In-Context Annotations for Refinding and Sharing

Ricardo Kawase, Eelco Herder, George Papadakis, and Wolfgang Nejdl

L3S Research Center, Leibniz Universität Hannover
Appelstr. 4, 30167 Hannover, Germany
`{kawase,herder,papadakis,nejdl}@L3S.de`

Abstract. Annotations support understanding, interpretation, sensemaking and scannability. As valuable as in paper-based contexts, digital online annotations provide several benefits for annotators and collaborators. To further explore the real benefits of online annotations, we implemented a simple Web Annotation tool, SpreadCrumbs, to support our studies. The tool provides a simple annotation mechanism, simulating real-world paper-based annotations. In addition, the tool supports search, sharing capabilities and social navigation. We conducted a series of user studies that empirically demonstrates the benefits of "in-context" annotations for refinding and sharing.

Keywords: Annotation, In-context, Trail, Social media, Social network, Online collaboration, User interface, SpreadCrumbs.

1 Introduction

The World Wide Web is arguably the biggest source of information nowadays. Whereas the exchange of ideas on the Web was predominantly one-way, the Web 2.0 now offers a new means of interactions and has shifted more power and influence to users. However, there are still a number of features missing that are essential for supporting information classification, retrieval, processing and understanding. Most of these issues have been reported already during the early inception of the Web, mainly from the hypertext community [1] [2]. In particular, frequently mentioned are: the lack of typed or annotated links; the absence of hypertrails; limited browser history mechanisms; and the lack of support for annotations.

In order to bring these missing features into the Web, a common workaround is to create applications that enhance the Web usability, such as search engines, tagging systems, annotation systems, social networks and others. The competitive character within the Web 2.0 has arguably led to a more powerful reincarnation of the rich features that once were part of the classic hypertext systems [3]; albeit as a collection of diverse, disconnected applications, interoperating on top of a common Web platform. Surprisingly, despite the prevalence of interactive applications and social networking, thus far Web annotation systems haven't seen a significant take-up [4].

Given the absence of any dominant mature annotation system, it appears that there is still no generally accepted, concrete method for straightforward online annotation. This is surprising, given the abundance of literature showing the importance of annotations for comprehension and their benefits for reading and writing proposes [5].

Similar to the paper-based environment, digital annotations are expected to be useful for supporting comprehension and interpretation [6]. Moreover, comments and references are known to stimulate associative thinking, which can be even better reproduced digitally, by what we call "hypertrails". For this reason, our research goal is to understand users' annotation behaviors and identify the benefits and drawbacks of online annotations and trails.

Based on insights gained from earlier work and an analysis of the reasons that hampered wide-spread adoption of earlier annotation systems, we created SpreadCrumbs [7, 8]. SpreadCrumbs is an online annotation tool that allows the users to place annotations within Web resources, either for themselves or for other users. In this work we introduce the application, its main functionalities and present a system evaluation.

The rest of this work is structured as follows: In section 2, we discuss related work in the field of annotations and annotations systems, followed by the description of Spreadcrumbs in section 3. In section 4 we present a preparative study on annotations for the experiments and studies, of which the setup is described in section 5. The results are presented in section 6, followed by our conclusions in section 7.

2   Related Work

2.1   Paper Annotations

We adopt the definition of annotations as set forth by MacMullen [9] and Marshall [10]: *An annotation is any additional content that is directly attached to a resource and that adds some implicit or explicit information in many different forms.*

Annotations may serve different purposes, such as: signaling a foreshadow, aiding memory and interpretation or triggering reflection. Additionally, annotations may occur in many different forms; for example: by highlighting, encircling or underlining text, we emphasize the importance of a certain part of the document; a strikethrough indicates that something is wrong, misplaced or not relevant; arrows signal relations between two or more elements.

In [3], the authors draw a comparison between the early Hypertext pioneers visions and the present-day Web applications, commonly known as Web 2.0. The results of their analysis show that most of these systems support both private and public annotations and provide support for collaboration. Even though these features are identical with the first ideas of the Hypertext, the annotations are limited, because they reside exclusively bound to individual Web 2.0 services providers and they are not "in-context" – More specifically, they are not visualized together and associated with the annotated content (the topic of interest), which the benefits will be exposed later.

2.2   Digital Annotation Systems

The Fluid Annotations projects [11] introduce an online annotation system that supports in-context annotation in an extension of the open hypermedia Arakne Environment [12]. Their studies focused on evaluations and the presentation of the annotations in terms of visual cues, interactions, accommodation and animated

transactions. Their main approach to in-context notes uses between-lines annotations. Their evaluations give valuable insights into the usability and manipulation of annotations. Nevertheless, we believe disrupting the original layout of the annotated content may be more confusing and disruptive than beneficial.

