Summary

The focus of the PhD project reported here was Web navigation. The term "Web navigation" was used in a broad sense referring to users' orientation in an information space, locating information and progressing from one information source to another. The Web has brought not only the opportunity of a nonlinear access to information sources but also the challenges of *cognitive overload* and *disorientation* (Conklin, 1987; Edwards & Hardman, 1988). Involving spatial features (syntax) in processing contents (semantics) was recognized as a distinguishing characteristic of Web navigation (Di Blas, Paolini, & Speroni, 2004).

As stated in *Chapter 1* – Introduction, the objective of the research presented here was to build a cognitive model that predicts and explains human performance in Web-assisted tasks. The research was driven by the following questions:

- What are the most important factors determining success in Web-assisted tasks?
    - How can one measure or estimate these factors in an automated way?
- What are the explanatory cognitive mechanisms for the identified factors?
    - How can one implement these mechanisms in a (computational) cognitive model?
- What kind of Web navigation support can be conceived based on the knowledge gained from the two previous questions?
    - What impact has this support on users?

Chapter 2 presented the main behavioral and cognitive perspectives on Web navigation as reported in the literature of the field. *Chapter 3* discussed methodological issues involved throughout this research project: attention was paid simultaneously to theory, method and real-world applicability; Web navigation was grounded in the theories of Cognitive Science, in particular Text Comprehension, and Information Science, in particular Human-Computer Interaction; methodological paradigms of experimentation, statistics and modeling have been applied; the whole research was placed as much as possible in realistic settings and practical needs of Web engineering and design has always guided research decisions.

This research has documented how real Web applications are used (*Chapter 4*). A large number of users have been invited to the Usability Lab, where they have performed Web tasks of various sorts from several domains (Juvina & Van Oostendorp, 2003). Their behavior was recorded by various means: Web-logging, screen captures, video, audio, and paper-and-pencil. A large amount of data has been analyzed and interpreted in order to discover the regularities that can inform scientists and practitioners. The most important findings in this category are:

- Using the Web can be seen as a dialogue: users inform Web applications about their choices and Web applications "reply" with content. Thus, interaction paradigms such as reading/writing and talking/listening can be applied in understanding Web navigation behavior.
- In order to adequately characterize Web navigation behavior, a complete set of criteria is needed, including objective (performance), subjective (satisfaction) and undesirable aspects of task execution (disorientation).
- Users re-visit pages not only because they forget what they have seen before, but also to get acquainted with the structure of the information space, which in turn helps them in preventing disorientation.
- Using the Web via a screen reader by visually impaired persons is more taxing from a cognitive point of view, and users are more vulnerable to dissatisfaction and disorientation if they have to visit a large number of pages to reach their goal.

Observational and empirical research has revealed the most important user-related factors that determine performance in Web-assisted tasks (*Section 4.2*). A large number of factors have been analyzed together in relation to a comprehensive range of outcomes of Web navigation tasks. A limited number of significant predictors were identified, and their relative contribution to explaining task outcomes was estimated. Since factors were studied together and the stepwise method of regression analysis was employed, it was possible to rule out factors that were only marginally significant or confounded with one another. A sequence of repeated studies have shown that a combination of two factors is the most important determinant of human performance in Web-assisted tasks: a structure-related factor (spatial ability) and a content-related factor (domain expertise). Spatial cognition is involved in representing the structure of the information space while domain knowledge is necessary for understanding and selecting relevant content.

Factors such as spatial ability can be measured only with specialized tests, which cannot be implemented in realistic Web applications. For this reason, using Web-logging data to calculate metrics of Web navigation behavior has been proposed (*Section 4.1.2*). Metrics referring to the structure of user navigation were called syntactic metrics, whereas metrics referring to the visited content were called semantic metrics. By summarizing raw Web-logging data, such as use of navigation actions, page visits and re-visits, links followed and duration of visits, *first-order metrics* have been computed, such as:

- Path length
- Amount of re-visits
- Back button usage
- View time per page
- Fan degree
- Number of cycles
- Net density
- Average connected distance

