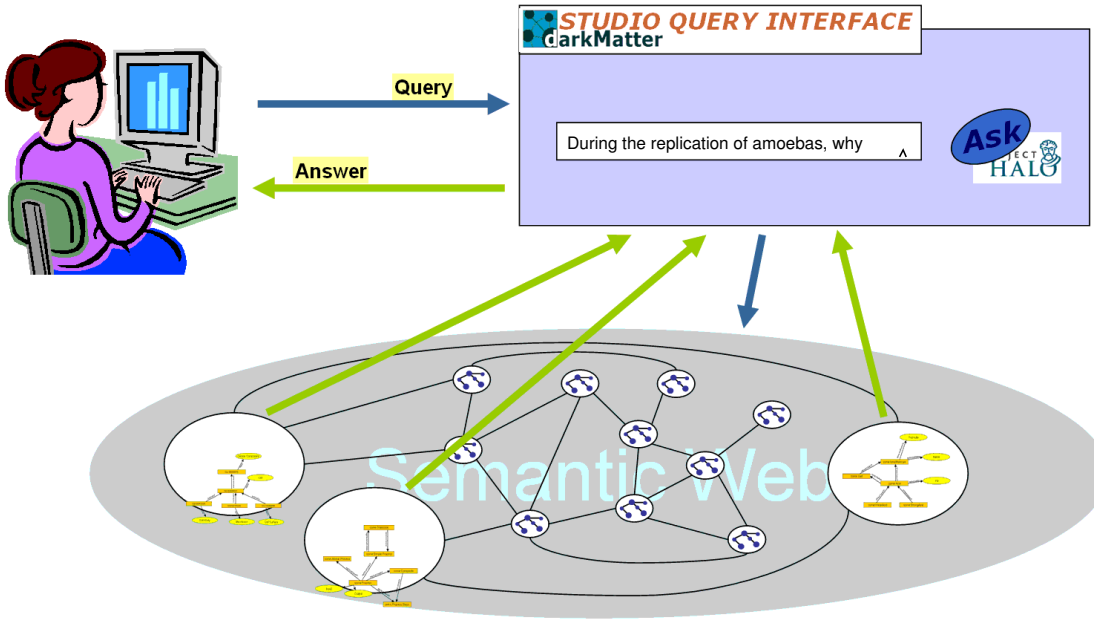

The "Digital Aristotle" vision

15

Halo 2: Knowledge Formulation

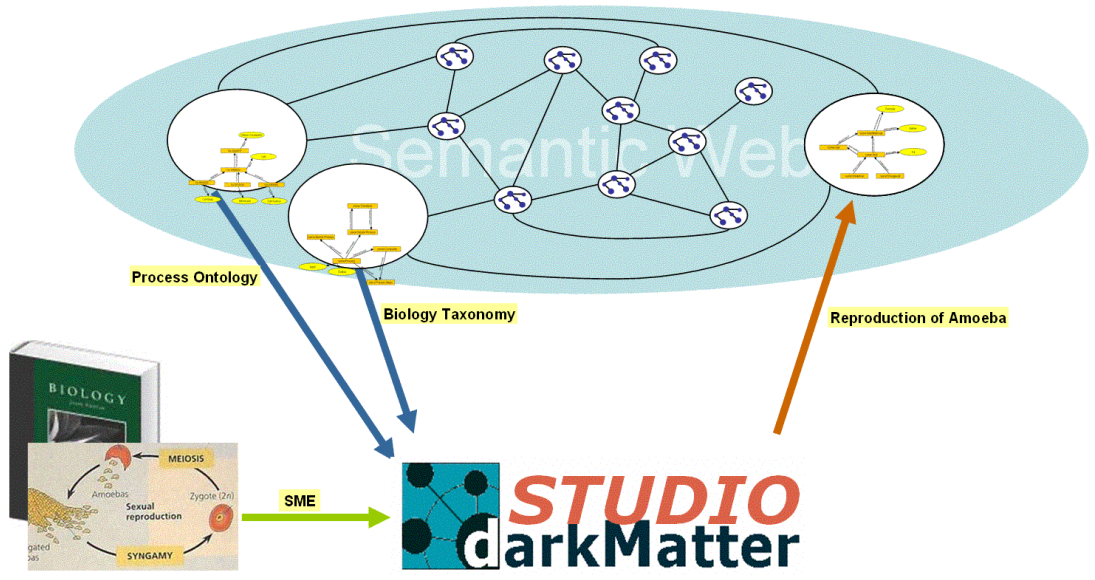

16

A Query in DarkMatter

Original formulation of a physics question

- A car accelerates from 0 km/h to 95 km/h in 6.2 seconds. What is its acceleration?

Formulation constructed in current version of DarkMatter

- A vehicle has a motion of a constant accelerated motion. It has an initial velocity of 0 km/h. The constant accelerated motion has a final velocity of 95 km/h. The constant accelerated motion has an initial time of 0 second. The constant accelerated motion has a final time of 6.2 second. What is the acceleration of the constant accelerated motion? 17

Queries formulated with the WYSIWYM te… (see the next slide) currently require the user in effect to put together el…nts from the ontology created by the knowledge formulator. The addition of more domain-specific patterns is expected to improve the compactness and naturalness of the formulations.

The WYSIWYM Query Formulation Tool

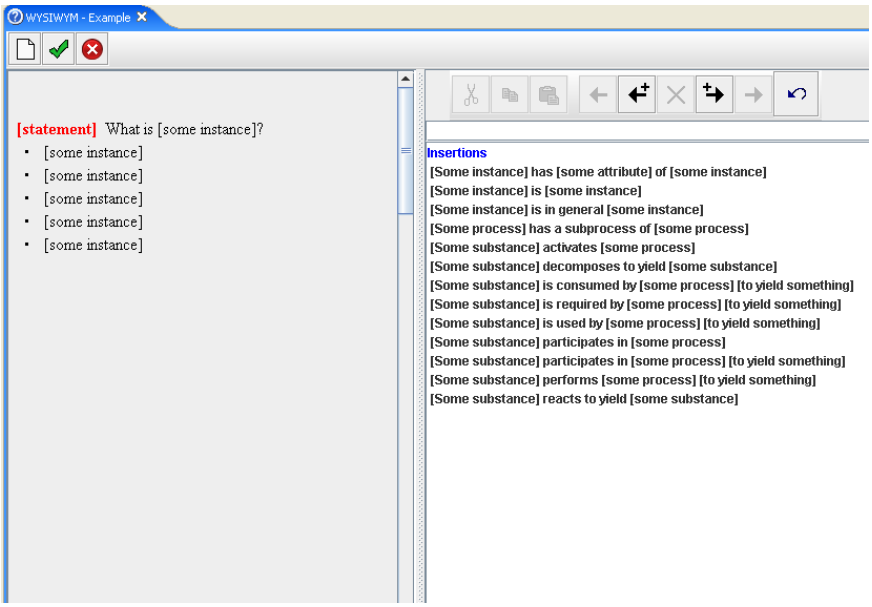

WYSIWYM, developed by Richard Power, colleagues at the University of Brighton and Open University, has been in… ated as part of the DarkMatter system. See *http://www.itri.bton.ac.uk/projects/WYSIWYM/wysiwym.html* for further information

