

Note to Self: Examining Personal Information Keeping in a Lightweight Note-Taking Tool

Max Van Kleek, Michael Bernstein

Katrina Panovich, Greg Vargas, David Karger

MIT CSAIL, 32 Vassar St. Cambridge MA 02139

{emax, msbernst, karger, kp, gvargas} @
csail.mit.edu

mc schraefel

Electronics and Computer Science

University of Southampton

Southampton, UK, S017 1BJ

mc+chi @ ecs.soton.ac.uk

ABSTRACT

This paper describes a longitudinal field experiment in personal note-taking that examines how people capture and use information in short textual notes. Study participants used our tool, a simple browser-based textual note-taking utility, to capture personal information over the course of ten days. We examined the information they kept in notes using the tool, how this information was expressed, and aspects of note creation, editing, deletion, and search. We found that notes were recorded extremely quickly and tersely, combined information of multiple types, and were rarely revised or deleted. The results of the study demonstrate the need for a tool such as ours to support the rapid capture and retrieval of short notes-to-self, and afford insights into how users' actual note-taking tendencies could be used to better support their needs in future PIM tools.

Author Keywords

Note-taking, personal information management, information scraps.

ACM Classification Keywords

H5.2. User interfaces: user-centered design.

INTRODUCTION

Despite the sophisticated personal information management (PIM) tools available on our computers today, in reality, many people still rely on Post-it notes, disorganized *todo.txt* files, and even random scraps of paper with barely legible notes-to-self, to maintain valuable information [2]. Recently, several new classes of PIM tools have emerged to help organize this accumulation of personal information. In particular, personal note-taking tools such as OneNote, EndNote, and ZOHIO Notebook, as well as a class of applications known as “snippet keepers” (such as Yojimbo) have gained popularity. Yet the lack of in-use studies of these tools has made it difficult to determine how people actually use them, or how well the particular features of these tools satisfy people’s needs.

The purpose of this paper is twofold: first, to examine how people use personal note-taking tools, and second to develop a basic note-taking tool that effectively addresses people’s needs. Toward these goals, we developed a browser plug-in called list.it, which offers simple textual

note-taking functionality. We recruited forty-two participants to use list.it for a period of ten days to manage two types of notes: their own and notes we prompted. We found that participants recorded notes extremely quickly and tersely, often combined information of multiple types, and rarely revised or deleted notes. Participants reported that they felt using list.it improved their information keeping practices.

RELATED WORK

Existing studies examining personal note-taking include those addressing the lifecycle of short micro-notes [4], specific types of self-notes, such as to-dos and reminders [1], and the study of factors influencing the preference of one note-taking tool over another for capture and retrieval [3]. The work in this paper is a continuation and expansion of our research into *information scraps* -- short notes to keep track of important ideas, names, numbers, or reminders for later, which began in an interview and artifact study we conducted in June 2007 [2]. This examination revealed a power-law distribution of types of information contained in notes, from a small set of common types such as to-do items, to dozens of infrequently found types such as cooking recipes, fantasy football lineups, guitar tabs, and other miscellanea. We identified the tools people most often used to manage their information, how language was used, and self-reported reasons why each note was created. Since our interviews and artifact studies were limited to interviews and post-hoc analysis, this study focuses on extending our examination to note creation and use *in situ*.

THE LIST.IT LIGHTWEIGHT CAPTURE TOOL

List.it was designed to be the most basic of textual note-taking tools, supporting the simple, fast creation and retrieval of notes. By restricting our design to include only features common to all textual note-taking tools, i.e., note creation, deletion and keyword search, we sought to reduce its learning curve, and improve its relevance to other note-taking studies. Due to the prime importance of speed and facility of note creation and retrieval [3], it was essential to incorporate a number of features to support quick interactions. The resulting design, visible in Figure 1, consisted of a simple list of notes residing in the user’s Firefox sidebar, a text field for incremental keyword search, and an input box for capturing new notes. To




Figure 1. The list.it interface. Top left: note capture and search; Middle left: example note; Bottom right: quick capture bar.

support quick navigation and use, these components could be accessed via the keyboard through user-customizable hotkeys. A popup note-input box (visible at the bottom of Figure 1) could be used to capture notes without opening the sidebar, to avoid having to divert one's attention from another web-based task. All data was kept in a local database and loaded quickly and accessed without internet connectivity; when connectivity was present, however, list.it synchronized notes with a server to enable a consistent view of notes across multiple computers.

