



UvA-DARE (Digital Academic Repository)

Topic driven access to scientific handbooks

Caracciolo, C.

[Link to publication](#)

Citation for published version (APA):

Caracciolo, C. (2008). Topic driven access to scientific handbooks Amsterdam: SIKS

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <http://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Interacting with the LoLaLi Map

In Chapter 3 we described the design and implementation of the LoLaLi map, now we address the issue of providing the end users of the LoLaLi map with a suitable interface to it. Recall that we assume that the *end users* of the LoLaLi map may have some, but not deep, knowledge of the domain; they are not assumed to have a background in knowledge engineering. Instead, they can be assumed to have a fair level of familiarity with computers and the Internet. This type of user should be enabled to conveniently visualize, browse and search the entire LoLaLi environment consisting of the map and the handbook linked to it.

Currently, many visualization tools for graph-like structures (in particular ontologies) are available (e.g., Jambalaya [Storey et al., 2002], OntoViz [Sintek, 2003], IsaViz [ISAVIZ, 2004], Kaon [Oberle et al., 2004]), but they usually address experienced audiences, such as knowledge engineers or other types of experts. Also, the theoretical work on this issue has mainly focused on experienced users [Ernst et al., 2003]. In contrast, the visualization of structures like the LoLaLi map for end users of the type we envisage has remained relatively unexplored. The challenge is to give our end users a sense of what the structure is like, to maintain a sense of orientation in it, where appropriate to present detailed pieces of information about each element in the map, and at the same time hide technical details from the user.

We start our work by making explicit the requirements that we impose on the interface, then, after reviewing related work, we present our proposal. The interface we propose aims at balancing a global view of the map with a focus on the node(s) being visualized in more detail. It is a proof-of-concepts, our interest is in both the visualization of the map and in the interaction with it. The visualization is then based on the notion of a focus node that allows the user to concentrate on a single node without losing sight of the surrounding area. We also report on user studies that were per-

formed to check the usability of the interface and to explore how the background of the users influences the understanding of the interface and browsing in it.

The remainder of this chapter is organized as follows: in Section 4.1 we give the requirements that a user interface for the LoLaLi map should fulfil. In Section 4.2 we introduce the issue of balancing global and local views for visualizing maps like ours. In Section 4.3 we describe the interface designed for the LoLaLi map and highlight its main features. In Section 4.4 we present the user studies we performed and discuss their results. Finally, in Section 4.5 we draw our conclusions.

4.1 Requirements for the User Interface

Our assumption is that our end users have a fair level of familiarity with computers and the Internet, and at least some background in the area covered by the map. The typical activity for such users is learning a subject, and in particular using reference tools, such as handbooks. When consulting a reference tool, users have to identify their information need and find out how to satisfy it, either by issuing a query, or searching for it in some other way. Then, they have to inspect and read the results of their search, and possibly confront them with other sources. Obviously, no editing action on the map is contemplated. The profile just outlined leads us to consider the following requirements:

Requirement A. Present each topic with its *context* in the graph.

Requirement B. Enable smooth browsing the map.

Requirement C. Enable intuitive search of the available resources.

Let us discuss these requirements in some more details.

Requirement A. Present each topic with its *context* in the graph. Graph structures are complex, and it is challenging to balance a broad, *global view* of their overall structure, and a *local view* that privileges the focus on single nodes. When inspecting a topic, the user should have a good view of the pieces of information attached to it, and at the same time, the user should not lose sight of the surrounding area, or at least of a relevant part of it. For each node in the graph, its direct ancestor and descendants, together with its siblings and the relation it has with them, are all relevant pieces of the context. Other fundamental pieces of information to show include the node's gloss and its links to the handbook. All this should be presented on the same screen, and the number of clicks required to inspect this information should be kept minimal.

Requirement B. Enable smooth browsing. A smooth browsing experience implies that the user always has a clear grasp of the location and the meaning of the information presented in order to identify the right region of the screen to look at and the links to

click on. The complexity of the graph, and in particular, the presence of multiple parents, may hamper smooth browsing. Also, the depth of the graph and the possibly large number of subtopics available may be an issue.

Requirement C. Enable intuitive search of the available resources. The LoLaLi environment includes a map, the textual resource attached to it (i.e., the handbook) and the links that bridge the map and the handbook. Therefore, the search facilities should enable the user to search across each and all these components in a unified way. These search facilities should not require extra effort from the user in terms of clicking and exploration of the interface, and they should always be available together with the browsing facilities. A “best match” search (like most Internet search engines, and as opposed to an “exact match” database-like search) should be enabled as the default search that allows the user to type in a keyword or simply a phrase with no need to specify explicit constraints or filters on the search.

4.2 Related Work on Tree and Graph Visualization

The focused study on how to visualize graphs began with the seminal paper by [Tutte \[1963\]](#), who suggested an algorithm that places each vertex in the center of its neighbors. Now that structures like trees and graphs are increasingly used for the organization of data and knowledge, this issue has received great attention.¹ Much of the work, however, concentrates on presenting global views, as is shown by attempts to define when a graph is *understandable* and *nice* [[Purchase, 1997](#), [Rosenstiehl et al., 2001](#)]. Common requirements to achieve that result include the following: edges and vertexes should not overlap; vertexes should be regularly distributed in the available space; if any symmetries of the data are present, they should be made visible. In the rest of this section, we survey approaches known in the literature and present them by adopting the distinction introduced in Requirement A above between global and local views.

4.2.1 Global views

Well-known global view oriented visualizations (for trees and/or for graphs) include hyperbolic trees, treemaps, cone trees and cluster maps.

Hyperbolic trees. [Lamping et al. \[1995\]](#) proposed a tree visualization method that allows focusing on one part of the data in the context of the whole unfocused data set. A hyperbolic tree displays a hierarchy of any size within a finite circular area so that

¹International conferences are devoted to this issue, such as InfoVis (standing for *Information Visualization*), which started in 1995 and especially deals with the visualization of hierarchies and databases, diagrams, and information spaces. In 2003 the conference became a contest, providing an evaluation benchmark for information visualization techniques and systems. Even more focused is the series of Symposia on Graph Drawing, held regularly since 1992.

the node that is being focused on is centered and nodes away from the central node are made exponentially smaller. An example of visualization based on hyperbolic trees is the Visual Thesaurus [Plumb Design, 2004] interface, where nodes are displayed around a circle (Figure 4.1). Hyperbolic trees allow for larger representations near




Figure 4.1: A screenshot from the Visual Thesaurus centered around the term ‘model.’

the focused area while still displaying the overall structure of the tree. Hyperbolic trees can also be animated and/or used to display diagrams in three dimensions, and can be modified to represent graphs. The problem with hyperbolic trees is that they do not scale well with the size of the tree, as the screen easily becomes too cluttered when increasing the number of nodes to accommodate. The dynamic layout of the nodes allows for a certain degree of flexibility in the amount of nodes to show, but the downside of this flexibility is that the visualization offers few points of (visual) reference, making it harder to discover paths in the structure. The fact that nodes in the tree move when clicking on them forces the user to get acquainted again with a new display. Finally, it may be hard to include more pieces of information about nodes than their very labels.