Minimizing Complexity and Cognitive Effort

Questions to be considered in each case

1. How great is the complexity and effort required of the user?
2. How might they be reduced?

Strategies for reducing complexity and effort

1. Require recognition rather than recall
2. Provide familiar, domain–specific interfaces
3. Have the system do the mapping of the input onto the concepts of the formal representation (even if imperfectly)
4. Support trial and error
   - Users generally prefer quick cycles of action and evaluation to careful thought

OntoIR: Query Screen

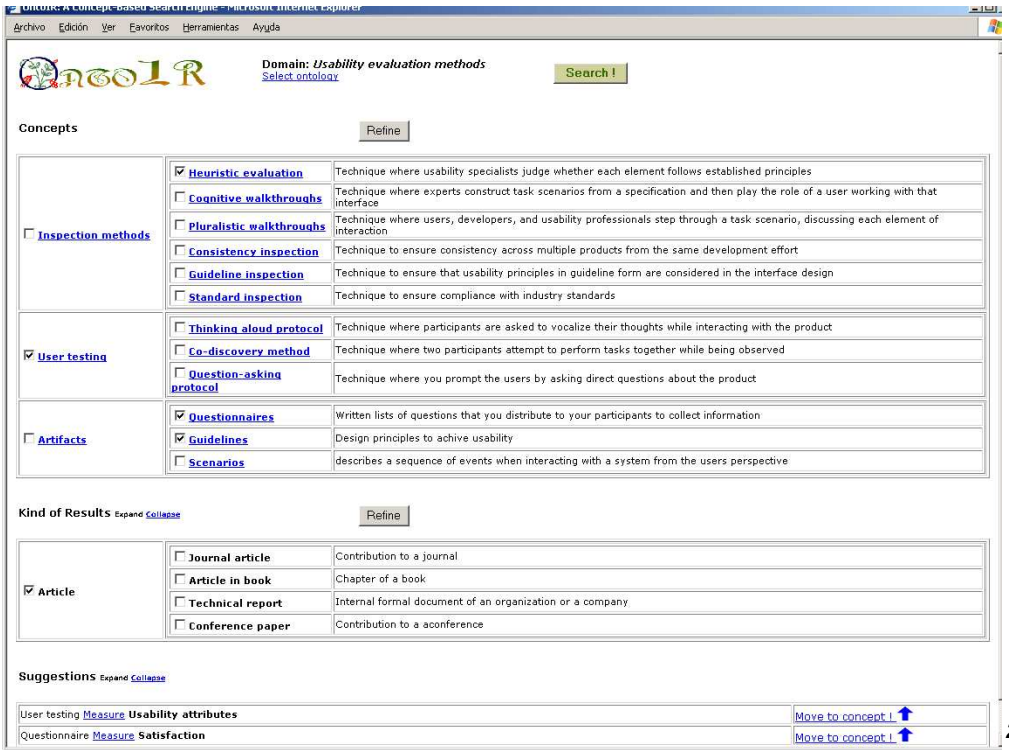

OntoIR presents aspects of a document collection selectively, shielding the user from irrelevant aspects of the ontology (scope, e.g., Gardia, E., & Sicilia, M. (2003). Designing ontology–based interactive information retrieval interfaces. *Proceedings of OTM Workshops*, pp. 152–165.

OntoIR: Search Result Screen

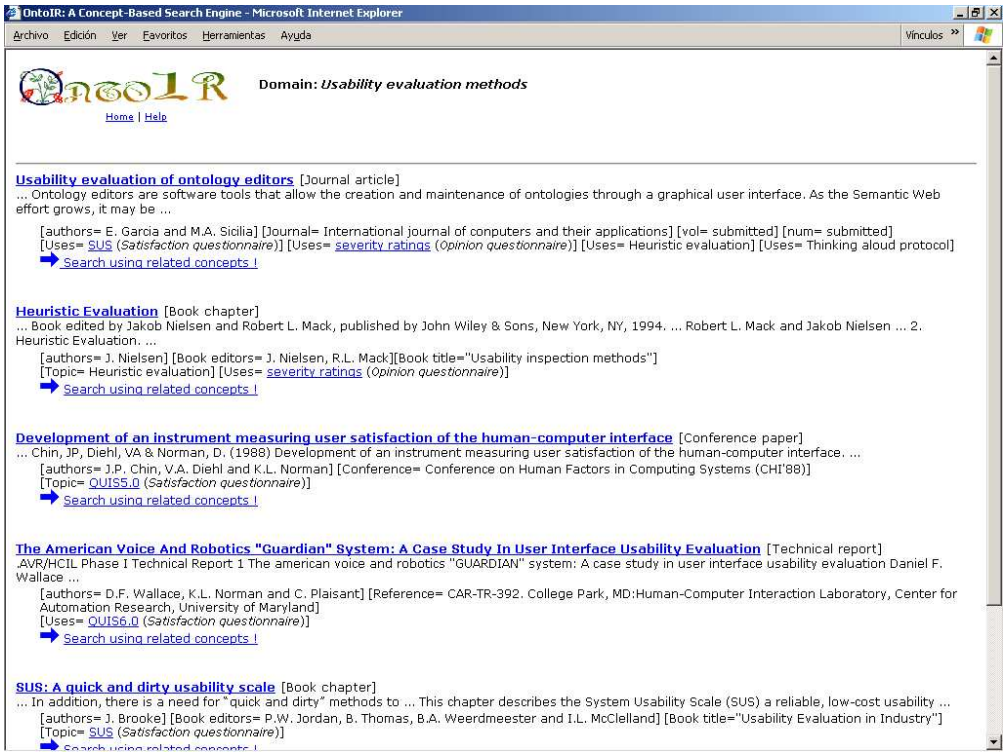

21

Considering Expected Benefits

Problem

- With semantically based systems, often no useful result (or no result at all) is returned
  - The available content is often limited
  - The semantically based retrieval mechanisms may not support the retrieval of partially relevant results

General strategy

- Allow easy refinement of queries and/or further navigation using the returned results as a starting point
- "Piggyback" on methods can be counted on to yield some useful result in almost any case

Search on TAP: Results

Cf. http://tap.stanford.edu/ (go to "Demos" to try the system) and McCool, R., Cowell, A. J., & Thurman, D. A. (2005). End–user Evaluations of Semantic Web Technologies. In A. Bernstein, I. Anagnostopoulos, D. Degler, & B. McBride (Eds.), *Proceedings of the iswc 2005 workshop on end user semantic web interaction*. Aachen, Germany: CEUR Workshop Proceedings, Vol. 174.