Despite the potential danger of influencing the note-taking practices we wished to study through the mere introduction of a new tool, we proceeded, first, because we felt that this design was simple and similar enough to existing tools to mitigate adoption issues. Furthermore, building this tool would give us greater control over the tool's design and allow us to achieve the desired degree of use logging (i.e., timestamps and durations for actions).

METHOD

Out of 112 initial list.it users, we recruited 42 to participate in our study. Through an instructional web site, participants were directed to install list.it on the computers they frequently used. Participants were then asked to try list.it for their own note-taking needs throughout the duration of the study.

In addition, on each of the ten days of the study, we delivered two note-taking prompts via e-mail, at 10am and 3pm, respectively. Each prompt consisted of a short note-taking exercise either consisting of a request to write a specific piece of personal information (such as something they had to do by the end of the day), or a role-playing scenario (in which the participant was asked to perform a particular note-taking action as if they were in the situation described). The types of notes participants

were asked to take consisted of one the following: a to-do item, a how-to, a wish-list, a link to a web site, and/or a summarization of some event. Prompts were delivered in an order such that conditions were fully counterbalanced.

Following the study, participants were asked to fill out a web-based exit survey in which they categorized up to 15 randomly selected non-prompted notes they took during the study. The survey also asked participants to interpret (in free response) the meanings of three preselected notes. Participants were given a small gratuity, but only for each prompt note they completed. We also distributed 3 gift certificates chosen via a lottery for participating in the study.

RESULTS

We collected and compiled statistics three days after the final note prompt was delivered. Unless explicitly stated, the analysis presented pertains only to notes taken without prompts. Forty-two participants captured at least one non-prompted note into list.it; the median was 11 notes and the maximum was 142. In aggregate, the number of undeleted notes in list.it grew by an average of 35 notes per day during the study. Thirty seven participants responded to our survey request.

Notes are Captured Quickly and Tersely

Participants spent little time composing notes. 30% of notes were captured in five seconds or less; 50% in 10 seconds or less; 95% of notes were captured in 2 minutes or less.

Notes were also typically very short, with a median length of 29 characters. (The length of this statement.) The mean note length was 62 characters ($\sigma=164$ characters). The median note was 7 words long; 7% of notes were only one word, and 43% of notes were 5 words or fewer. 80% of notes contained no line breaks, and 78% did not contain punctuation. We expected to find two styles of note-takers, either terse or long-winded; however, the distribution of median note length over participants was approximately normal, suggesting the lack of such a division.

We observed two general strategies for shortening note text: omission of non-key words, and abbreviation of common words (e.g., "tomorrow" as "tom.") and names. An example of such shortening can be seen in the following prompted note response:

CAMPING TRIP. Get: backpacking tents, ask michael, if not buy @ REI, propane stoves x 2, check gatage [sic]

Some notes were extremely terse, consisting of a single word or phrase. These notes were apparently used as memory triggers to remind them of information not explicitly stated [1, 3]. Upon asking participants to interpret some of their memory trigger notes, they responded as follows:

website → “Get bits for new website; update and transfer old website data to new website.”
scholo → “I was leveling my warlock in World of Warcraft [...] part of it involved running the instance Scholomance (“Scholo” for short).”
jhsieh → “I need to contact this person soon”

Notes Are Rarely Revised or Deleted

Notes were generally changed early on or not at all. After capture, 75% of notes were never edited again; 19% of notes were edited exactly once. Among edited notes, 39% were changed within 5 minutes of creation, while 76% were changed within a day.

Examining the edit distance between an original note and later revisions, 40% of edited notes changed by only one or two characters. Such edits typically involved typing corrections and adding characters for emphasis or metadata; for example, “clean kitchen” to “!clean kitchen.” The prevalence of typo correction was unexpected, given that participants seemed to spend little time creating the note in the first place. But for the most part, it was more common for participants to append information to a note than to delete or revise existing text.