Treemaps. The treemap visualization proposed by [Bederson et al., 2002, Johnson and Shneiderman, 1991] displays hierarchies as series of embedded boxes. Each box




Figure 4.2: A treemap representation of some of the LoLaLi map data.

contains its children, usually by alternating at each level between laying them out vertically and laying them out horizontally (Figure 4.2). The relative size of each box can be made proportional to some property so that the parent's property is represented as the sum of the properties of the children. For instance, in a hierarchical file system, the size of a directory is represented by the sum of the sizes of all files (recursively, directories) in it. In this sense a treemap is a space-constrained tree visualization algorithm that turns a tree into a planar map. Treemaps are often used to visualize tree structures of directories and files, but they are also used for applications in domains as different as zoology taxonomy, bioinformatics, and finance [Baehrecke et al., 2004, Cable et al.,

2004, Plaisant et al., 2003]. The treemap visualization method is good at showing high-level structure, and it is also more suitable than a hyperbolic tree for visualizing more details about nodes in the tree. Also, since nodes can be characterized by their size and color, it is easy to recognize patterns across the structure. The disadvantage of such a visualization is that it does not maintain consistent layout when navigating through different levels of the structure, and it does not easily represent multiple parenthood.

Cluster maps. Fluit et al. [2003] proposed a cluster-based visualization technique that aims at visualizing populated, lightweight ontologies.² A cluster map visualizes the subclass relations between classes and the instances of each class, together with a label stating the name of the class and the number of its instances. A cluster map is especially good at showing the level of overlap of instances between classes, and at giving an immediate visual feeling for the size of a class (Figure 4.3). The advantage of




Figure 4.3: An example of cluster map visualization, taken from [ADUNA, 2005].

²The phrase “lightweight ontology” is used to refer to simple taxonomies with few axioms and only instances and definitions associated to each element of the taxonomy. Database schemas can be considered lightweight ontologies. Ontologies that have a richer structure or that are extensively axiomatized are referred to as heavyweight ontologies.

such a visualization is that it easily represents graph-like structures and is scalable to a large number of instances. Its major merit is, in fact, its ability to represent associations of classes and elements belonging to classes. The disadvantage is that it is not as scalable to a large number of classes as it is to large number of instances, and when the structure only includes classes this visualization becomes similar to a diagram-like visualization, like the one provided by Graphviz (see the figures in Chapter 3).

Cone trees. Robertson et al. [1991] visualize trees in 3D. In the representation they propose, the root of a tree is located at the tip of a transparent cone and the children of the root, and the subtrees stemming from them, are recursively arranged around the base of the cone. This allows a denser layout than traditional 2-dimensional diagrams, but nodes may be obscured, which forces the use of shadows or projections to make all nodes visible and give the perception of the whole shape. Cone trees are good at showing the global view of the structure, but the visualization becomes less clear when there are many nodes and many levels.

4.2.2 Local views

Visualizations that are local view oriented privilege a detailed view of single nodes of the structure. An example of that is the visualization used by many DOS and Unix-based operating systems to display hierarchical file-systems through command lines. Each level of the tree is represented by the list of its contents and the user navigates up and down the tree by means of typed commands. This way, the user can inspect one level at a time in great detail. This visualization method has the advantage of simplicity and clearness, since it provides a well focused view on each single level of the tree. Interfaces based on this approach can be defined as strongly local view oriented. They utilize a “stimulus reinforcing” approach: by repeatedly accessing directories and files, the user can reinforce and internalize her mental map of the structure. This process is facilitated by the stable schema of representation of the inspected level. Once the structure is known by heart, the user is able to quickly navigate through it. The downside is that the structure is implicit in the interface and active querying is required by the user to learn the structure and inspect its content.

Most current user interfaces for thesauri and ontologies display all pieces of information for a single node, plus links that point one level up or down. An example of this approach is the interface for the Medical Subject Headings (MeSH) [MESH, 2007] of the National Library of Medicine; Figure 4.4 shows a fragment of its structure as an indented tree. A directed graph can be presented as a tree by duplicating a child node as many times as the number of its distinct parents (i.e., tree branches are duplicated). In Figure 4.4 the duplicated child node is ‘Myocardial Infarction.’ The main advantage of this practice is that by displaying the graph as a tree, the complexity of the structure is simplified. The disadvantages are that much duplicated information is shown, and in order to identify and fully inspect nodes with multiple parents, the user has to compare all the branches shown. This operation can be facilitated by color-coding, like



Figure 4.4: A snapshot of the interface of the Medical Subject Headings (MeSH). The heading on focus is 'Myocardial Infarction.'

in the MeSH browser example (where the heading under inspection is marked by an arrow head to its left and printed in red), but this, however, does not fix the inherent weakness of this approach. To grasp the graph structure the user has to look at all the

shown branches and compare them. While this might be a workable option with two branches, it quickly becomes more difficult when a node has more than two parents. In addition to these difficulties, this type of visualization also fundamentally misrepresents the graph because at any level it only shows one parent per node. In this way, only one of the paths connecting a node to a given ancestor is shown, while there could be an entire line of multiple parents along that path.

Summarizing, the main difficulties and limitations of the visualization approaches we presented concern the representation of multiple parenthood, since the presence of multiple parents may make the screen cluttered and scarcely readable. The solution of transforming the graph into a tree, by duplication of branches, does simplify the structure, but it also introduces much repetition and does not scale well. Interfaces based on a hypertree visualization are better solutions, but the fact that they dynamically lay nodes out requires more effort on the user side in order to place herself in the context of the map. Hypertrees are usually very pleasant to watch and interact with, but for the reasons given above we do not consider them to be a viable solution for visualizing resources to be consulted often or that carry the type of information carried by the LoLaLi map. Moreover, the fact that the tree moves and redisplay nodes when clicking on any of them can be distracting and consequently not apt for a resource such as the LoLaLi map.

4.3 A User Interface for the LoLaLi Map

The interface we designed in collaboration with the User Centered Design group at Elsevier Science³ took the principle of information visualization called “focus + context” [Card et al., 1999] as its starting point. According to this principle, the most important data must occupy the focal point, in full size and detail, while the area around the focal point (the context) is displayed to help the user make sense of how the important information relates to the entire data structure. Often, the principle of “focus + context” is applied in such a way that regions far from the focal point are displayed smaller (as in a *fish-eye view*).

According to the “focus + context” principle, our interface divides the screen into a *map area*, the center of which is occupied by a *focus topic*. Parent and child topics of the focus topic are located respectively on the left and right side of it. Below the map area is the *data area*, where additional pieces of information, including glosses, and links to the handbook are shown (Figure 4.5). The search box and the button to start browsing from the top node are always available. The list of results is presented in the same layout as the rest of the interface (more details about the implementation of the search facility in [van Hage, 2004]). Then, the user is able to start exploring the surrounding area by clicking on topic titles or by rolling the mouse over topics in the map area.

³A special thank to Stephan Stipdonk for his invaluable contribution to the making of the interface.

The interface is implemented using Flash [Flash, 2005], which allows for animated movements that create a seamless browsing experience. The map is stored in RDFS format [RDFS, 2004] accessed by Sesame [SESAME, 2005], with a MySQL database [MySQL, 2007] as the back-end.



Figure 4.5: A snapshot of the LoLaLi interface.