- The "Search on TAP" system illustrates how the form in which results are presented can convey picture of the different types of results that are available and how they were derived
- But a small–scale user study revealed that users still had difficulty assessing the range of information available via the system
- This user study also uncovered a bug in the screen design:
  - In the earlier version of the system that was tested, the order of the two columns was the opposite of that shown here
  - As a result, users tended to overlook them, as if they were advertisements

Conveying Adequate Mental Models
Introduction to Mental Model Exercise

The following sequence of slides shows screens from the current demonstration prototype of SmartWeb, illustrating its responses to each of three queries.

Although the basic input modality of SmartWeb is speech, these queries were typed in for the purpose of this demonstration.

Readers are asked to put themselves in the position of a user who has entered these queries, considering the following questions:

1. Are there any major differences in the ways in which these three queries were processed?
2. Can I predict what kind of answer I'm going to get to different types of question?

They can then look at the overview of SmartWeb's actual processing on the succeeding slide and consider how much of the information shown there out to be conveyed to the user

The screen shots were kindly supplied by the authors of: Ankolekar, A., Cimiano, P., Hitzler, P., Kiesel, M., Kroetz, M., Ladwig, G., Lewen, H., Oberle, D., Sintek, M., & Studer, R. (2006). SmartWeb: Multimodal mobile access to the semantic web. *Posters of the Third European Semantic Web Conference*, Budva, Montenegro.

SmartWeb: Stadium Query (1)

"Welcome to SmartWeb. Please ask a question" ... "Data received" [left over from previous query.]

"Where is the Gottlob Daimler Stadium?"

SmartWeb: Stadium Query (2)

"No semantic analysis available"

"Stuttgart" ... "Results received"

SmartWeb: Route Query

"How do I get from Stuttgart to Berlin?" ...

"... Route requested; goal: Berlin, starting point: Stuttgart"

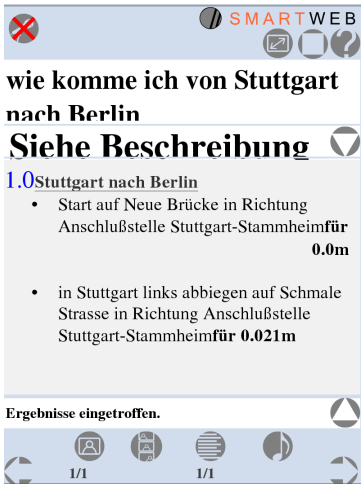

"See description" [shown on main screen]

SmartWeb: Montenegro Query

"Where is Montenegro?"

"No semantic analysis available"

"Albania"

SmartWeb: Processing Overview

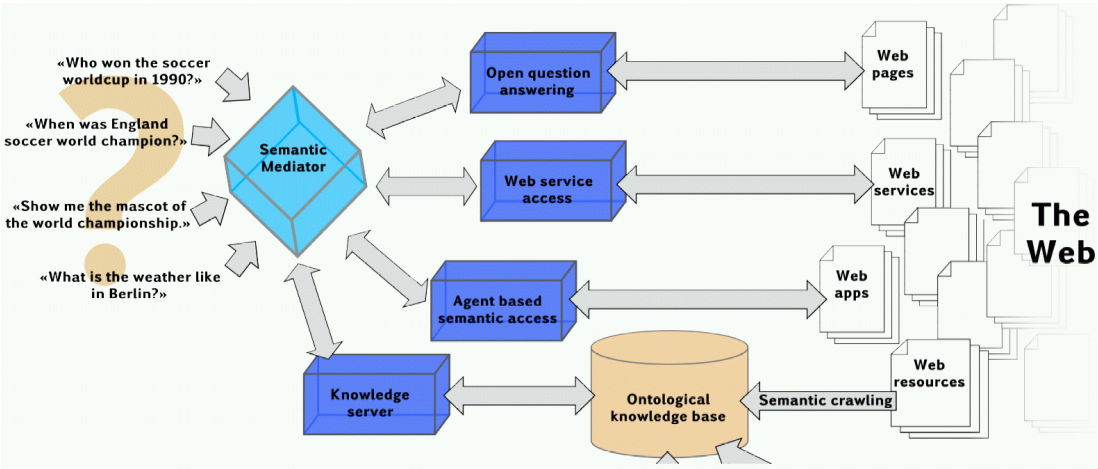

- The structure of SmartWeb is designed to ensure that a broad variety of questions can be answered, including questions of types not specifically anticipated by the system designers
- This underlying complexity yields system behavior that may be hard for users to understand and predict unless they are given some guidance

Basic Idea of Mental Models

- Users sometimes(!) like to have a vague(!) idea of what's going on
  - ... called the *mental model* of the system
  - ... typically nontechnical, incomplete, and changing over time
- When is a mental model needed?
  - Not so much: When everything goes smoothly
  - More:

    When you want to ...