Another annotation system is MADCOW [13, 14] a digital annotation system organized in a client-server architecture, where the client is a plug-in for a standard web browser allowing users to annotate Web resources. Although MADCOW supports different representations for annotations, previous work comparing paper-based and digital annotations [15] suggests that paper-based annotations should not be mimicked by similar representations but by providing the means to achieve the same goals. In addition, the placeholders of the annotations are inserted between the HTML content which can be disruptive, distractive and may lead to the problem of orphan annotations. Finally, usage complexity will impact the dissemination of any new technology, and in particular, will always be an obstacle for the non engaged users. The annotation interface in their work has not been evaluated.

A more full-fledged annotation tool is Diigo . Using the Diigo toolbar, users can highlight text or attach 'inline sticky notes' to Web pages. Despite the wealth of features, Diigo cannot boast a big user population. According to online user comments, this is due to both usability issues and the fact that all annotations are public by default. We understand that sharing annotations is one of the main possible advantages of digital annotations systems; however in light of Diigo, we believe that a 'shared' annotation must not be mistaken for a 'public' one. The benefits of reliable collaborators are not fully applicable in the 'public' scenario; we elaborate, further on this point in sub-section 2.3.

In summary, there are numerous and similar annotations systems - most of them are discontinued works which have neither developed further nor been presented in further studies.

2.3   Social Navigation

Social navigation support (SNS) describes techniques for guiding users through specific chosen resources [16]. In AnnotatEd [17] the authors introduce two types of SNS: traffic-based and annotation-based. Our model more is related to the annotation-based style, in that every annotated page becomes a step in a trail.

Annotation-based social navigation support has been shown to be more proficient and reliable than traditional footprint-based social navigation support [18]. When the annotated resource reflects the interest of the annotator, it appends more value to the SNS. Annotation based SNS assists users in gathering information by making it easier to re-access the information and by showing the collective wisdom of the collaborators.

Allowing users to "attach" their personal insights to a resource increases the reliability of annotation-based navigation support. Previous study of annotation-based SNS shows that users are particularly interested in being informed about resources annotated by others. Annotated resources are significantly more likely to be visited by users, specifically after being annotated [18].

3   SpreadCrumbs

SpreadCrumbs is an in-context Web annotation system, which has been implemented as an extension of the Mozilla Firefox Web browser. The underlying assumption of SpreadCrumbs is that users can annotate Web resources with keywords or sentences and create hypertrails through a set of annotations. These annotations can not only be used for one's own reference, but can also be shared within a social network. The design of SpreadCrumbs has deliberately been kept minimalistic. Following the approaches seen in related work, we chose the basic visual metaphor for the annotations: Post-it notes.

The Post-it representation has an optimized approach to simulate the most common paper based annotations forms namely underlining, highlighting and notation in margins. The idea is not to mimic different representations but to provide a way to achieve the same goals: signalling for future attention, comprehension and summarization. In addition post-it notes are extremely efficient as "in-context" landmarks which are the main purpose of the research.

Furthermore, by bringing the annotation behaviour to the digital online environment we also add valuable features that are not applicable in the paper-based scenarios. The most prominent are the re-finding and the social sharing possibilities. The content of an annotation is easily searchable within the tool and shareable with other users.

3.1   The Browser Add-On

The SpreadCrumbs Browser add-on is a Javascript implementation based on AJAX principles. We used the AJAX and Javascript library from Yahoo, The Yahoo! User Interface Library (YUI). The library provides functionalities for drag & drop and other manipulations used in SpreadCrumbs. A simple client server architecture stores all the data on the server providing the user the possibility to access her data anytime from any computer where the client application is installed.

Once the client add-on is installed to the browser the user can access the sidebar. Through the sidebar the users have access to straightforward ordinary actions like creating account, profile management, login and logout. Additionally, the user has direct access to a contact managing webpage and a tabbed annotation-browser-window. From the right-click context menu an option is available to annotate the page, the same as from a small annotation button near the address bar.

3.2   Networking

As a non-mandatory step, new users may add their social network contacts to become collaborators in SpreadCrumbs. From the sidebar the users have access to the 'contact manager' webpage, from which they can import their contacts from their Facebook Network using Facebook Connect technology. Once the contacts are imported they become part of the user's SpreadCrumbs network and the user is able to share annotations with her contacts. If at some point these contacts join SpreadCrumbs and grant permission to Facebook Connect; their accounts will be synchronised and all the annotations previously shared by some other user will be retrieved.

3.3 Annotating

Annotations (which we will refer to as 'crumbs') are added via the right-click context menu by the option "Add Crumb", which results in the opening of a pop-up window that contains three fields: the receivers of the annotations, a topic and the content. By default, annotations are private. An auto-completion drop-box helps the user in adding receivers from her contact list.