Second-order metrics were computed as linear combinations of the first-order metrics by the aid of principal component analysis. They were completely specified (numerically) by first-order metrics. However, interpreting their meaning and labeling them was based on their correlations with user characteristics and task outcomes. The interpreted second-order metrics were labeled *navigation styles*. Two of these navigation styles are described below:

- **Flimsy Navigation** is a parsimonious navigation style. The navigation path was not very elaborate with most of the navigation taking place around the homepage. Time was spent more with processing content than with figuring out the hyperstructure that showed where the relevant information was.
- **Laborious Navigation** involves intensive use of navigational infrastructure provided by the site. Users seemed to employ a trial and error strategy. They followed links just to see if they were useful or not. They figured out quite fast when paths were not leading towards their goal and returned. Re-visits were quite numerous but they were not redundant: once a page was re-visited a different link was followed, it was just another trial.

A *semantic metric* called *Path adequacy* was calculated based on navigation data and the task description that subjects were provided with at the beginning of a task. A navigation path was considered to be a concatenation of semantic objects that the user has encountered in her/his way. Page titles and link labels have been used as semantic objects. Link labels were better than page titles in characterizing the user’s navigation path and computing semantic metrics because they convey more specific information. A navigation path was used in simulations of Web navigation as an indicator of contextual information involved in selecting specific navigation actions. Path adequacy was determined as a coefficient of semantic similarity between a navigation path and a task description. Semantic similarity was calculated with Latent Semantic Analysis (LSA). Path adequacy calculated at the end of a particular task was significantly correlated with *spatial ability* (r=0.36), and *task performance* (r=0.47). Path adequacy calculated at each step of a navigation session was used in simulations of Web navigation as a coherence criterion involved in selecting specific navigation actions.

Thus, different types of knowledge about users can be inferred based on the kind of information that is extracted from Web-logging data: *syntactic* (structural) information indicated mainly users' navigation styles, for example, if they re-visit pages rather than viewing new pages, if they return using the back button or just by following links, etc.; and *semantic* information indicated if users were effective in pursuing their goals independent of their navigation styles. These navigation metrics can be used in building adaptive Web applications, such as recommender systems (Juvina & Van Oostendorp, 2004).

Well-established theories of text comprehension (Kintsch, 1998), memory and cognition (Anderson, 1983), and working memory (Baddeley, 1986) have been used in the model development process. While using these theories to explain Web navigation behavior, specific aspects have been noticed. For example, in Web navigation a spatial representation of the information space is much more important than in the case of reading plain text. Users build and update a mental representation of the information space being navigated. This representation has a spatial character (in the sense that positions and distances are relevant, but not in the sense that is visual or three-dimensional), and it is relatively independent of the contents (semantics) being represented. Users make assessments of relevance and decisions to select particular contents based on both prior knowledge they have about those contents, and knowledge they gain from the local context of those particular contents (i.e., what contents they link to). Even in the case of using the Web by visually impaired persons (VIPs) spatial aspects are essential. For VIPs spatial aspects of Web use could refer to: temporal position of a particular information element in a sequence, category membership, etc.

A cognitive model (labeled CoLiDeS+) has been presented (*Section 4.3*) in which assessments of relevance are made based on both prior knowledge, modeled by an LSA semantic space (on the basis of the concept 'information scent'), and contextual information, modeled by users' past selections (on the basis of the concept 'path adequacy'). CoLiDeS+, an augmented version of CoLiDeS (Kitajima, Blackmon, & Polson, 2000), has been shown to account for important aspects of user navigation behavior such as: considering contextual information when judging goal-relevance and employing navigation strategies. This was accomplished by including the user's navigation path in the model and allowing the model to backtrack and reconsider past selections (e.g., next-best strategy). The model has been empirically tested for how well it simulates the actual user behavior and whether it is useful in generating Web navigation support. Although it does not simulate the user behavior particularly well, CoLiDeS+ was shown to perform better that its previous version (CoLiDeS). CoLiDeS+ was used to generate navigation support and this support had a positive impact on user behavior and task outcomes. A number of limitations of CoLiDeS+ have been identified, such as: low accuracy in simulating the users' behavior, caused mainly be its reliance on LSA to compute semantic similarities, and a high amount of hand-coding required for running simulations.