    ... predict what the system will do in a given situation

    ... understand some unexpected system behavior

Conveying Mental Models

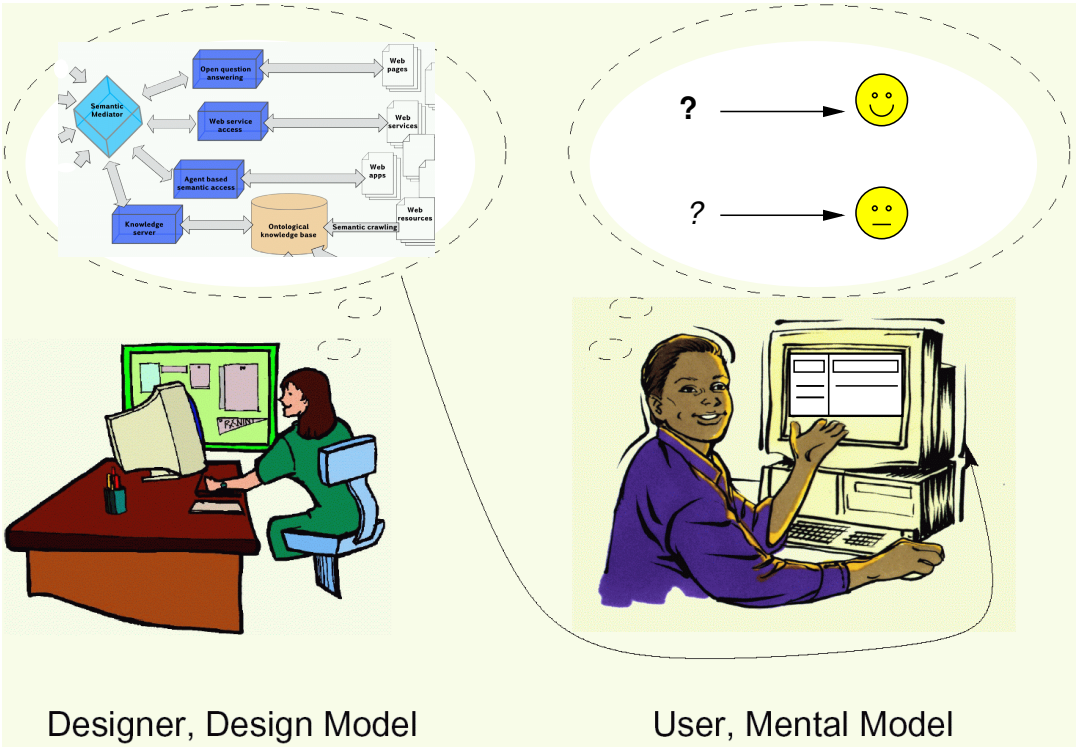

Designer, Design Model User, Mental Model

How Not to Convey a Mental Model

1

- Welcome to SmartSemanticSearch
- Just ask whatever you like
- I will use semantic web technology to understand your question and give you a much better answer than you could get with Google

2

- Welcome to SmartSemanticSearch
- The ontology used as a knowledge base was created on the basis of the *XYZ* and *UVW* ontologies and populated via ...

Ways of Conveying Mental Models

Suggesting what the user can do

- Appearance of interface elements for input
- Examples of possible inputs
- ...

Suggesting what the system has done

- Layout of the presentation of the results
- Indications of information used to derive the responses
- ...

Providing Adequate Motivation
Introduction

A widespread concern in the semantic web community is that some applications presuppose that nontechnical users will do a good deal of work (e.g., semantic annotation) in order to make content available for processing by semantic web techniques

The examples in this section concern two systems (among others) that (a) provide more or less immediate benefit to the user who adds semantic content and (b) also offer delayed benefits to the user and to other persons

The Document Manager from the SemIPort project helps the user to organize scientific articles by placing them in nested groups and by specifying relations among them

The zoomable interface can be seen as a domain–specific visualization of the ontology that internally represents the information supplied by the user

The immediate benefit of the annotations added by the user is the improved access that he or she has to the documents

Less direct benefits include:

1. For the user: the ability, while querying the computer science portal *io–port*, to send the current work context to the portal so that search results can be reordered on the basis of their relevance to the work context
2. For other users who are looking for documents that are related in particular ways to a document that they have found: the use by the central io–port server of all relevant annotations that have been made by users who have uploaded their annotations to the central server

Responses from users indicate that they would not take the trouble to make the annotations only for the sake of the indirect benefits

SemIPort Document Manager (1)

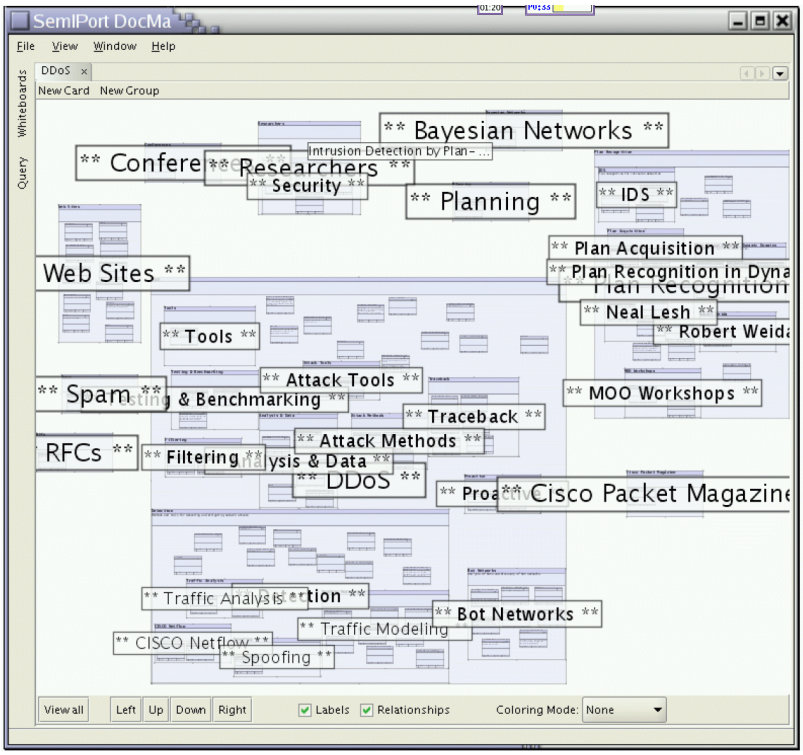

SemIPort Document Manager (2)