Once the annotation is created, a post-it note appears in the screen, originally on the clicked spot but easily relocated by drag & drop (Fig. 1). When any of the involved users in the annotation accesses the annotated website, post-it note will be displayed. Additionally, if the user keeps her connection to Facebook through SpreadCrumbs, the receivers of the annotation will get a notification on Facebook and a notifying e-mail about the new annotation.
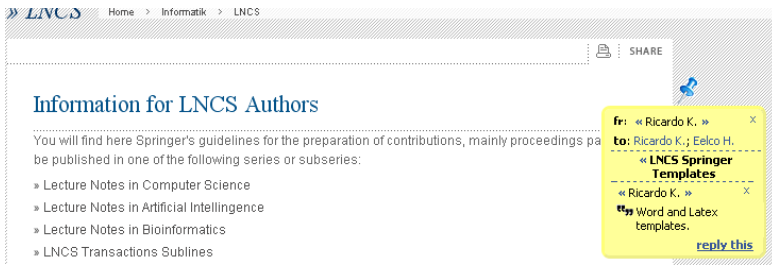
Fig. 1. Page annotated with SpreadCrumbs

3.4 Reacting

Each annotation is an entity in a thread (a crumb in a trail) and diverse actions can be taken over it. When visualizing an annotation, any of the involved users has the ability to interact with it: moving it around, closing it, following trails and replying.

Connect and Disconnect. Each user has her individual status in the context of one annotation. The status "Connected" is the normal status to visualize the annotations; "Disconnected" means that she will not visualize the annotation anymore once she comes back to the website; and "Stand by" means that she will not visualize the annotation again until some modification has occurred in the annotation thread.

Replying. The reply link on an annotation brings up the same window pop-up as adding an annotation offering to the user just the content field to be filled. Once confirmed, the reply is attached to the first post-it note and the same notifications actions are triggered. Any user involved in the annotation is able to add a reply to the running thread, which is visible to all participants. This action simulates a micro in-context forum on each annotated web page.

Following Trails (SNS). What makes SpreadCrumbs unique is that the annotated pages are not simply a loose collection, but the resources become interconnected. Each annotation is associated with links that can be followed from the crumb: the user trail and the topic trail. Near the name of each user who annotated the page and near

the topic text there are two small linked arrows indicating the path to the previous and to next annotation in the hypertrail. Following the previous/next link next to the name of a user will redirect the current user to the next/previous annotated page where both users share another annotation.

Following the topic trail will lead the user to web pages on which the user has annotations with the same topic description. A simple illustrative example: one user privately annotates five different pages with the topic "Conference" adding specific content for each annotation. Once it is done, each conference page annotated has a link connecting to each other. A temporal defined (and connected) collection of web resources was created and at any time the user is able to remove, edit or add new stop points in this trail. The final output is a simulation of the Memex idea where the resources are now annotated and associated in accordance with the user's preferable organization.

Providing sharing capabilities of these trails, SpreadCrumbs grants Social Navigation Support in a very concrete and defined manner. Differently from others SNS systems, the resources are not only a collection of links but they have a well-defined temporal order, each resource becomes interconnected and they hold in-context insights from the annotation authors.

3.5 Browsing Annotations

The SpreadCrumbs' sidebar contains a browser pane with three different tabs that shows the three facets of the organizational dimensions of a trail: topics, pages, people. Additionally, a small pane in the bottom shows detailed information on the selected trail.

The tab 'topics' shows the trails grouped by topic description. The user visualizes distinct items that represent the different trail-topics she created. From this pane, the user is able to access the annotated page, edit the topic description and change her status in the topic. By clicking or selecting one of the topic-trails the bottom pane loads and displays all the crumbs belonging to this trail assembled by page. In this pane the user has the same possibilities to directly access the annotated page, to edit the crumb and to reply it.

The second tab, page, shows the trails grouped by the resource annotated. The visualization has the title extracted from the Webpage and the trail last modified date as well. The user has the possibility to edit the name of the page, if she wants to. It is important to notice that although trails mainly contain the same page title in this facet they will not be grouped together since the grouping is based on the URL location of the annotation. By clicking or selecting one of the page-trails the bottom pane loads and displays all the crumbs belonging to this trail assembled by the different topics existing on the selected page with same management capabilities.

Finally, the people tab shows items that represent the trails from the user's contacts. The item visualization shows the name of the contact and her last activity on the trail. It also indicates whether the contact is already connected to SpreadCrumbs' network or not (due to the fact that is possible to share annotations to imported contacts that are not subscribed to SpreadCrumbs). By clicking or selecting one of the people-trails, the bottom pane works in the way as the topics tab previously described.

4   Understanding Annotations

The first step in our studies is to fully comprehend the desired annotation features needed on the web. In order to understand the real use of annotations and Web annotations we conducted a field-study examining the paper-based annotations of 22 PhDs students and pos-Docs in their own work environment [15]. For each participant, we looked at the last 3 research papers or articles that they have printed and read. In total we have collected 66 articles, covering a total of 591 pages of text. This is a mere preparative step to the incoming experiments presented on Section 5.