Notes were not commonly deleted – only 28% of the notes created in list.it were deleted by the end of the study. Among deleted notes, 10% were deleted within an hour of being created, while 26% were deleted within a day. Thus, some notes were intentionally created with short lifespans. In fact, one participant reported his reason for deleting the note as “Note did serve its purpose.” We hypothesized that such notes often served as memory triggers, and thus would be inherently shorter. A t-test comparing the length of notes deleted within 24 hours of creation to notes kept longer confirms that short lived notes were indeed more terse ($t(165)=-2.26$, $p<0.05$, $\mu_{<24hrs}=44.4$, $\sigma_{<24hrs}=68$ characters, while $\mu_{>24hrs}=73.6$, $\sigma_{>24hrs}=98$ characters). There was inter-participant variation in deletion strategy: 16% of participants deleted over half the notes they created, while most participants deleted fewer ($\mu=21\%$, $\sigma=22\%$ notes deleted).

Refusal to Fit PIM Stereotypes

We found that notes often combined multiple traditional PIM types such as to-dos, contact information and URLs. We asked participants to label a random subset of their notes by primary type, and to-dos were by far the most common response. However, inspection revealed that many of these self-labeled to-dos contained associated information pertaining to the task to be done that might traditionally be considered a different PIM type. For example, the to-do item “Sept 4 12-1pm CCI meeting NE25-746.” could be considered a calendar event with location information. Similarly, participants labeled 5% of their randomly selected notes as “bookmarks”, each of which contained one or more URLs. However, many notes that participants labeled as other types, such as how-

tos and wish-lists, also contained URLs. This may indicate that people considered notes as bookmarks primarily when they were created for the purpose of link archiving. Corroborating our previous findings [2] these data suggest that people’s notes did not naturally fall into established PIM data types.

Metadata added to aid re-finding

Some notes contained extra terms distinct from the main content most frequently added to the beginning or the end. For example, in “write python calculator for 20.110? to do classes”, it seems likely that the terms “to do” and “classes” were *not* themselves note content. We hypothesize that such terms were added as metadata to assist later re-finding and search. In support of this hypothesis, we find many searches (“today,” “to-do,” “9.18”) that were identical to these appended terms. Although we cannot report exactly what fraction of notes were intended to be stumbled upon, and what fraction were intended to be the targets of searches, we have evidence that suggests both intentions were pervasive. In addition, several participants adopted syntactic conventions to distinguish certain terms from others. Several users prefixed words with “@”, while one participant surrounded words with asterisks “***”. Still another told us of her convention of pre-pending note contents with exclamation marks to indicate importance: “!! means really important!”

Search is Infrequent and Targeted

With respect to re-accessing notes once they were taken, we expected that browsing would be a common method of re-finding due to the relatively small number of notes people took. For the 7% of notes that were one word, browsing was the most likely re-finding strategy, since the note contained no other information than the search term itself. For other notes, since we could not reliably discriminate browsing from other types of client usage, we relied on self-report. Participants reported 34% of notes were intentionally re-found at least once, while 21% were referenced without explicit searching, e.g., by browsing or being “run across” unintentionally.

Although most participants (72%) invoked keyword searches at least once, overall use was infrequent. We recorded 335 total instances of searches, with a median search string length of 5 characters; however, 32 of the 42 participants each searched fewer than 10 times. As this lack of search use is likely explained by the relatively small number of notes accumulated during the study, we will continue tracking this over a longer term of tool use. However, an unusual use of search was observed in the two participants who most heavily used search (42 and 34 searches each). These participants seemed to use search primarily to filter their list of notes. Among the queries issued by these participants, 76% constituted repeated queries for metadata terms such as “today” and “todo”.

Among all participants, 22% yielded exactly one note, suggesting that they knew exactly which keywords to look for, and that were using search as a mechanism to quickly get to a particular note they remembered taking.