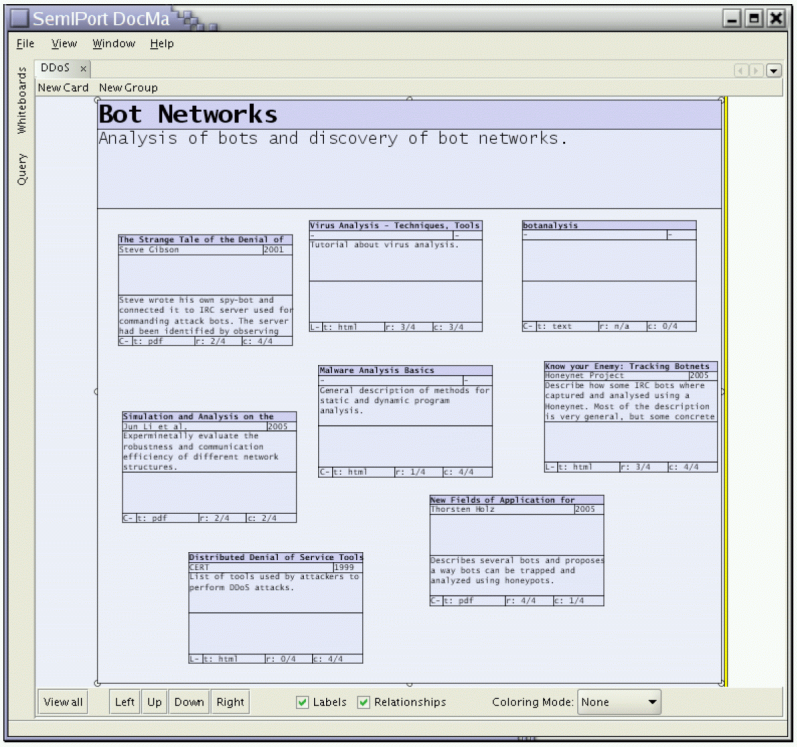

Specifying Relations Between Documents (1)

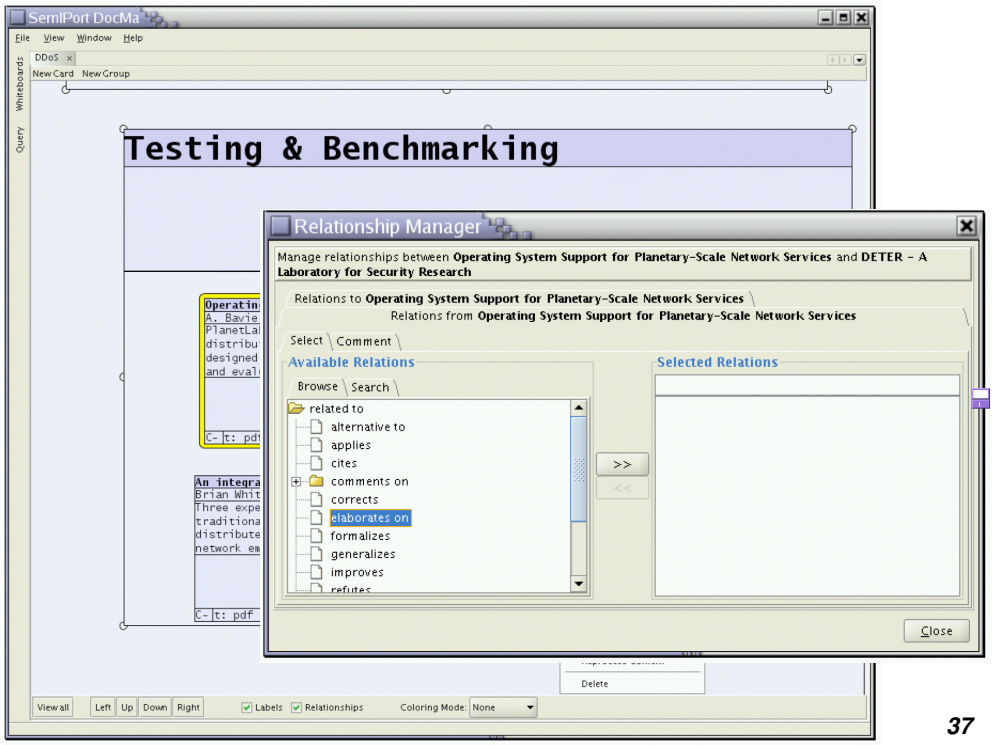

37

Specifying Relations Between Documents (2)

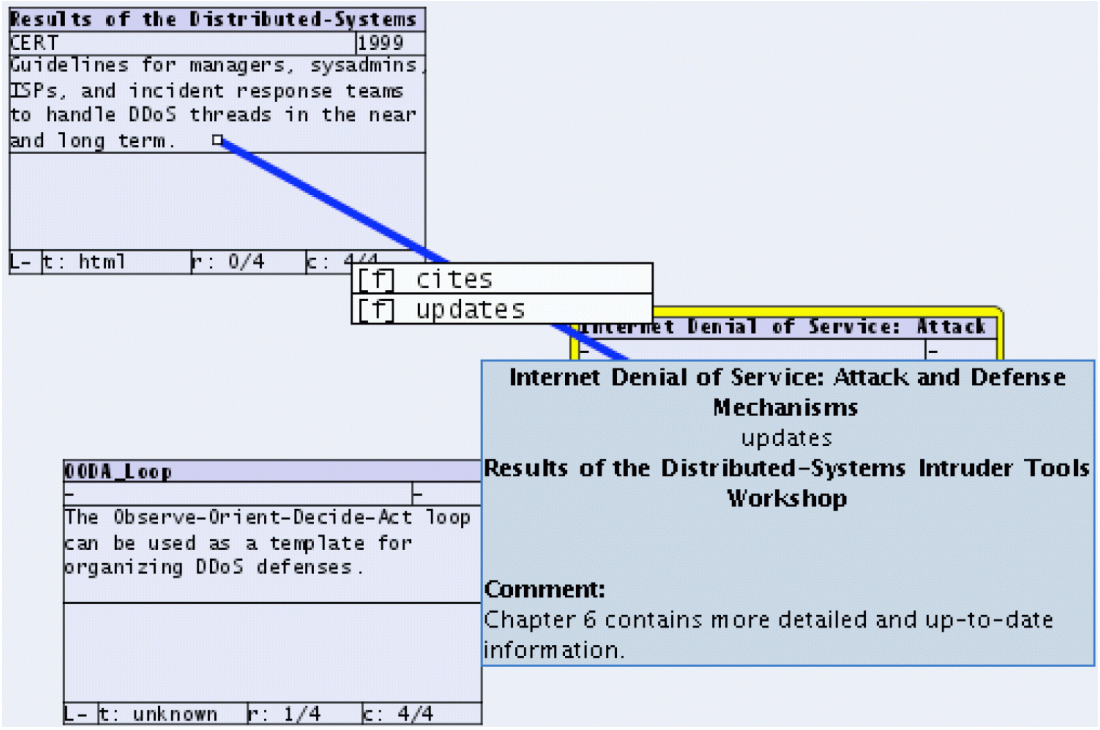

Mangrove: Overview

- The Mangrove annotation tool is likewise based on the idea of offering immediate reward for work that contributes to an evolving knowledge base
- As the second of the following slides shows, annotations that a user makes are processed immediately, and the user is given feedback on the ways in which they can now be used
- The subsequent three slides illustrate three of the services at the University of Washington Computer Science Department that make use of annotations made with Mangrove