The user can inspect the map in one of the following three ways:

1. by clicking on any of the topics (on its title) in the map area; in this case, the topic clicked on becomes the focus topic (Figure 4.6);
2. by rolling the mouse over a topic title in the map area; in this case, additional data about it is shown, keeping the focus topic fixed on the screen (Figure 4.7);
3. by using the search facility and then clicking on any of the returned nodes (Figure 4.8).

The action of clicking on one of the parents or children of the focus node makes the node occupy the focus area; then the map shifts “forward” (if one clicks on a child, cf. Figure 4.5 and Figure 4.6) or “backward” (if one clicks on a parent, cf. Figure 4.5 and Figure 4.7). The roll-over on a child topic makes its gloss appear in the data area, together with its relation with the focus node. In case it has other parent nodes beyond the focus node, those extra parents also appear in the focus area (Figure 4.7⁴). Symmetrically, when rolling the cursor over one of the parents of the focus node, all

⁴In the data area shown in the figure, modal logic appears labeled as ‘subclass of logic’ whereas in Chapter 3 we talked about a *subtopic* relation. The reason for this is given in Section 4.4.1, last paragraph.

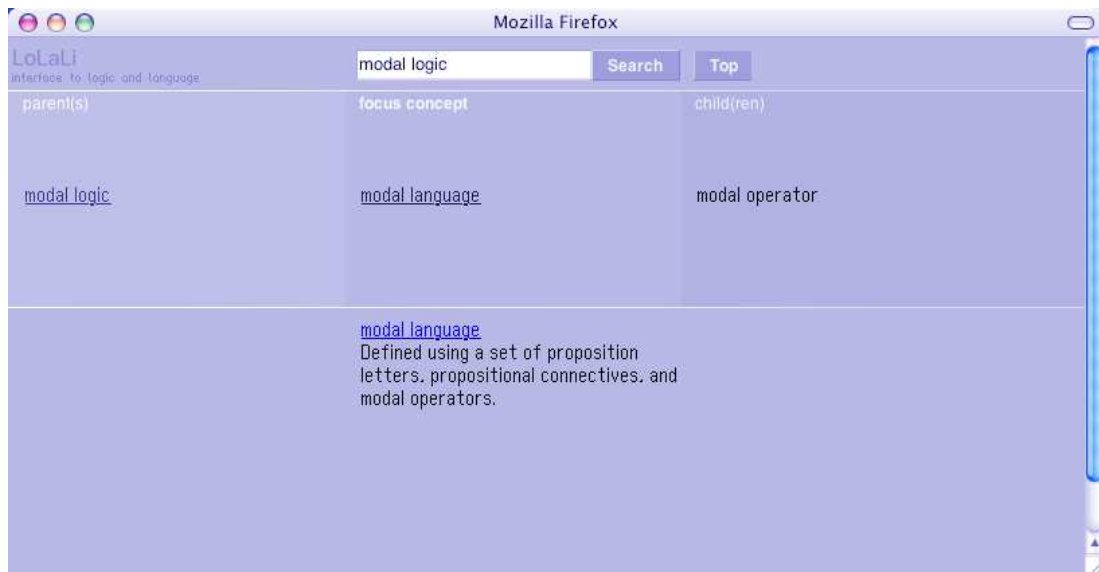


Figure 4.6: LoLaLi interface: after clicking on a child topic.



Figure 4.7: LoLaLi interface: rollover effect on modal logic.




Figure 4.8: LoLaLi interface: search results.

siblings of the topic that has the focus (with respect to the parent) appear in the focus area, while the gloss of the parent, and the relation of the focus to the parent appear in the data area. This rollover behavior is meant to allow users to quickly get information about a topic and assess its meaning and position in the map. The combination of displaying topics on rollover and the animated movement is meant to create a seamless browsing experience when moving from one node in the map to another. Summarizing, the visualization we provided has the following characteristics:

1. It makes the graph planar, i.e., without crossing of edges. In this way, it simplifies the complexity of the graph without misrepresenting it.
2. There is always a node in the focal position. In this way the structure is not presented as a whole in which users have to find their way (as in global views), but the user is always situated somewhere *inside* the map. The movement involved when changing focus concept creates the experience of navigating one coherent structure.
3. The interface limits the user's memory overload, as the user does not have to memorize the entire structure (as in the only-local views). Also, the overall layout of the map does not change during navigation, and the location of nodes relative to one another is fixed.
4. It is easily scalable, in terms of number of topics, relations, links and length of the glosses.

5. The interface maximizes the amount of information displayed for each node (as in local views), while showing essential information about the nodes in the surrounding area.

The LoLaLi interface does not exploit user familiarity with “collapsible” tree interfaces (such as the MeSH browser), but the motion effect created when clicking on concepts gives a feeling of continuity, as opposed to the feeling of “jumpiness” typical of hyper-text documents. Because at any one time only a small part of the map is shown, it is crucial that moving in the map is as easy and natural as possible so the user does not feel inhibited. The seamless movement through the map also creates the sense of being in a larger coherent structure [Friedrich and Eades, 2002].

The roll-over effect is used to limit the amount of information displayed at one time in the screen, as sometimes less data has to be shown in order to fit on the screen without overloading the user [Tufte, 1988]. Also, the information shown on roll-over allows the user to decide which way to go before actually clicking. This mechanism gives users full control of the information flow, and they can absorb information at their own pace.

4.4 User Studies

Given our interests in end users and our aim of supporting their information seeking process, we ran a user study to assess the usability of the interface described above. More specifically, we wanted to identify our design’s most important usability problems [Wilson, 2000]. Usability is the quality concerning how easy user interfaces are to use; usability tests are primarily, if not uniquely, concerned with ‘effectiveness,’ and so they imply a quantitative evaluation. For example, if the usability test concerned an interface for an email program, a quantitative goal could be stated as follows: “After inspecting the interface for five minutes, the user is able to send a message to a group of people in the address book in less than two minutes and making less than three errors.” User studies are necessarily more qualitative, although they do take into consideration the quantitative data collected during the studies.

The user studies we conducted were aimed at answering a number of questions:

1. As for the usability of the interface, we were interested in knowing what, if any, are the main problems with the interface as a tool for navigating the LoLaLi map. And, can users find their way around the map using the interface described in the previous section?
2. As a secondary goal, we aimed at testing if and how the usability of the interface changed depending on the level of education of users with comparable background and on the level of familiarity with computers and the Internet.

To collect information, we designed a questionnaire. The questionnaire was not meant to be filled in by the users. Instead, two interviewers asked the questions directly to

one participant at a time and the questionnaire served as a reminder for the question to ask. During each interview one interviewer was active, asking questions, explaining the tasks for the user to perform and taking notes, while actions on the screen were recorded by suitable software. The second interviewer was always present, to take notes and occasionally ask additional questions.

The questionnaire included closed (i.e., evaluative) and open (i.e., asking users to freely provide comments) questions. The evaluative questions were formulated according to a classic Likert-style format [Likert, 1932], and formulated in a way so as to avoid as much as possible the risk of influencing the users. As for the answers to the questions, we report all answers and comments provided by the users and wherever useful, we compile summary tables with all figures divided by groups (see next section) and the totals.

The interviews were conducted in Dutch with the Dutch speakers, in English with the others. As a sanity check for the questionnaire, before proceeding with the real user studies, we administered the questionnaire to two PhD students in the area of logic and physics respectively.