We found 1778 annotations and an average of 3.08 annotations per page. **Table 1** below shows the average of each type of annotation based on Marshall's proposed classification [10] by forms and functions.

Although most of the annotations consist of highlighting activities, we identified in our previous study that it does not imply that mimicking this feature is the most appropriate approach to be followed. We identified that paper-based highlights are used for signalling and attributing different levels of importance and to help memorization during the reading activity. However, digital highlight is usually a non-persistent activity to help focusing on the text and re-finding – users highlight the text with the mouse cursor while reading. Excessive amounts of digital highlighting turns out to be more distractive than helpful. The conclusion of this work led us to the consideration that annotation systems should emphasize re-finding, visual overviews, grouping, sharing and collaborating rather than to try and mimic the 'old-fashioned' paper-based annotation.

Table 1. Collected annotations classified by type

Annotation types		
Highlighting/Mark sections headings	153	8.6%
Highlighting/Mark text	1297	73%
Problem solving	2	0.1%
General notes (Notes in the margins)	326	18.3%

5   Web Annotations in Practice

The next step after understanding annotations is to observe and analyze Web annotations in practice. To evaluate the usability and performance of Web annotations, we ran a series of laboratory experiments and processed the usage logs. The aim of our experiment was threefold: 1) examine the possible benefits of annotations over bookmarks for personal information management, and 2) evaluate social navigation support in an arbitrary scenario. In this section we describe the experiments.

5.1   Annotations vs. Bookmarks

The goal of our first study was to quantitatively estimate the efficiency and ease-of-use of annotations as a means for personal information management and refinding. As a reference point, we compared Annotations with (Social) Bookmarks. Our participant pool consisted of 24 males and 10 females, with an average age of 28. The participants were randomly and equally split into two groups: the first group created annotations using the Delicious social bookmarking service, the second group made use of SpreadCrumbs.

After a short introduction to the basic features of the tool (either SpreadCrumbs or Delicious), each individual session was conducted. We asked the participant to find answers for ten random questions. This task was presented as 'just' an exercise in order to get used to the system. In reality this first task was a preparational step for the second round of the experiment. All questions were specific information-finding tasks that could be solved by a brief internet search with any popular search engine. We ensured that the questions were sufficiently obscure, to minimize the chance of participants knowing the answers themselves.

Five months after the initial round of the studies, the participants were invited to participate again. This time, their task was to relocate the answers that they had previously found during the first task. The long time interval ensured that the participants remembered neither the answers they had provided nor the resources they had used to find the answers. In total, 30 out of the initial 34 participants were involved in this phase of our study (21 males and 9 females, average age 28 years).

The participants were divided into three equivalent groups of 10 people, each one corresponding to a specific refinding methodology and corresponding tool. As a base line, the first group used a search engine in their efforts to carry out their tasks (in other words, they had to search again for the same information). The second group used bookmarks to refind the information. This group consisted of those subjects that used Delicious in the previous session and had the URLs of the visited resources at their disposal. The third group consisted of the SpreadCrumbs users. The members of this group had the in-context annotations at their disposal.

We ensured that all participants accomplished all of their tasks under the same conditions and that their performance is compared on an equal basis. After the appropriate Web resource was found, thus completing the 'searching stage', the participant had to locate the answer in the page and highlight it using the mouse – the browsing stage. There were no instructions or restrictions on how to proceed at this stage: the participants were allowed to perform this task the way they would in a non-controlled environment. The necessary data for estimating and evaluating the average and overall browsing time per individual were collected using screen capture and data-logging software that recorded all participants' actions.

5.2   Shared Annotations

To evaluate the usability and benefits of annotations we asked the same 34 participants from the previous study to play a role in a scenario on collaborative decision making. The participants were asked to plan a trip to London, by reviewing the options, as

collected by their 'partner' (the experimenter). Via either SpreadCrumbs or Del.icio.us, the participant received a number of annotations/bookmarks on suitable hotels, restaurants, museums and musicals in London. The participants evaluated the given options – by visiting the bookmarked sites and/or by reading the annotations – and finally decided for one option in each category.

Upon completion of all tasks, the subjects were asked to fill out two questionnaires, one regarding the information refinding experience and another one investigating their opinion on the tool they used.

6 Results

From the refinding task we collected a total of 297 successful activities, evenly distributed across the conditions. With this data we extend our analysis [19], investigating further aspects that may have contributed to the results. First, the most appropriate metric for expressing the overall performance of each group is arguably the *average tim*e taken to complete the *browsing phase* – therewith ignoring the time it took participants to locate the page in the *searching phase*. In our case, the available sample of 99 browsing times produces the following mean values: 46s for Search Engine, 38s for Bookmark, and 21s for Annotation. It turns out therefore that the performances of the first two groups differ slightly, whereas the performance of Annotation is substantially better, corresponding to a time that is almost the half of the other two groups. This suggests that in-context annotation boosts refinding to a great extent.