Mangrove: Annotation Tool

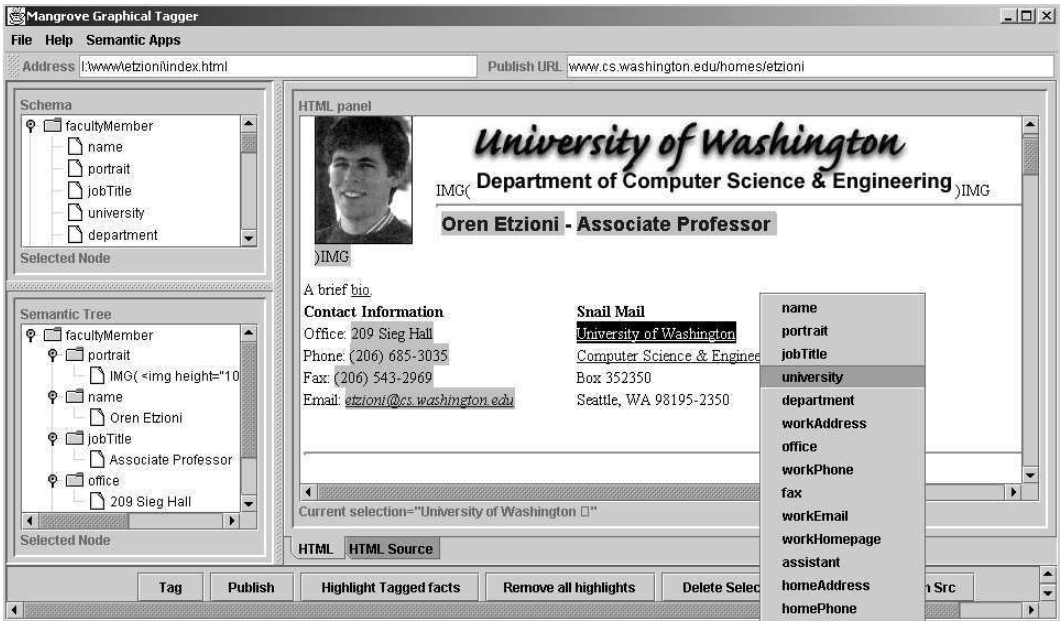

See McDowell, L., Etzioni, O., Gribble, S., Halevy, A., Levy, H., Pentney, W., Verma, D., & Vlasseva, S. (2003). Mangrove: Enticing ordinary people onto the semantic web via instant gratification. *Proceedings of ISWC 2003*, Sanibel Island, Florida, pp. 754–770.