4.4.1 The Setting

Below we explain and motivate the decisions taken for the organization of the user studies, corresponding to the following issues:

- What type of interface prototype to use?
- Who should participate in the tests?
- How many participants to involve?
- What kind of questions and/or tasks should the users answer and/or perform?

What type of interface prototype to use? The prototype to use for the user studies could be either a paper mock-up, or an interactive prototype. We opted for the latter type, because of the intrinsically dynamic nature of the interface we developed, and because an interactive prototype seemed better at uncovering problems than a paper prototype. As for the features that the prototype should include, Nielsen [1989] distinguishes three types of prototypes: horizontal, vertical and scenario. A *horizontal* prototype only contains a shallow layer of the surface of the user interface (for example, a menu option, with no operation associated), a *vertical* prototype fully implements only a small number of features of the interface, while a *scenario* prototype fully implements only the operations necessary to perform a certain task. We decided to test a vertical prototype, which would enable us to see how users would use the interface to browse and search the LoLaLi map, and, partially, to inspect links to the handbook. Such a prototype should also give insight into currently missing features of the interface.

Who should participate in the tests? As for the users to invite to the study, it was natural to address an audience of university students (cf. Section 4.3). We contacted students in an area related to that covered by the map, but distinguishing two levels of educational profile, namely masters and undergraduate students.

Note that if we had been interested only in assessing the soundness of the interface, we could have performed a “heuristic evaluation,” consisting of an evaluation done by a pool of human computer interaction and design experts. There is evidence suggesting that if the pool only includes experts, close to 100% of all interface problems are found [Jeffries et al., 1991]. Instead, we preferred to have a sample of the type of people that would actually use the interface so as to be able to have insight into their view on it. We did not expect to find all problems of the interface, but only the major ones, arguably the best strategy for this type of problem [Bailey et al., 1992].

How many participants to involve? Nielsen and Landauer [1993], Virzi [1992] and Lewis [1994] published influential papers on the topic of sample size in usability testing. The authors presented a mathematical model for determining the sample size for usability tests, with empirical evidence for the models and several important claims: 80% of the usability problems are detected with four or five participants, while the proportion rises to 90% if there are ten participants. The “small sample” claims and their impact on usability methodology have been popularized in Nielsen [2000]’s widely read “useit.com” online column. Recently, it has been argued that a sample size larger than five may be required to detect a satisfactory number of usability problems. Taking into account these considerations, we aimed at two groups with six to ten participants.

What kind of questions and/or tasks should the users answer and/or perform?

The LoLaLi environment (map and handbook) is to be used as a reference tool to support our users, and ideally our user studies should include these types of real life tasks. Then, the tasks our intended users are usually called to perform include the following: “Write an essay on topic Y,” “Read about topic X and answer the following questions,” and “Discuss similarities and differences between topic Z and W.” Unfortunately, this type of end-to-end tasks does not fit the context of our work: it would be problematic to assess, and it would hide the operations of searching and browsing the map—the very basic operations that the LoLaLi environment should support. Therefore, we opted for subtasks of a real life task, consisting in searching about a topic (“what is . . . ?”), both at a shallow level, that could be answered by the map, and a deeper level, answerable by following links to the handbook. Also, users could be interested in finding out connections between important topics in the area. Then, the tasks we included in the questionnaire are of the type: “Browse/search for topic modal logic and do the following/answer the following question.” The drawback of this type of tasks is that it does not encourage users to pursue their own curiosities and information needs.

The user studies involved 14 people: 6 masters students (5 students in the Master of

	<i>UG</i>		<i>MSc</i>	
	<i>M</i>	<i>F</i>	<i>M</i>	<i>F</i>
English	-	-	2	1
Dutch	7	1	1	1
Other	-	-	1	-

Table 4.1: Composition of the groups of undergraduate students (UG) and masters students (MSc) on the basis of gender (M/F) and native language.

Logic program, MoL, at the University of Amsterdam, one Dutch student in Humanities Computing), and 8 undergraduate students (8 first year students enrolled in an Artificial Intelligence course, one of which was a student in Philosophy and the others in Artificial Intelligence). Table 4.1 summarizes the composition of the participants on the basis of their university status, gender and native language. In all tables in this chapter we use the short forms MSc for masters students and UG for undergraduate students.

At the time we conducted the user studies, the content of the LoLaLi map was being converted into RDFS (cf. Section 3.7.2). All hierarchical relations were initially rendered as a property called ‘unspecified subtopic’ that was incrementally specified over the course of the project. At the time of the user studies, the relations used were: subclass, part-of, instance, mathematical result, computational tool, notion, historical view. Theoretical considerations discussed in Section 3.4.1 (concerning the many flavors of the is-a relations), and the observations made in the course of the user studies (concerning the relations unspecified subtopic and notions) make us select the relations described in Chapter 3.

4.4.2 The Questionnaire

After we analyzed the goals of our user studies (along the lines described above), we singled out the following questions to which we wanted our user study to provide answers.

1. On finding information:
 - (a) how do the users perceive the browsing with the interface?
 - (b) what kind of difficulties do the users encounter?
 - (c) is searching difficult?
 - (d) do the users have a preference for search over browsing?
 - (e) how do the users perceive the links to external resources?
 - (f) what expectations do user have about them?
2. On the map structure:

- (a) is the structure of the LoLaLi map understandable in its main features?
 - (b) what is the value of relations in terms of understandability of the structure?
3. On some features to be considered for later inclusion:
- (a) would users find back button, bookmarks, and grouping of subtopics useful?

We phrased our questions in as neutral a way as possible, so as to avoid pushing users towards “socially desirable” answers. At the very beginning of the questionnaire, the users were asked some questions about their computer literacy and information gathering strategies. The complete questionnaire is included in Appendix A.

4.4.3 Results

We now report on the results of the user studies following, section by section, the questionnaire administered to the participants. The sections covered are:

- Computer literacy
- Information gathering strategies
- Browsing
- Searching
- Features of the LoLaLi map
- Links to the handbook
- User preferences
- User wishes

Computer Literacy

Both groups (undergraduate students and master students) present a high level of computer literacy (Table 4.2). Even though 4 undergraduate students declared they had no formal computer training at all, this does not correspond to a lack of education, as is suggested by the high number of programmers among them. The two groups are also comparable in terms of years of computer usage (in each group, 4 people have 10 years or more experience), as well as in terms of frequency of Internet usage (almost all on a daily basis). The proportion is slightly different if we look at the years of Internet usage: among the undergraduate students, 6 people have 4–6 years experience, one person less than one year, and one more person between 1 and 3 years. Among master students, 2 people have 7–9 years experience, while 4 have 4–6 years. As for the top three activities with computer, master students mainly write documents and papers, use email and surf the web. Undergraduate students program, surf the web, and use email; less frequent activities are video editing, graphics and computer games.