People Use the Design Affordances of list.it

When asked why participants chose to record particular notes using list.it, participants most commonly cited quick capture (35% of 290 polled notes) over browser integration (18%), note visibility (13%), searchability (7%) and other reasons (27%). Several participants commented that list.it was most useful for short notes: "List.it seemed most useful for small lists and brief notes [...] due mostly to its simplicity. Overall, I think I'll continue to use it [...] for jotting down quick notes and reminders." The dominance of notes participants labeled as to-dos (69%) further suggests that list.it's affordances were suited to to-do list management.

When asked where a note might have ended up without list.it, several participants remarked that the note in question may not have been captured at all:

- *I wouldn't have saved it, I don't have anything else to quickly take a note like that.*
- *I probably would not have taken a note at all, and I probably would have forgotten to do it.*
- *[I would have written it] probably on a piece of paper that would then get lost.*

Interestingly, the note referenced in the final quote was successfully re-found by the participant using list.it. Thus, list.it allowed this user to capture and re-find information that might otherwise have been lost.

DISCUSSION

Our study produced substantial evidence of the need for rapid capture of information scraps. The speed with which notes were captured indicates that every second counts. Users compressed information, removing all redundant syntax and even omitting semantic content; the one obvious benefit being speed. Users placed information into list.it that was perfectly suited to another application such as their calendars; given that the calendar is better suited to the domain and will even remind the user of the appointment, the most apparent benefit of list.it is its rapid entry. Users specifically reported that the lower time investment associated with list.it led them to capture information that would otherwise have been forgotten.

What are the ramifications of this demand for speed? Given that elementary GUI operations like launching an application or selecting menus and fields can add orders of magnitude to the interaction time, we see evidence that text-based, non-GUI interaction is highly desirable for PIM, as argued previously [5]. We also suggest that PIM approaches based on natural language should instead consider "Unnatural Language Processing" aimed at interpreting the highly compressed language people choose for recording information.

We also observed that users often do not respect the traditional boundaries of PIM — for example, by mashing contact information into calendar appointments and calling it a to-do. This may be yet another instance of users optimizing for rapid capture: the time cost of interacting with *multiple* traditional PIM applications is even more substantial than that needed for one. But we believe another issue is in play: that they feel the information is a unit, and do not wish to partition it among multiple disconnected applications, where it will be harder to view and retrieve as a unit. This indicates a significant need for a more flexible data model and user model in PIM systems.

CONCLUSION

In this paper, we reported the results of a field study that lent insights to the practice of digital information scrap management. We proposed that users of list.it exhibited needs in the capture and retrieval of short notes to self that matched several of list.it's affordances well, particularly speed and flexibility of use. Many users captured more information than was expected in their prompt exercises and reported successfully saving information that otherwise would have been lost. A week after the conclusion of the study, 16 of 42 participants continued to use the tool, demonstrating its efficacy. Thus, we have evidence that a simple tool that offers basic text capture and search can be well suited to a task that is both common and important: managing the small information scraps that fall between the cracks of traditional information management tools.

REFERENCES

1. Bellotti, V., Dalal, B., Good, N., et al. 2004. What a To-do: Studies of Task Management Towards the Design of a Personal Task List Manager. In *Proc. CHI '04*. ACM Press, 735-742.
2. Bernstein, M., Van Kleek, M., Karger, D.R., & schraefel, mc. 2008. Information Scraps: How and Why Information Eludes Our Personal Information Management Tools. *Transactions on Information Systems* 26(4).
3. Kalnikaité, V. & Whittaker, S. 2007. Software or Wetware?: Discovering When and Why People Use Digital Prosthetic Memory. In *Proc. CHI '07*. ACM Press, 71-80.
4. Lin, M., Lutters, W. G., and Kim, T. S. 2004. Understanding the Micronote Lifecycle: Improving Mobile Support for Informal Note Taking. In *Proc. CHI '04*. ACM Press, 687-694.
5. Van Kleek, M., Bernstein, M., Karger, D. R., & schraefel, mc. 2007. GUI — Phooey!: The Case for Text Input. In *Proc. UIST '07*. ACM Press, 193-202