Mangrove: Feedback From Services

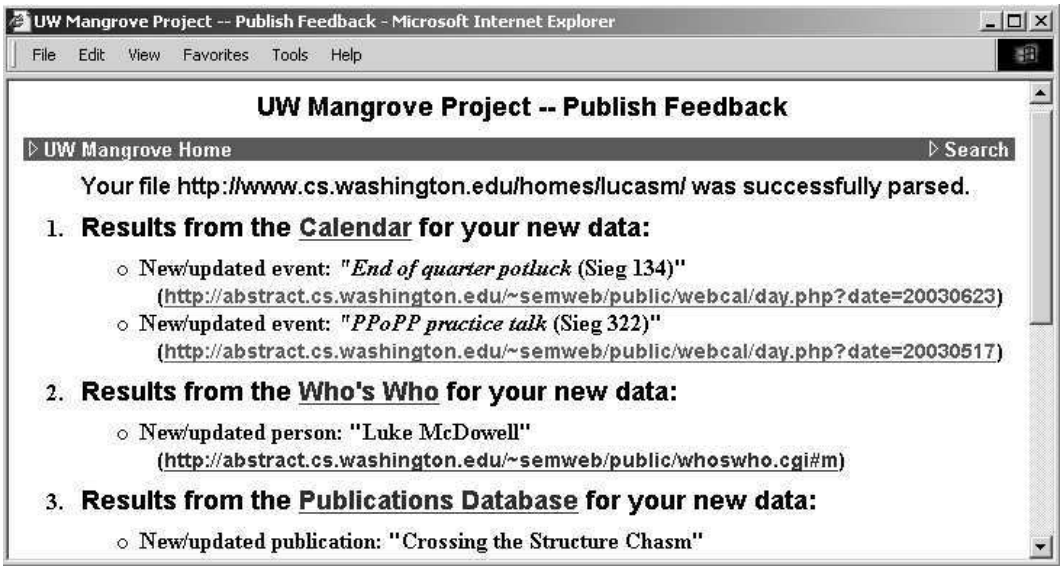

41

Mangrove: Semantic Search Results

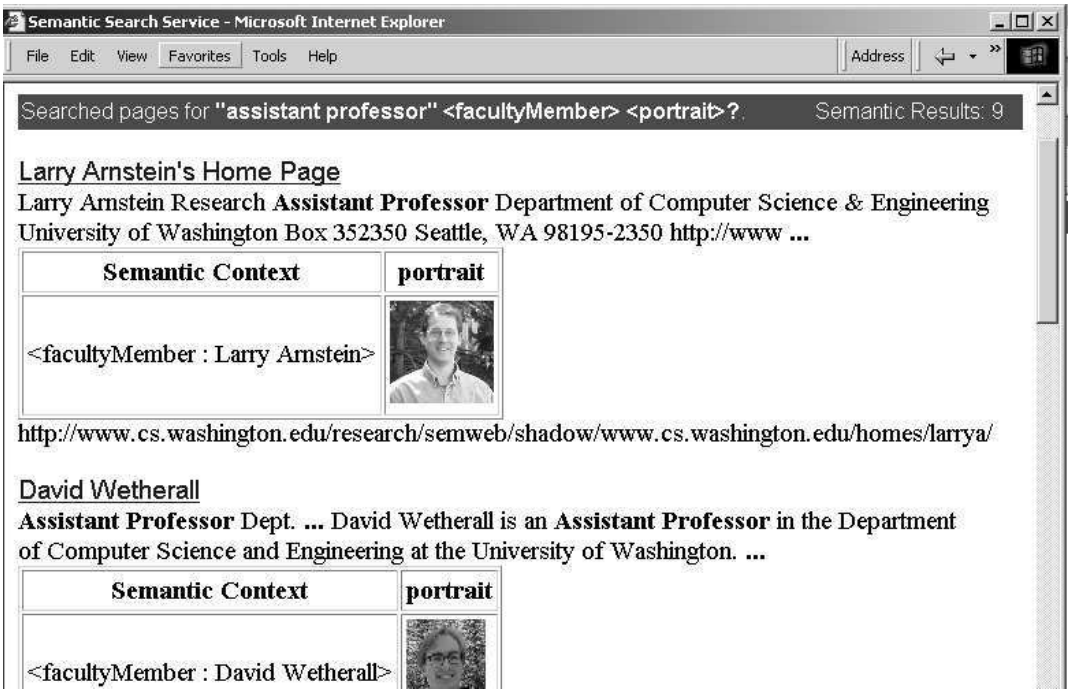

Mangrove: Calendar

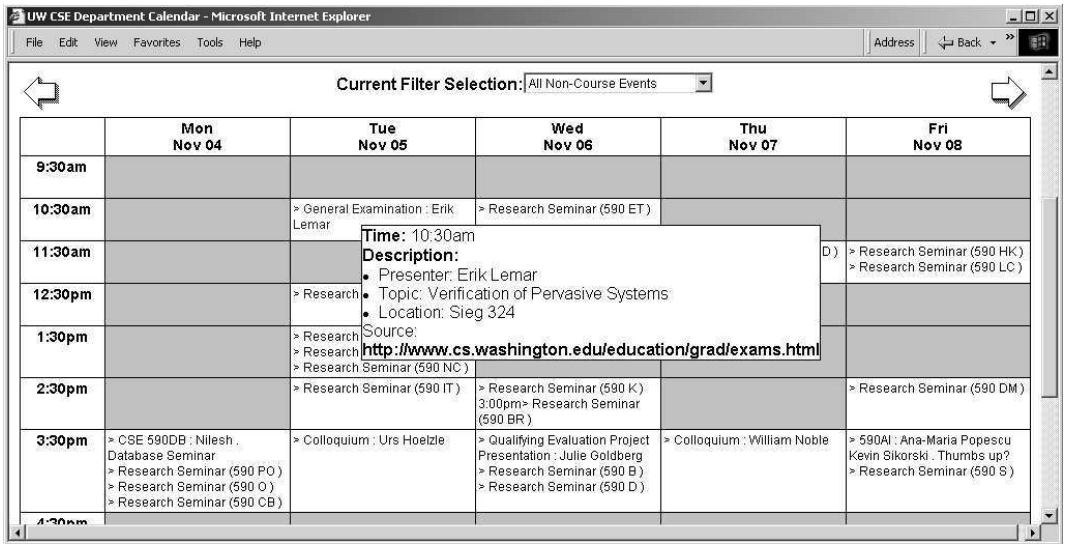

43

Mangrove: Who's Who

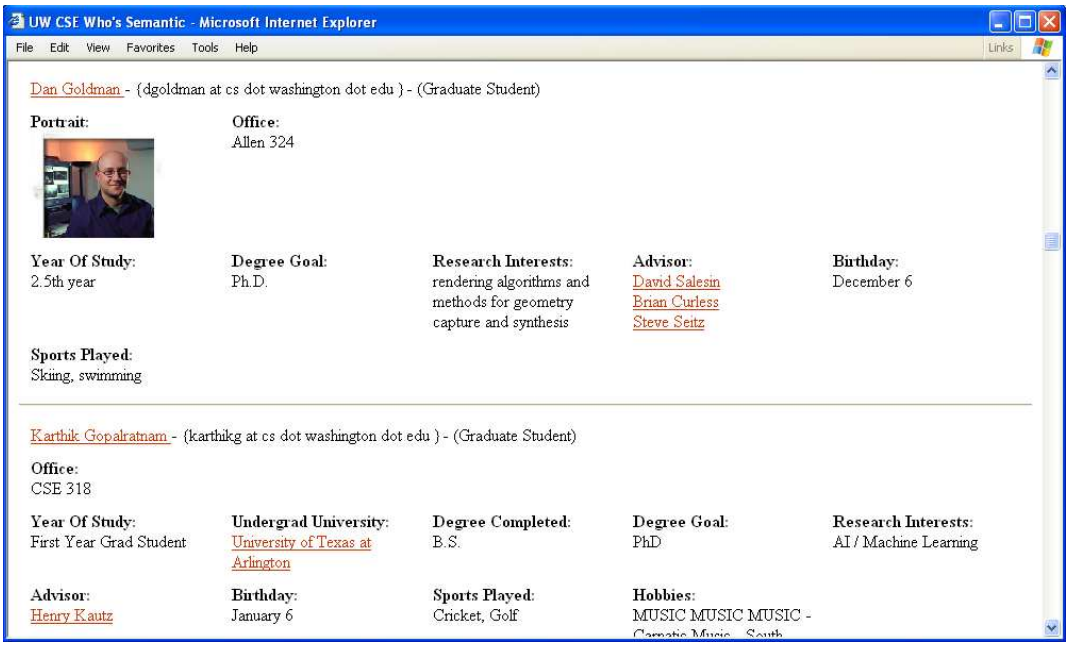

44

Overview of Types of Benefit (1)

- The following slide shows how the benefits that a user derives from work can be seen as being more or less direct along two dimensions
  - Temporal: How far away (and thus how uncertain) the benefits are in time
  - Social: The extent to which the benefits depend on the actions of other persons – or even accrue only to other persons
- One general strategy, illustrated by the two systems just discussed, is to ensure that at least some of the benefits lie within the less problematic areas of the space
- This strategy is also applied in the *Community Navigator* of Takeda and Ohmukai (2005)

Further references

- Beenen, G., Ling, K., Wang, X., Chang, K., Frankowski, D., Resnick, P., & Kraut, R. E. (2004). Using social psychology to motivate contributions to online communities. In J. Herbsleb & G. Olson (Eds.), *Proceedings of the 2004 Conference on Computer–Supported Cooperative Work* (pp. 212–221). New York: ACM.
- Cosley, D., Frankowski, D., Terveen, L., & Riedl, J. (2006). Using intelligent task routing and contribution review to help communities build artifacts of lasting value. *Human Factors in Computing Systems: CHI 2006 Conference Proceedings*.
- Grudin, J. (1994). Groupware and social dynamics: Eight challenges for developers. *Communications of the ACM, 37*(1), 92–105.
- McDowell, L. (2004). *Meaning for the masses: Theory and applications for semantic web and semantic email systems.* Dissertation, Department of Computer Science and Engineering, University of Washington.