This fact is also advocated by the outcomes of the corresponding independent two-tailed t-tests between groups Search Engine and Annotation (Table 2B) as well as between Bookmark and Annotation (Table 2C). Both tests produce a significant result ($p < 0,01$) with a medium effect size *r*.

By contrast, when comparing the performance of the first two groups, the value resulting from t-test is well above the threshold (Table 2A). This result does not match our initial expectation that Bookmark would outperform Search Engine due to the wealth of cues associated with them - the comments that were attached to bookmarks as well as the keywords of the tags that were drawn from the questions or even the answers. However, this equivalence can partially be explained by the *theory of context-dependency* [20]. This theory states that all context knowledge acquired in the refinding process, including even the non-semantically related elements located within the target information, serves as relevant cues for refinding information. According to the theory, search engine users acquired a comprehensive enough context while searching and browsing the search results, which assisted them in the browsing stage, whereas the bookmark users had to acquire the context during the browsing stage itself. It should be stressed that the performance of the search group would have been significantly worse if we also took into account the searching stage, which is minimized for bookmarks' users.

Table 2. Results of t-test between Search Engine and Bookmark (A), between Search Engine and Annotation (B) and between groups Bookmark and Annotation (C)

	A		B		C	
	Search Engine	**Bookmark**	**Search Engine**	**Annotation**	**Bookmark**	**Annotation**
Mean	46:30.9	37:45.5	46:30.9	20:41.2	37:45.5	20:41.2
Variance	02:19.5	00:57.8	02:19.5	00:17.4	00:57.8	00:17.4
Observations	99	99	99	99	99	99
df	98		98		98	
t(98)	1,21		4,07		3,88	
p	0,11		< 0,01		< 0,01	
r	0,12		0,38		0,36	

The increase in performance of annotation users is also verified by the distributions of the browsing times of each group, which are presented below in Fig. 2A. This diagram depicts a well formed normal distribution corresponding to Annotation that has its mean close to the beginning of the x axis. On the other hand, Search Engine and Bookmark exhibit two almost coinciding distributions that are positively skewed and have their means located further away from the y axis. The skew shows that annotations have particularly been useful for some specific questions – for the majority of questions the browsing time is more or less similar.
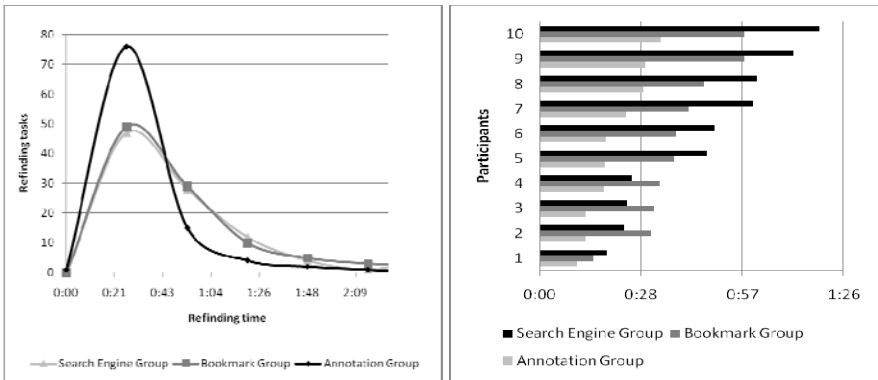

(a)   (b)

Fig. 2. Distribution of refinding tasks by time (A). Average browsing time (x-axis) per participant in ascending order (B).

Finally, in Fig. 2B. we illustrate the average browsing time of each participant for each one of the three groups. Participants are actually presented in ascending order, starting from the one with the shortest time and ending with the least efficient one.

This diagram makes clear the fact that the poor performance of Search Engine and Bookmark is not caused by a few, slow performing participants. Conversely, the majority of them had large average browsing times, in contrast to the users of SpreadCrumbs that in general relocated the required information quite quickly.

Thus far we have focused on the effect of the diverse tools on users' efficiency, thus ignoring other important factors, like the size of the visited pages as well as the use of browsers' "find" functionality. Due to their significance, we have devoted the following two subsections to their analysis.

6.1   CTRL+F

As it has already been mentioned, the browsers' "find" (CTRL+F) functionality plays a major role in relocating a specific piece of information within a web page. In order to quantify its degree of use, we measured the percentage of tasks of each group that were carried out with its help and present the outcomes in Fig. 3A. As expected, it is evident that CTRL+F has been extensively used by the subjects of Search Engine and Bookmark, whereas participants of Annotation resort to it to a lesser extent. They actually use it solely in the cases of modified web pages that result in misplaced or orphan annotations, as in these cases the annotations are of little help and the user has to resort to other means for pinpointing the desired information.