Information Gathering Strategies

The way people from the two groups approach new fields is comparable (Table 4.3). Undergraduate students tend to go for a combination of web search engines (7 people)

<i>Questions</i>	<i>Answers</i>	<i>UG</i>	<i>MSc</i>	<i>Total</i>
Which operating system do you use more often?	Windows	5	2	7
	Linux	2	1	3
	Macintosh	-	1	1
	Windows <i>and</i> Linux	1	2	3
What formal computer training do you have?	none	4	-	4
	college courses	2	-	2
	undergrad. courses in computer science	1	1	2
	minor in computer science	-	3	3
	major in computer science	-	2	2
How many years of computer usage?	4–6	1	1	2
	7–9	3	1	4
	10	4	4	8
How many years of Internet usage?	1	1	-	1
	1–3	1	-	1
	4–6	6	4	10
	7–9	-	2	2
Frequency of Internet usage	weekly	1	-	1
	daily	7	6	13
Which are the most frequent activities on screen?	writing papers	3	6	9
	e-mailing	5	5	10
	surfing Web	6	5	11
	programming	7	2	9
	playing games	1	-	1
	graphic, video	2	-	2

Table 4.2: Questions on computer literacy of the participants in the study. Multiple answers were allowed.

and other resources, i.e., people, textbooks, and libraries. Master students use web search engines, ask people, and consult textbooks.

Glossaries are never used by undergraduate students, and by only one master student, yearly. Dictionaries are slightly more frequently used: by undergraduate students monthly (2), weekly (4), daily (1), by master student, monthly (2), and weekly (2). For obvious reasons (cf. Table 4.1), many master students consult bi-lingual dictionaries. As for encyclopedias, only one master student and 2 undergraduate students declare to use them (daily, and monthly and daily respectively). These resources tend to be consulted online, and the three people consulting them on CD-ROM are all undergraduate students.

Only three master students take notes, either on file (1) or on paper (2), while undergraduate students do it on paper (3), on file (3), or depending on the moment (2). Both groups use paper textbooks, but also consult electronic texts: master students in

<i>Questions</i>	<i>Answers</i>	<i>UG</i>	<i>MSc</i>	<i>Total</i>	
What is your preferred way to approach a new field?	web search engine	7	4	11	
	people	6	4	10	
	textbooks	3	3	6	
	libraries	2	-	2	
Do you use reference tools?	glossaries	never	8	7	15
		yearly	-	1	1
	dictionaries	monthly	2	2	4
		weekly	4	2	6
		daily	1	-	1
	encyclopedias	monthly	1	-	1
daily		1	1	2	
In what electronic format?	online	5	5	10	
	CD-ROM	3	-	3	
How do you take notes?	on paper	3	2	5	
	on file	3	1	4	
	it depends: paper, file	2	-	2	
What digital resources do you use more often?	e-assignments	1	1	2	
	handouts, slides	3	4	7	
	journals, papers	1	5	6	
	web sites	1	1	2	
	depends on web search engine	2	-	2	
	other	3	-	3	

Table 4.3: Summary of usage of electronic resources.

particular for assignments, handouts and journal articles; undergraduate students also use message boards and web sites. Interestingly, but not surprisingly, two undergraduate students claim that the kind of electronic document they consult depends on what the web search engine they use returns.

Browsing

The aim of the questionnaire's section on browsing is to see how the users orient themselves when browsing the interface, what problems they encounter, and whether they like the experience. In this part of the questionnaire, users are asked to reach a given topic by browsing (hints are given), and comment on the experience. An example of a browsing task is the following: "Browse until you get to the topic recursion theory, knowing that it is a topic in the area of theory of computation." Users were given 5 such questions.

Both groups perform the browsing tasks with no major problems. All of them, especially the undergraduate students, try first to retrieve the topics by searching (this is very evident for the first question), but once better instructed they performed the




Figure 4.9: Logic (1) is a subtopic of artificial intelligence, computer science and linguistics.

browsing tasks smoothly.

A remarkable difficulty experienced by the users concerned the position of the topic logic (1). All topics indicated by the questions belong to the branch of logic (1), which is an unspecified subtopic of mathematics, computer science, and linguistics (Figure 4.9). This implies that users had to choose which of those topics to inspect in order to find logic (1). Only one master student stated openly the problem (claiming that “logic is related to all of them, I could not decide where to go”), though all participants appear to have some moments of uncertainty before deciding where to look for logic (1). Many of them inspected one of the topics mathematics, computer science, linguistics, with an evident correlation between their choice and their background, and went back to it every time they needed to consult again the same topic.

Then, after performing the browsing tasks, users were asked to comment on the experience and mention at least three problems they encountered. We can divide the reported problems into two classes: problems concerning the structure/content of the map, and problems concerning the interface.

With regards to problems related to structure and content, master students mentioned the issue of the position of logic (1) sketched above (1 person); the fact that they had to think of the best strategy to apply, i.e., whether to start again from the top node or to click back until the right path is found (1); and that browsing is easier when topics are already known, though “you get used to it fast” (1). Among the undergraduate

students, one person admitted to having problems understanding the organization of the map, two people commented that browsing is easier when you have some prior knowledge of the area. As for problems concerning the interface, master students reported problems with scrolling the lists of topics, namely that they could not scroll with the rolling wheel (2 people), problems with the particular mouse used during the interview (2), and problems with the data appearing for the focus topic when the mouse is on a parent topic (1). Undergraduate students did not report problems with the interface, but three of them pointed out that they usually only search.

When asked to give a general comment on the interface, both groups were positive. Master students unanimously said the interface helped them to make a mental picture of the structure of the map and gave a generally positive comment on the browsing experience (5 people said “I like it very much,” one person “rather so”). Nevertheless, one person found it annoying that terms disappear when rolling the mouse, and that more options should be in the center associated to the focus node. Undergraduate students had problems with the mouse used during the interview (did not roll smoothly enough), and complained about the focus topics being underlined, which they took for the indication of the presence of a link (see Figure 4.5.)

Searching

The users participating in the study were given a number of tasks in which they had to search for a given topic, then perform an action or answer a question. The aim of this section is to test the usability of the search facilities provided, and the user’s comprehension of the information shown in the interface, such as the typed relations and the glosses.

1. Search for the topic operator and divide its subtopics into groups. The topic operator is presented with the following subtopics: quantifier, arity, modal operator, truth functional operator. Since both typed relations and glosses are shown in the interface, we expected that the users would group together modal operator and truth functional operator, and leave arity apart.

Searching did not present any difficulties for any of the users. As for the grouping, undergraduate students gave a variety of answers, from “I cannot make any distinction” (1), to “I can distinguish different types of subclasses.” One person had the intuition to group together topics that are “at the end of the tree,” thus making an intuitive distinction between topics with subtopics and topics without subtopics. Two people admitted that the differences among the subtopics were unclear to them. Among the master students, 2 people found the task clear and grouped the subtopics of operator as expected, 2 people said that only arity was clearly different from the others, and 2 people could not make any grouping.

After a few interviews we realized that the types of the relations were not clearly visible in the interface, which explained the difficulties of the participants in understanding the question. After realizing this, the interviewer pointed the relations out to

the users, who were then able to suggest sensible groupings among the subtopics. The improvement was equal for both groups. Unfortunately, a few people were interviewed before the source of difficulties was realized, and this is likely to have introduced noise in the study.

2. Find the topic logic (1) and distinguish different types of its subtopics. At the time when the test was performed, the topic logic (1) had some 30 subtopics, as shown in the fragment below, where SC stands for the relation `subclass`, PO stands for the relation `part-of`, MR stands for the relation `mathematical result`:

```

logic (1)
  SC temporal logic
  SC modal logic
  SC first order logic
  PO recursive function theory
  MR Löwenheim-Skolem-Tarski theorem
  MR Gödel's 1st incompleteness theorem (1931)
  . . . . .