Takeda, H., & Ohmukai, I. (2005). Building semantic web applications as information/knowledge sharing systems. In [illegible] Dzbor (Ed.), *PROCEEDINGS of the workshop on end users aspects of the semantic web*. Heraklion, Greece.

Overview of Types of Benefit (2)

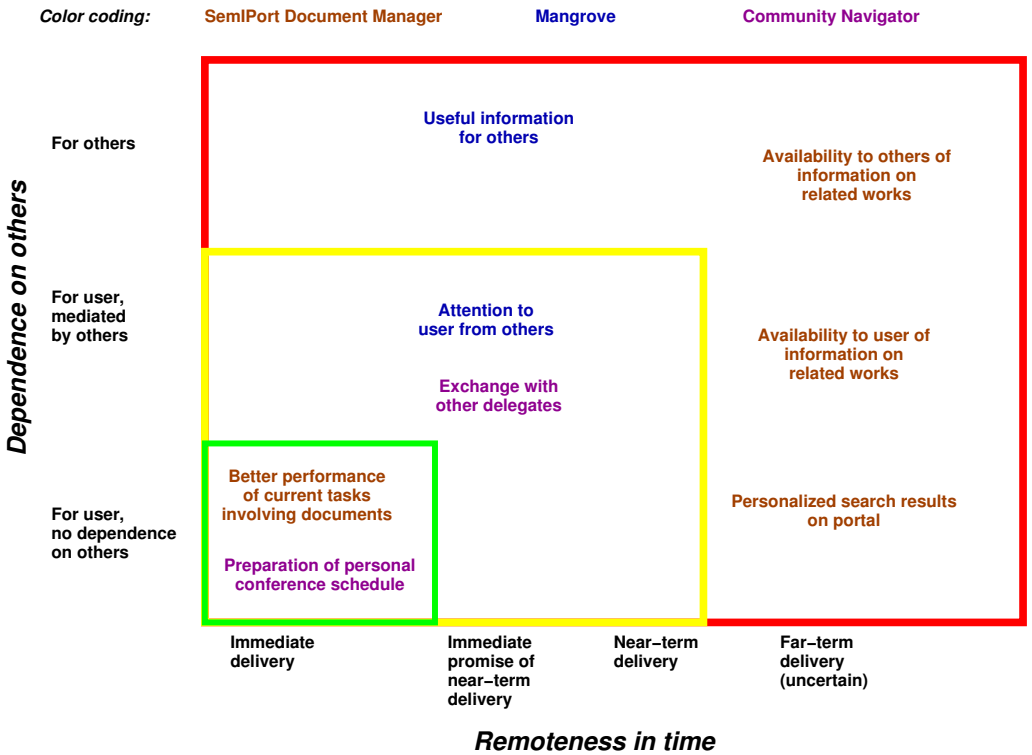

Existing Areas of Research

Social psychology

- Examples
  - Collective effort theory
  - Goal setting theory
- Utility
  - The theories generate unobvious predictions
  - ... but these often aren't confirmed in practical settings

Groupware, online communities

- Utility
  - The ideas have been tested in practical settings
  - ... but these settings are somewhat different from those of semantic web applications

Some Ideas From Research

How to (maybe) motivate users to contribute

- Emphasize the uniqueness of their possible contribution
  - "Only you can do this"
- Remind them of the benefits (for themselves, for the group)
  - ⇒ May backfire
- Publicize their contributions
  - With or without quality ratings

Caveat

- Try it out in your setting first!

Cf. Beenen et al. (2004), Cosley et al. (2006) (references given above)

How to Involve Users in Research
Introduction

- This final section turns to the general question of how research on semantic web technologies can benefit from increased involvement of users
- It specifically addresses concerns that have been voiced by members of the semantic web community
- Further information concerning the involvement of users in design and evaluation can be found in the following sources, among many others:
  - Jacko, J., & Sears, A. (Eds.) (2006). *Human–computer interaction handbook (2nd ed.).* Mahwah, NJ: Erlbaum.

    A comprehensive handbook whose second edition will be appearing within the next few months
  - Stone, D., Jarrett, C., Woodroffe, M., & Minocha, S. (2005). *User interface design and evaluation.* San Francisco: Morgan Kaufmann.

    One of several good introductory textbooks
  - Mayhew, D. (1999). *The usability engineering lifecycle: A practitioner's handbook for user interface design.* San Francisco: Morgan Kaufmann.

    A more practically oriented, how–to–do–it guide to all phases of the usability engineering lifecycle

Why Not to Conduct User Studies (1)

1. The results of user tests are always positive anyway, so why bother?
   - The negative results don't get published, but they may be the most valuable ones for you
2. We love our systems and know that they're great
   - A generally useful motto is "Users aren't like you"
   - On the other hand, don't assume that they are different in every way (e.g., incapable of understanding even in general terms how a system works)

Why Not to Conduct User Studies (2)

3. By the time you get the results, it's too late to change anything anyway
   - It's better anyway to involve users throughout the design and development lifecycle (see slide below)
4. The semantic web doesn't exist yet, so how can we test semantic web technologies?
   - It's not just a matter of testing (see the slide below)
   - You can study parts of larger systems

     Making the parts worth using in themselves is a good idea anyway

Why Not to Conduct User Studies (3)

5. If you test your system with users, they may waste a lot of time dealing with uninteresting software bugs
   - This fact constitutes one advantage of the use of low–fidelity prototypes (e.g., Flash mockups or paper prototypes)

How to Exploit Knowledge About Users

Phase	Activities
Analysis of requirements	Conduct observations, interviews, ...
Interface design	Apply design principles and guidelines, psychological knowledge, ...
Iterative testing with prototypes	Use cheap mockups where possible
Summative evaluation of final version	(Long–term) field studies, lab studies, ...

Who Is Supposed to Do All This?

- Involve in your project one or more persons with some training and/or experience in user–centered design
- Can’t afford it?
  - Get suitably trained students
  - ... or researchers who can also contribute technically

Conclusion
Reminder of the Questions Addressed

1. What have users GOT TO do with it?

- Communicate with the system without being overwhelmed by ontologies
- Form an adequate mental model of what's going on
- See enough reasons to contribute for the common good
- (and many other things not covered in this talk)

2. What have users GOT TO DO WITH it?

- We need their participation at all stages if we want to develop applications that will be usable and *used*