Judging from the wide use of the CTRL+F strategy, it is reasonable to assume that "find" helps participants to perform better in refinding information. To verify this assumption, we estimated the average browsing time that corresponds to subjects using it and compared against that of those that did not use it. This comparisons were made in the context of all three groups, and their outcomes, presented in Fig. 3B., suggest the opposite: participants that took advantage of this functionality needed *significantly* more time in completing their tasks than those that did not. Hence, although this functionality is supposed to constitute a quite handy tool for locating information, in practice there is no evidence supporting its beneficial contribution to re-visitation efficiency. A likely explanation, which we investigate further in the next subsection, is that CTRL+F is mainly used when users encounter pages that do not allow for quick visual scanning due to their length.
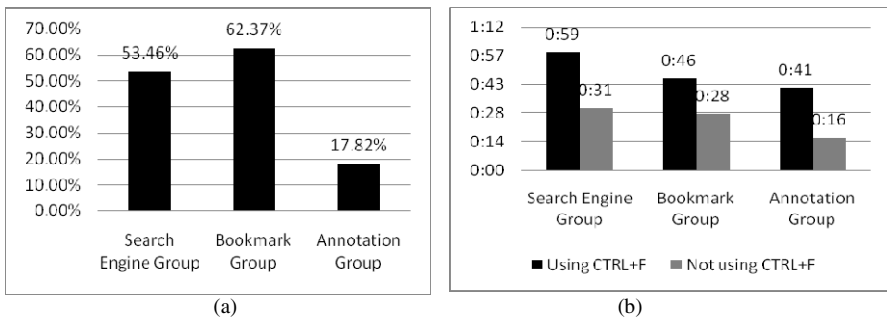

(a)   (b)

Fig. 3. Percentages of usage (y-axis) of the browser's "find" functionality of each group (A). Average times (y-axis) of each group distinguishing tasks where the browser's "find" functionality was used (B).

In the following table (Table 3), we briefly present the results of an independent t-test that show the significant differences in the performance between the subjects that do and those that do not use the "find" functionality. A more detailed investigation of the correlation between this functionality and the size of the web pages follows in the next subsection.

Table 3. Results of t-test between the tasks which had and hadn't the assistance of the "find" functionality for each group

Search Engine Group	t(94)=2.54, p<0.01, r=0.25
Bookmark Group	t(84)=2.54, p<0.01, r=0.27
Annotation Group	t(18)=3.08, p<0.01, r=0.59

6.2   Page Content Size

As mentioned in the previous subsection, there is likely to be a correlation between the length of web pages and the usage of CTRL+F as a means for refinding. To investigate this correlation, we divided the participants of each group into two subgroups; one that used CTRL+F, and one that did not. For each of these subgroups we estimated the average size of the accessed web pages in terms of number of words and calculated the average browsing time – see Fig. 4. The figure shows that – in particular in the Bookmark and Annotation groups – there is no interdependency between page size and the usage of CTRL+F. However, it shows that the browsing time is significantly higher in the CTRL+F condition, which suggests that the find functionality does *not* sufficiently leverage the detrimental effect of long and possibly unstructured pages.
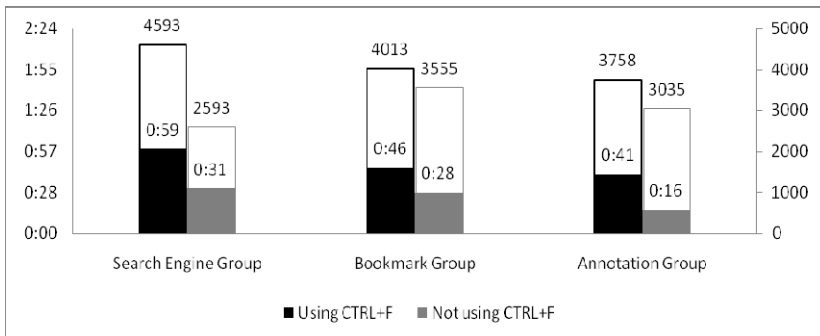

Fig. 4. Average times (y-axis left) of each group and average Web page sizes (number of words) (y-axis right) distinguishing tasks where the browser's "find" functionality was used