```

We expected that the users would use the labels for the typed relations shown in the interface to group together topics, but people encountered the same difficulties discussed about the previous task, related to the visibility of the relations in the interface.

3. Explain the relations between logic (1), proof theory, cut elimination theorem.

For this task, users had to retrieve the topics proof theory, cut elimination and logic (1). and explain the relations between them. The users were also asked to indicate what features of the interface had given them the hint. In the LoLaLi map they were presented the three topics are related in the following way: cut elimination theorem is a `mathematical result` of proof theory, which is `part-of` logic (1) (Figure 4.10). Master students appeared rather comfortable with the sequence logic (1), proof theory,




Figure 4.10: Topics to retrieve for the task on searching: explain the relations among proof theory, logic (1), cut elimination theorem.

cut elimination theorem, and 5 of them gave the expected answer. Only one person reported problems with the relation `mathematical result`, claiming that the relation `part-of` is clearer than the other. In general, the relation `part-of` appeared not problematic, though they might call it “branch of” or “subset of.” When asked, all of them had a clear picture of the position of those topics in the graph (gained during previous browsing in the area). Undergraduate students also showed a good intuition

of the relations among the presented topics: 5 people gave the expected answer, 2 people commented that the relation `part-of` is clearer than `mathematical result`, one person said the opposite.

4. Explain why truth function has two parents. The topic truth function has two parents because it is a subtopic of truth and semantics. Users had to explain why, according to them, that happened. All users referred to the fact that “truth function had something to do with both,” and many were able to express their understanding of the relations among the parents and the child topic.

Features of the LoLaLi map

This section of the questionnaire aims at testing what difficulties users encounter with some features of the map. In particular, the section concentrates on relation types, on multiple parenthood, and on how these features are rendered in the interface.

1. Comments on the typed links. First of all, the users were asked to give a general evaluation of the typed links as seen so far. The meaning of the typed relations were generally clear to master students. Typed relations were clear to 4 people (2 said “very much clear,” 2 said “more or less clear”), while only one person said the relations are “hardly” understandable. Two people noticed they had some difficulties with the relation `unspecified subtopic` and with the relation `notion`. One of the undergraduate students rated the relations as “very much clear,” two as “more or less clear”, three as “rather clear”. One person admits to understanding the general idea of typed relations, but to having some difficulties with specific relations. In particular, difficulties occurred with the relation `notion` and with the relation `instance`, that seems to be understood as something in between `part-of` and `subclass`.

2. Assign a typed relation to pairs of topics. The users were asked to assign a typed relation to the following pairs of topics:

- modal logic – logic (1)
- theory of computation – computer science
- Gödel theorem – logic (1)
- operator – logic (1)
- Frege on quantification – quantification
- formal semantics – semantics

In Table 4.4 we report the entire list of results: the pairs proposed to the users are listed in the first column, the labels chosen by the users are reported in the second column (a ● marks the label adopted in the LoLaLi map), the following two columns show the answers given by the undergraduate students and the master students, respectively, while the last column reports the total number of answers given. One master student consistently did not provide answers to any of the questions.

<i>Question</i>	<i>Answers</i>	<i>UG</i>	<i>MSc</i>	<i>Total</i>
modal logic - logic (1)	SC ●	6	3	9
	SC or IN	-	1	1
	SC or PO	-	1	1
	PO	2	-	2
	no answer	-	1	1
theory of computation - computer science	PO ●	1	1	2
	subject area	-	1	1
	PO or SC	-	1	1
	MR	4	-	4
	Notion	1	-	1
	SC	1	-	1
	unsp. subtopic	1	2	3
no answer	-	1	1	
Gödel theorem - logic (1)	MR ●	4	4	8
	unsp. subtopic	-	1	1
	HV	2	-	2
	PO	1	-	1
	no answer	1	1	2
operator - logic (1)	Notion●	2	4	6
	SC	1	-	1
	unsp. subtopic	1	-	1
	PO	1	-	1
	used in logic	1	-	1
	MR	1	-	1
	IN	1	-	1
	no answer	-	2	2
Frege on quantification - quantification	HV ●	3	1	4
	IN	-	1	1
	unsp. subtopic	1	2	3
	Notion	2	-	2
	SC	5	2	7
	PO	1	-	1
	no answer	-	2	2
formal semantics - semantics	SC ●	1	1	2

Continued on next page...

... continued from previous page:

<i>Questions</i>	<i>Answers</i>	<i>UG</i>	<i>MS</i>	<i>Total</i>
	unsp. subtopic	1	1	2
	PO	2	-	2
	no answer	-	2	2

Table 4.4: Summary of results for the question: “Assign a typed relation to pairs of topics.” More than one answer was possible.

Let us take a look at each question individually. The pair modal logic – logic (1) is well understood, as 9 out of 13 people indicate `subclass` (to be read as: modal logic is a subclass of logic (1), and two more were undecided between `subtopic` and `IN` and `PO`; the same reading applies also in the following), while 4 people give other answers.

The pair theory of computation – computer science appears to be more difficult. Only one master student and one undergraduate student answered that theory of computation is a `part-of` computer science, while the remaining 14 people gave 7 different answers. Masters students mentioned: “subject area” (1), unspecified subtopic (2), `part-of` or `subclass` (1); the answers of the undergraduate students include unspecified subtopic (1), `notion` (1) and `subclass` (1).

The pair Gödel's theorem – logic (1) is also well understood, since 8 people labeled Gödel's theorem as a `mathematical result in logic` (1). Interestingly, 2 undergraduate students labeled it as `historical view`, probably because the name of the topic includes the name of the author, which gives a sort of historical flavor to it. This fact raises the issue of what is *historical* in our work, which is of course related to the definition of historical in science, and to the kind of study the user we address is interested in.

For the pair operator – logic (1) we can observe a clear difference between master students and undergraduate students, as four master students opted for `notion` and two had no idea, while almost each undergraduate student proposed a different label.

The pair Frege on quantification – quantification is an interesting one, as for ten people (6 undergraduate students, 4 master students) is either an unspecified subtopic or a subclass and only four people indicated `historical view`.

Finally, the answers for the pair formal semantics – semantics are like those for the pair modal logic – logic (1). a common confusion of `subclass` and `part-of`, among both master students and undergraduate students.

3. Explain why modal logic has three parents. Using the search functionality, the users had to retrieve the topic modal logic and explain why it is presented with three parents. Master students were all comfortable with the multiple parenthood: only 1 person said it is “rather clear,” while 4 people said it is “clear” or “very much clear.” All undergraduate students affirmed it was “very much clear,” or “rather so.” As an

explanation of the multiple parenthood, people from both groups mention the fact that modal logic has three parents because it has “something to do” with all three parent topics.

4. Describe what happens when you roll over semantic ambiguity in child position.

To answer this question, the users had to find the topic semantic ambiguity, make it occupy the child area and roll the cursor over it. When doing so, both parents of semantic ambiguity appear in the focus area (where usually there is only one topic at a time). This question was meant to see how visible and understandable the mechanism of roll over was to the users. Among undergraduate students, two of them said they would not have noticed that if the interviewer had not pointed it out to them. For three people it was “very clear.” For all master students what happens was “clear,” but 3 of them were annoyed by the fact that data disappears when moving the cursor away.