An ACT-R model of Web navigation has been presented (*Section 5.2*) in order to demonstrate the possibility of overcoming some of the CoLiDeS+ limitations. This model offered computational solutions to implement key features of Web navigation behavior as reported in literature and found in our empirical research. Some of these features were shared with previous models – selections based on goal relevance ('information scent') (SNIF-ACT and CoLiDeS); backtracking, threshold values and opportunistic strategies (MESA); and back coherence ('path adequacy, CoLiDeS+). Other aspects were implemented here for the first time – the intertwining between conservative and explorative strategies, and 'post-valued recall'.

It has been suggested that a cognitive model of Web navigation can be used as generator of Web navigation support, particularly when the model is specified in computational terms and can be run as a computer program (Juvina & Van Oostendorp, 2005). In a series of empirical studies, several ways to deliver model-based navigation support have been tested (*Sections 4.4, 4.5, and 5.1*). Suggestions of goal-relevant links via voice have been shown to increase task performance. In addition, users with low spatial abilities had a higher performance increase than users with high spatial abilities. It seems that the offered navigation support prevented users from spending time and cognitive resources on those navigating actions that are not directly effective but are usually employed in order to accurately represent the information structure. Users engage in apparently useless navigation actions in order to get acquainted with the context of a particular piece of information, which is eventually useful in judging the value of that particular information. It follows that users with low spatial abilities are probably less able to represent the information space and this is why they benefit considerably when the cognitive model is doing this job for them. However, suggestions via voice were not well received from a subjective point of view, users found them annoying and manipulative.

Graphical suggestions in the form of small red arrows pointing at the relevant link (Juvina & Herder, 2005) were not only effective but also well received from a subjective point of view. Men receiving support showed a decreased level of perceived disorientation as compared with men in the control condition, whereas such a difference was not found in women. Navigation support in the form of graphical link suggestions changed the structure of users' navigation behavior. In the support condition, participants used the back button less and the average connected distance in the navigation path was higher than in the control group. Thus, link suggestions caused the subjects to navigate in a more linear manner and reduced the amount of backtracking. High performers tend to take fewer suggestions than average and low performers. However, within each performance level, taking more suggestions is associated with increased task performance.

In the case of using the Web via screen readers, suggestions were implemented by changing the order of items on Web pages, in the sense that relevant items ('hyperlinks') were placed higher so that they are read sooner in a sequence of links. This manipulation was not successful, most probably because changing the order of items on Web pages breaks the coherence established by the content authors.

This research does not support extracting the relevant information from its original context as in the case of search engines. Instead, it has been shown that emphasizing the relevant information in its original context helps users discern between relevant and irrelevant information, and compensates for their deficient spatial abilities (when this was the case for the user).

Performing ecologically valid research was an important desideratum of this project; real websites and realistic tasks have been used as much as possible (*Chapter 6*). A trade-off between modeling accuracy and practical relevance has guided the research. CoLiDeS+ was conceived with the aim of building model-based navigation support. For this reason, some of the methodological criteria of cognitive modeling were relaxed. The simulation of user behavior was not complete. For example, the model did not have a mechanism to identify target content pages. Such a mechanism would have been extremely difficult to build and it was not necessary for the intended use of the model. The model was meant to work alongside the user and suggest links that are relevant to a given user goal. The user was assumed to take those suggestions or not and stop when the target page was reached. Performance of cognitive models in the field of Web navigation depends on progress made in other fields such as machine learning and natural language processing. For cognitive modeling, working in a task domain as weakly structured and knowledge-intensive as the one proposed here was a great challenge. A cognitive model of Web navigation needs to handle natural language, large knowledge networks, and a great deal of sub-symbolic computations. These aspects are not part of the traditional work in cognitive modeling research (Gluck & Pew, 2005) but are becoming increasingly prominent in the cognitive modeling community (Pirolli, 2005). A foreseeable issue for future research will be handling the computational complexity required by up-scaling cognitive models to be included in adaptive Web applications as generators of navigation support.