In order to examine the effect of page size on browsing time, we clustered the pages that were used by our subjects according to their number of words and estimated the average time related to each cluster. In Fig. 5 we present the performances associated with each application and page cluster. It is obvious that no

trend can be identified for Search Engine and Bookmark, as they depict a high variance, whereas Annotation has a more stable but still volatile behavior. By manually observing the structure of the visited web pages, we came to the conclusion that the size of the page has no significant influence on the browsing stage, as opposed to the web page structure. This is in line with our qualitative analyses, presented below, whose outcomes indicate that, in the absence of in-context annotations, a well structured content together with an elegant layout, and a nice organization that is rich in anchors; substantially assist users in refinding information.
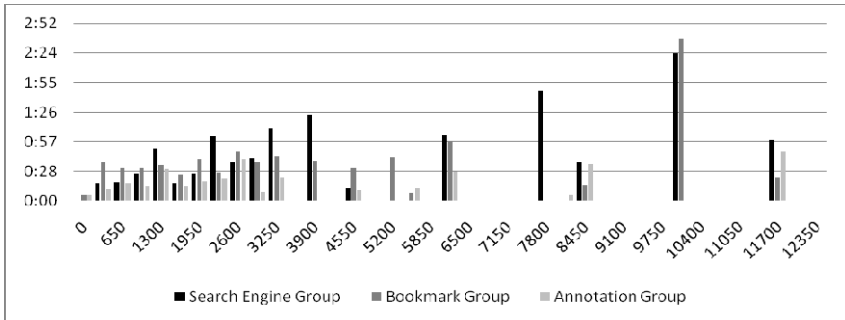

Fig. 5. Average times (x-axis) of each group by Web page size (y-axis)

6.3   Shared Trails and Annotations

In the second part of the studies (described in Section 5.2), 50% of the users who received the suggestions from their 'partner' via Delicious did not read or even did not notice the additional comments on each bookmark, which were displayed just below the page title and the URL. One participant explicitly told us that she noticed them only in the middle of the task. Another participant said that she noticed the comments, but did not read all of them because she thought they were irrelevant.

By contrast, all the participants who received the suggestions via SpreadCrumbs indeed noticed and read the comments, which were displayed as post-it notes. They all accessed the bookmarked pages and read the shared comments in the context. During the interview after the task, some of them confirmed that their choices were influenced by those comments.

The results show that if annotations are meant to provide additional information and to influence the receiver's opinion or choices, they should be presented as such, in the context. A text snippet below the title, as provided by many social bookmarking sites, is clearly not sufficient to catch the receiver's attention and may be overlooked during a collaborative knowledge building process.

6.4   User Feedback

After completing the set of tasks, each participant was asked to fill out a questionnaire, with the aim of distilling opinions on the tool used as well as the experiments in general. The answers were given by selecting the appropriate value on

a 7 point Likert scale. The user experience survey consisted of 13 questions, taken from established surveys on user satisfaction, frustration and disorientation. The Crombach *a* of 0,762 indicated a good reliability and the results were grouped nicely into the three factors.

Participants from the Bookmarking and the Annotating group reported less frustration than participants from the Search group. Further, the participants from the Annotation group reported a marginally significant lower value of difficulties in finding the right information. Whereas most other questions did not result in significant differences in answers, the overall trend indicated positive effect of bookmarking - and of annotation in particular - on the subjective user experience.

It is also worth mentioning that five participants of the annotation group marked the same page, a page that had been changed during the time interval between the first and the second session of the study. As a result, the annotations they had posted were misplaced in all the five cases, which caused a slight delay in the refinding task. Two of them suggested a more intuitive way of attaching annotations that involved arrows. Even though this way could well solve the issue of misplaced annotations, it will still be of no help for orphaned ones, which is in the cases where the annotated information has been completely removed. This issue is actually considered as one of the most complicated and challenging problems of the in-context annotation approach [21, 22].

7 Conclusions

In this work, we presented the SpreadCrumbs Web annotation tool and demonstrate how it is able to overcome the limitations of previously existing annotation mechanisms. With SpreadCrumbs, users can place Post-it-like notes at any location of a Web page. From our user studies and a literature survey we identified that user needs for making annotations in the Web environment do not differ significantly from their needs in the paper environment [23]. In addition, Spreadcrumbs supports different user tasks, not only private annotations, but also personal reminders, refinding enhancement, and social bookmarking with extensive support for social navigation and collaboration.

We also presented empirical results that show the important role of annotations in the digital environment, the outperformance of in-context annotations over bookmarks in terms of supporting information refinding, the analysis and the impact of in-context annotations on social and collaborative scenarios and finally the usability and users' opinion feedback.

Although we have seen the importance and benefits of annotations, no annotation system is widely adopted. This implies that there are still several issues to be studied and solved. The main challenge for annotation systems is on the user interface level. It is necessary to balance the classic tension between full-fledged features and ease of use. Particular attention should be paid to the question to what extent annotation systems should provide and emphasize social bookmarking features.