Links to the Handbook

The prototype includes some links to the handbook, accessible from the data area below the focus topic. Links to the handbook are marked by the title of the chapter from which they are taken, together with the available information about the position of the link target in the chapter. Our subjects were asked to inspect one of these links (suggested by the interviewer), state their expectations about the text they would find and comment on the differences between their expectation and what they actually found.

Master students gave articulate descriptions of their expectations, saying they would expect to find different definitions of the focus topic, related topics, topics “using” or “used by” the focus topic and other information about it. They also mentioned they would like to have a clearer description of what they would find when following the link, and that the source of the text should be highlighted so as to be aware of what text they were exploring. They also expected a shorter text than the one they were shown, and two people stressed the importance of using the same layout for the visualization of the text and of the map. Undergraduate students stressed their preference for a shorter text, well described before clicking on the link, and with a layout consistent with the starting page. Both groups agreed on the usefulness of having both glosses and links to larger pieces of documents than the gloss, and motivate this by saying that a gloss is useful to decide whether it is worth inspecting the link.

User Preferences

This groups of questions is meant to discover the users’ opinion about search and browsing facilities. Table 4.5 summarizes the answers given.

1. Did you like browsing the map? The users were asked a similar question at the beginning of the test, during the session dedicated to browsing. Now they have performed tasks involving both browsing and searching, therefore they can compare

Questions	Answers	UG	MSc	Total	
Did you like browsing?	very much	6	3	9	
	yes, but	miss back button	1	-	1
		no keyboard shortcut	-	1	1
		too much to scroll	-	-	-
rather so	hardly	-	2	2	
		-	1	1	
Did you like searching?	very much	7	5	12	
	more or less	-	1	1	
	rather so	1	-	1	
Preference for	searching	4	1	5	
	browsing	-	1	1	
	combination	3	4	7	

Table 4.5: Summary of the questions about searching and browsing.

the two experiences and comment on that. Undergraduate students were generally positive: the majority said they like browsing “very much” (6), “rather so” (1) and one person said “Yes, but miss the back button.” Among master students we find slightly more variability, ranging from “very much” (3), “rather so” (2), “hardly” (1); finally one person remarks that he appreciates the organization with relationships, but would like to have more commands from the keyboard, and that there are too many children to scroll.

2. Did you like searching the map? Master students were very positive about searching: only one person said he liked the experience “more or less,” the remaining 5 people said they like it “very much,” and one among them stressed the fact that “search is essential.” Undergraduate students were even more enthusiastic about searching: 7 of them said they liked it “very much,” only one person answered “rather so.”

3. Preference for searching over browsing, or vice versa? Among the undergraduate students, 4 people expressed a strong preference for search over browsing, and one of them remarked that searching is always the first thing, then browsing follows. Three people preferred a combination of search and browsing. Master students mentioned: a preference for a mixture of search and browsing (3), a strong preference for browsing (1), a preference for searching, or a combination of searching and browsing (1), a preference for searching followed by browsing (1). One person did not express a preference, but underlines that it depends on the context and the aim.

User Wishes

Finally, we asked users to express their wishes about features not currently included in the interface or in the organization of the map.

1. Would you like to see a bookmark functionality? The majority of the undergraduate students answered positively (7). Five of them also stressed the fact that it would be useful for larger maps than the one shown to them, one person observed that it would be useful for off-line use. Only one person said it was not necessary. Also among master students we have a majority of positive opinions about bookmarks, though more articulated: “Yes, but it depends on the number of topics in the map” (1), “yes” (2), “Yes, but I would not use it” (1), “I do not know” (1), “no” (1).

2. Would you like back buttons? The question was about a back button internal to the navigation system, since the back button of the browser brings the user outside of the LoLaLi map site. All master students thought that it would be a useful option, in particular 4 people mentioned that they would like to have the history for both browsing and searching. Undergraduate students were slightly less enthusiastic about it: 3 said they would not like having the back button, 3 people said they would, two people wished to have a search history, one person noted that after understanding the interface, the back button is not essential anymore (because already clicked topics are distinguished by the others by means of an underline). One person spent some time pondering about what exactly a back button would mean in the context of a graph structure.

3. Would you like to group children, for example according to the relation type? In general, the users seemed to be comfortable with this idea. Among the undergraduate students, there were 4 positive answers, one person preferred the alphabetical list, another person preferred a complete view of all children, finally one person desired a system to group children, but also wanted the possibility of viewing the entire list or having an alphabetical listing. Among master students, 3 people said they would like very much to have the grouping, 3 people said they would like to have it (one person notes that sometimes it could be good to look only at subclasses), but one of them observed that relations should be better explained. Finally, one person observed that the alphabetical list is fine if you know the area, otherwise it is better to get the complete list of all children.

4. Would you like more sophisticated search? When asking this question we left open what we meant by “sophisticated.” Users were explicitly told that it would not imply to learn a more complex query language, but that the sophistication would only involve the type of query they could answer. Master students answered yes, they would like to have a more sophisticated search system (3), “Maybe, but not immediately visible to the user” (2), “It depends on the use and size of the map” (1). Undergraduate students answered: “It depends on the size, if it grows, yes” (3), “Basic is good enough” (1), “It could be useful but I would not use it” (2), “Yes, but without having to learn an extra language” (1). Finally, one person simply answered no.

5. Comments, suggestions? Although the question was totally open, people tended to give more comments than suggestions. Undergraduate students were generally very positive: only one person commented that it took him awhile to realize it was a graph. Master students also gave positive comments; in particular one person claimed he would use the map provided with links for every day work. Two people noted that scrolling of the list of child topics (cf. Figure 4.7) is not handy.

4.4.4 Discussion

The preliminary survey on *computer literacy* and information seeking behavior confirmed that the two groups have comparable familiarity with computers, although master students have more experience with the Internet, on average. The two groups also have similar habits concerning the exploration of new fields: they both have a strong tendency to use web search engines, but master students also ask people more than the undergraduate students. Both groups make little or no use of glossaries or encyclopedias, but half of the undergraduate students use a dictionary weekly, while only one third of the master students does. If in electronic format, the dictionary is usually used online by both master and undergraduate students.

The *browsing* actions were performed with no major problems by both groups. One question required users to go to the topic logic (1), which is under computer science, linguistics, and mathematics. They tended to inspect only one of those paths and went back to it every time they needed to consult the topic again. We noticed that those who had training in mathematics looked for logic (1) under mathematics, and those who had a training in computer science under computer science. This suggests the importance of the user background for the use of the interface, and the consequent need to provide for accurate explanation hints.

When asked to comment on the browsing experience, only one master student was able to verbalize his difficulty with the position of logic (1) sketched above (he had tried more than one path). Undergraduate students seemed to have noticed the same behavior, but could not verbalize the issue. Other comments concerned the fact that when asked to browse for a new topic one had to think of the best strategy to apply, i.e., whether to start again from the top node or to click back until the right path is found. Finally, both undergraduate students and master students noticed that browsing is easier when topics are already known, though “you get used to it fast.”