References

1. Halasz, F.G.: Seven issues: Revisited. Closing plenary address. In: Proceedings of ACM Hypertext Conference, San Antonio, Texas, (December 18,1991)
2. Vitali, F., Bieber, M.: Hypermedia on the web: What will it take? ACM Computing Surveys 31(4es) Article No. 31 (1999)
3. Millard, D.E., Ross, M.: Web 2.0: hypertext by any other name? In: Proceedings of the Seventeenth Conference on Hypertext and Hypermedia, HYPERTEXT 2006, Odense, Denmark, August 22 - 25 (2006)
4. Karger, D., Katz, B., Lin, J., Quann, D.: Stickey notes for the semantic web. In: Proc. Intelligent User Interfaces (2003)
5. O'Hara, K., Sellen, A.: A Comparison of Reading Paper and On-Line Documents. In: Proceedings of the ACM Conference on Human Factors in Computing Systems, CHI 1997 (1997)
6. Marshall, C.: Toward an Ecology of Hypertext Annotation. In: Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia, Hypertext 1998 (1998)
7. Kawase, R., Nejdl, W.: A Straightforward Approach for Online Annotations: SpreadCrumbs - Enhancing and Simplifying Online Collaboration. In: WEBIST, pp. 407–410. INSTICC Press (2009)
8. Kawase, R., Herder, E., Nejdl, W.: Annotations and Hypertrails with SpreadCrumbs: An Easy Way to Annotate, Refind and Share. In: WEBIST 2010: Proceedings of the 6th International Conference on Web Information Systems and Technologies (2010)
9. MacMullen, W.J.: Annotation as Process, Thing, and Knowledge: Multi-domain studies of structured data annotation. SILS Technical Report TR-2005-02. UNC School of Information and Library Science (2005)
10. Marshall, C.: Annotation: From Paper Books to the Digital Library. In: Proceedings of the ACM International Conference on Digital Libraries, DL 1997 (1997)
11. Zellweger, P., Mangen, A., Newman, P.: Authoring fluid narrative hypertexts using treetable visualizations. In: Proceedings of ACM Hypertext (2002)
12. Bouvin, N.O.: Unifying strategies for Web augmentation. In: Proceedings of ACM Hypertext 1999, pp. 91–100 (1999)
13. Bottoni, P., Civica, R., Levialdi, S., Orso, L., Panizzi, E., Trinchese, R.: MADCOW: a multimedia digital annotation system. In: Proceedings of the Working Conference on Advanced Visual interfaces, AVI 2004, Gallipoli, Italy, May 25 - 28 (2004)
14. Bottoni, P., Levialdi, S., Labella, A., Panizzi, E., Trinchese, R., Gigli, L.: MADCOW: a visual interface for annotating web pages. In: Proceedings of the Working Conference on Advanced Visual interfaces, Venezia, Italy, May 23 - 26 (2006)
15. Kawase, R., Herder, E., Nejdl, W.: A comparison of paper-based and online annotations in the workplace. In: Cress, U., Dimitrova, V., Specht, M. (eds.) EC-TEL 2009. LNCS, vol. 5794, pp. 240–253. Springer, Heidelberg (2009)
16. Brusilovsky, P.: Adaptive hypermedia. User Modeling and User Adapted Interaction 11 (1/2), 87-110 (2001); Claypool, M., Le, P., Wased, M., and Brown, D Implicit interest indicators. In: Claypool, M., Le, P., Wased, M., Brown, D. (eds.) Proceedings of 6th International Conference on Intelligent User Interfaces, pp. 33–40 (2002)
17. Farzan, R., Brusilovsky, P.: AnnotatEd: A Social Navigation and Annotation Service for Web-based Educational Resources. In: Proc. of E-Learn 2006, AACE, Honolulu, HI, USA, October 13-17, pp. 2794–2802 (2006)

18. Farzan, R., Brusilovsky, P.: Social navigation support through annotation-based group modeling. In: Proceedings of 10th International User Modeling Conference, pp. 463–472 (2005)
19. Kawase, R., Papadakis, G., Herder, E., Nejdl, W.: The impact of bookmarks and annotations on refinding information. In: Proceedings of the 21st ACM Conference on Hypertext and Hypermedia, HT 2010, Toronto, Ontario, June 13 - 16 (2010)
20. Pitkow, J.E., Kehoe, C.M.: Emerging trends in the WWW user population. Commun. ACM 39(6), 106–108 (1996)
21. Cockburn, A., McKenzie, B.: What do Web users do? An empirical analysis of Web use. Int. J. of Human-Computer Studies 54(6), 903–922 (2001)
22. Wang, S.: Annotation Persistence Over Dynamic Documents. Doctoral Thesis. Massachusetts Institute of Technology (2005)
23. Fu, X., Ciszek, T., Marchionini, G.M., Solomon, P.: Annotating the web: An exploratory study of web users needs for personal annotation tools. In: The 68th Annual Meeting of the American Society for Information Science & Technology (ASIS&T), Charlotte, NC, USA (2005)