When asked to give a general comment on the *interface*, both groups agreed that the interface helped them make a mental picture of the structure of the map and thereby of the underlying domain. They all gave positive comments on the browsing experience. Nevertheless, they stressed the annoyance caused by terms disappearing when rolling the mouse, and suggested that more options associated with the focus node should always be visible in the center of the display. It turns out that the motion effect when a topic shifts to the focus area is well understood, while the rollover effect seems to be more controversial. One type of rollover, over child topics (cf. Figure 4.7), seems to be easy to grasp (only two people said they did not notice the appearance of the parent

topics in the focus area). In this case, the user is shown the gloss of the topic over which she rolled the mouse, its relation type with the focus topic, and other parents of the child, if any (they will appear in the focus area). More problematic seems to be the rollover on parent topics (more child topics appear in the focus area), for which unfortunately, there was no dedicated question in the questionnaire.

The interface failed to make immediately clear the types of the relationships, as we noticed already during the interviews. This might be one of the reasons why, when we asked users whether they found the hierarchical relations clear, it appeared that even though the *typed relations* were clear or mostly clear to all participants in the tests, some confusions remained. The question on attributing a relation to pairs of topics reveals that the relations of `subtopic` and `part-of` are intuitively understood, but they can be confused with one another. This observation confirms [Storey, 1993], who suggests that an accurate understanding of even the most intuitive relations cannot be taken for granted. The domain-specific relations of `historical view` and `Notion` were especially difficult for undergraduate students but for different reasons. On the one hand students had different ideas about what “historical” is, depending on whether they consider historical something that was developed in the past and no is longer used, or something that was developed in the past and is still currently in use. On the other hand, the label “notion” seemed too general to many users. These observations made us to consider the appropriateness of the names used for these relations, and in particular their aptness to the background of the users. Finally, the relation `unspecified subtopic`, although it appeared in very few questions, seemed to generate some confusion.

The undergraduate students enjoyed the browsing experience (possibly more than master students) and search was equally appreciated by master students and undergraduate students. When asked to express a preference between searching and browsing, the users, especially undergraduate students, claim to prefer searching, and especially master students consider the usefulness of browsing and searching (or in combination, usually search first, then browse) depending on the different tasks they have to perform.

As for the features not currently included in the interface, a history button would be appreciated by all, while a bookmarking functionality is much wanted especially by undergraduate students (master students would appreciate it, but say they can also do without). We obtained more articulated answers when asking about the possibility of grouping children (for example according to relation types). In principle it would be appreciated, and our test subjects mentioned different views such as alphabetical order, ordered by relations, and so on. The interviewees also pointed out that the grouping views should not be exclusive, so that the user could go from one view to the other. Master students are in favor of a grouping based on relations. Finally, more sophisticated search functionality does not seem to be a priority for any of our users. They also said that in any case the first visible option must be as simple as possible.

4.5 Conclusions

In this chapter we have explored the issue of providing an interface for the end user of the LoLaLi map. We identified our typical end user as a university student in areas related to the one treated by the *Handbook*. We defined a set of requirements for such an interface, taking into account the operations that the user should be allowed to perform and the importance of obtaining the right balance between a local and global view. The graphical user interface described in Section 4.3 and developed in collaboration with the User Centered Design group at Elsevier Science respects those principles. The interface is organized into a map area, presenting a fragment of the map (i.e., a focus topic that occupies the center of the screen, plus parent and child topics to its left and its right, respectively), and a data area where data about topics in that fragment is shown. The interface is dynamic so as to allow smooth navigation from topic to topic. By showing only a relatively small fragment of the map at once, it makes the graph planar, because it hides the complexity of the graph without changing its structure. Also, by showing a small focus area, it is possible to show the hierarchical nature of the data in a consistent manner. The complexity is put back into the graph by means of the rollover behavior of the cursor. The possibility of showing larger portions of the data in one screen is sacrificed in favor of a display that does not force the user to re-scan the entire screen (as opposed to interfaces based on hyperbolic trees), and this allows easy inspection of the pieces of information associated to the topics.

In this chapter we also reported on the user study we performed (Section 4.4). The high level aim of the user study was to identify the most important usability problems with the proposed interface. The two specific questions we aimed to address concerned:

1. the usability of the interface (what, if any, are the main problems with the interface as a tool for navigating the LoLaLi map?), and
2. the influence of the background of the user on the usability of the interface

as expressed in questions 1 and 2 in Section 4.4.

Concerning the first question, the evidence we obtained suggests that the visualization approach we adopted suits our purposes. The interface is usable, and we obtained useful insights into what should be improved and how. Also, the rollover behavior of the cursor should be combined with some static features in order to allow users to easily inspect the map. All users found the search functionalities accessible from the interface extremely natural because of their familiarity with web search engines. In fact, the users' familiarity with the free-text type of search, together with the type of background they have, strongly suggests that more refined search tools are not essential to the usability of the LoLaLi map. Interesting work for the future would be to more deeply integrate searching and browsing in a way such that search results are displayed in a coherent fashion with the rest of the interface (for example, instead of presenting a

list of results, results could be presented together with a small fragment of map around them, and also making the gloss available to the users).

The graph structure of the LoLaLi map was well understood and not at all problematic. As for the relations, the main ideas behind them were understood by all participants, but some confusion remained. The lesson learned from the user studies was then to modify the set of relations used in order to minimize the source of confusion: the results of this process were presented in Chapter 3. We avoided the label “unspecified subclass” (that was meant to be temporary) and preferred to use the *subtopic* relation instead of the *subclass*. We also preferred to change the name of the relation “notion” into the longer but more expressive “feature and internal machinery.” Further user studies are advisable as a future work in order to investigate the understandability of the revised set of relations and of the revised interface.

As for the second question, the user studies showed the importance of the background of the users for the understanding and navigation of the map, for example when selecting a path to follow. One way to address this issue consists in improving the visual features in the interface, possibly together with good context sensitive help. Other ways would directly take into account the background of the user, for example by allowing personalization of the kind of help to be shown, or by defining different user interfaces for different groups of end users.

The fact that users do not consider bookmarks and back buttons important can be taken as a confirmation that the map is easy to navigate and that users have little or no difficulties in finding their way through it. On the other hand, they considered with attention the option of filtering the list of child topics: in general, people like to have a complete view of the entire list of children available, but this is obviously cumbersome to achieve when the list of children increases. A semantically oriented way to apply such a filter would take into consideration the available relationships and would allow the user to visualize only those topics related by one (or more) relationships. This option is certainly worthy of further investigation, but other selection criteria could also be valid and perhaps more informative to the user. The key question then becomes how to combine the strict symbolic information encoded in maps like ours with a notion of ranking/preference that could enable agile use of the structure.

The evidence provided by our user study confirms the usability of the graphical user interface proposed in this chapter. Looking forward, they suggest that we carefully consider the visual clues within the interface, and that master students are the best target for the LoLaLi map visualized by the proposed interface.

One of the sections of the questionnaire concerned the linking from the map to the text. Even though that feature was only partially implemented in the system used for the user study, we received encouraging feedback. Both undergraduate and masters students consider this a feature to develop and possibly to enrich with “metadata” about the content of the page they are going to inspect.

In the next two chapters we work on the selection of appropriate link targets for the topics in the map. We first inspect and compare methods for topic segmentation (Chapter 5), then we work on connecting the extracted segments with the map (Chapter ??).