

Examining Personal Information Keeping in a Lightweight Note-Taking Tool

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 resembling many lightweight tools available today, to capture personal information over the course of ten days. We examine the information they kept in the tool, how this information was expressed, and aspects of note creation, editing, deletion, and search. We find that notes were recorded extremely quickly and tersely, combined information of multiple types, and were rarely revised or deleted. Participants felt that our tool filled unmet needs and allowed them to capture information they would not have otherwise written down.

Author Keywords

Information scraps, personal information management.

ACM Classification Keywords

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

INTRODUCTION

Personal information management (PIM) tools have traditionally failed to support an important class of PIM activities surrounding *information scraps* — small fragments of information that fall outside our typical tools and into alternatives such as Post-it notes and freeform todo.txt files [2]. Recently, a new breed of PIM application has gained popularity in managing information scraps, exemplified by note taking tools such as Evernote, Google Notebook, and Post-It Digital Notes. However, little is known about these tools and the information kept in them.

We believe that a common set of design affordances characterize note taking tools: quick capture, free text input and search embedded into users' workflow. We have built and deployed a simple note taking tool according to these

design principles to understand this class of activities and to demonstrate the value of supporting them. Our tool, list.it, supports quick capture and retrieval of free-text notes.

We recruited forty-two participants to use list.it for a period of ten days to manage their own notes and to capture prompted notes we suggested. We find that notes were recorded extremely quickly and tersely, combined information of multiple types, and were rarely revised or deleted. We report on the demonstrable benefits participants reported in their personal information practice using list.it.

RELATED WORK

Previous work has examined the lifecycle of short micronotes [4], the keeping habits for specific types of self-notes such as to-dos [1], and the key factors influencing choice of tool [3] in information scrap keeping decisions.

The work in this paper is grounded in a previous study we performed to examine the phenomenon of information scraps [2]. Our interviews and artifact examinations revealed a power law distribution of note types among those artifacts collected, from popular items such as to-dos to a large number of uncommon types such as guitar tabs. We characterized the freeform digital and physical tools used to manage scraps, scraps' abbreviated language and underspecified data. Since our interviews and artifact examinations could not examine longitudinal practice or scrap creation and use *in situ*, this study focuses on such issues.

THE LIST.IT LIGHTWEIGHT CAPTURE TOOL

List.it was designed to support the most important, minimal functionality essential for capturing short textual self-notes, and to provide logging which would permit us to later examine note lifecycle and use. We focused on providing a simple, fast mechanism for creating new notes and a similarly quick and simple retrieval mechanism. We decided to omit more complex operations like tagging, foldering and re-ordering notes in order to focus on the smallest set of functionality possible to achieve our goals and thus increase the generalizability of our results. List.it was built as a Firefox extension to embed into users' typical workflow.

Our design, visible in Figure 1, consists of a simple list of notes residing in the user's Firefox sidebar, a text field for

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2009, April 4–9, 2009, Boston, MA, USA.

Copyright 2009 ACM 978-1-60558-246-7/08/04...\$5.00

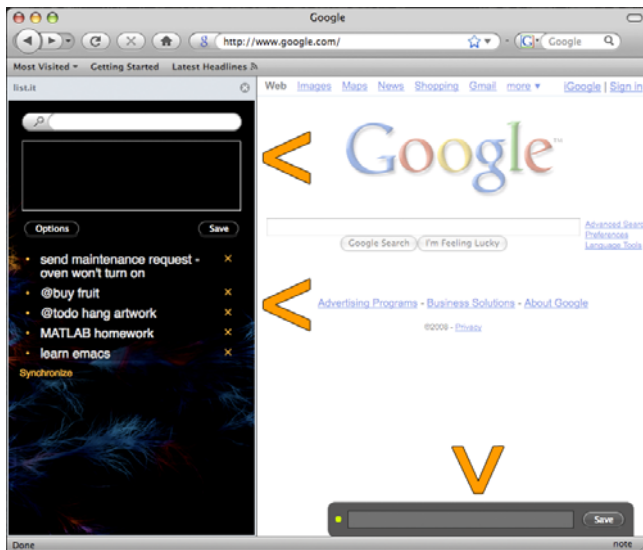


Figure 1. The list.it interface. Top left: note capture and search; Middle left: example note; Bottom right: quick capture bar.

incremental keyword search, and an input box for capturing new notes. We additionally provide a popup note input box which can be summoned via a hotkey and dismissed with note capture by pressing the enter key. All data is kept locally and thus can be accessed when not online; notes are synchronized with a server to allow for logging and multi-computer usage. Many actions in the client are logged.

METHOD

Out of 112 list.it users, we recruited 42 users via the internet to participate in our study. Participants were directed to install the list.it client software on all the machines they used frequently. Then, we (through a web site) instructed users to try to use list.it for their own information-keeping purposes throughout the duration of the study.

On each of the ten days of the study, we delivered two note-taking prompts, at 10am and 3pm. Each probe consisted of a short note-taking exercise that varied along the following dimensions: capturing personal information vs. a roleplay scenario, information type (to-do, how-to, wish-list, URL, summarization of an event), and expected re-finding date (next day vs. three months from now). Exercises were delivered such that each condition was counterbalanced. Prompts in the roleplay condition included a scenario and instructed participants to write down a particular piece of information as if they were in the situation described.

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 for each prompt note they captured, as well as entry in a lottery for a gift certificate to be rewarded to the most active list.it user.

RESULTS

We collected and compiled statistics three days after the final e-mail note prompt was sent to participants. Where analysis does not mention the prompted notes, we exclude such notes from our statistics. Forty-two participants captured at least one non-probe 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, and the standard deviation was 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 classes of individuals, characteristically terse and long-winded notetakers; however, the distribution of median length over users was approximately normal.

We observed three strategies for managing note length: no compression (full length), inclusion of only key words, and abbreviation. The most common strategy was to use only key words, as in this prompted note response:

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

Participants abbreviated common English words such as tomorrow (“tom”) and the names of people and places – suggesting that these abbreviations are sufficient to disambiguate within their personal context.

Some notes were extremely short because they were used primarily as memory triggers. In memory triggers, participants used a single word or phrase to remind them of information kept in their heads [1, 3]. Upon asking participants to interpret some of one word notes (visible to the left), they responded:

- 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, and in order to get my epic land mount, I had to do a super long crazy quest chain. Part of it involved running the instance Scholomance (called “Scholo” for short).”*
- jhsieh → *“I need to contact a person as soon as I have time.”*

Notes Are Rarely Revised or Deleted

Notes were generally changed early on in their lifetime or not at all. After capture, fully 75% of notes were never

edited again. Another 19% of notes were edited exactly once. 39% of edited notes were changed within 5 minutes of capture, and 76% were changed within a day.

Examining the Levenshtein edit distance between an original note and later revisions, 40% of edited notes had changed by only one or two characters. Such edits typically involved typo corrections and metadata addition; for example, editing “granola bars and bluets make a tasty treat” to “granola bars and *bleuets* make a tasty treat,” (emphasis ours) and “clean kitchen” to “!!clean kitchen.” The prevalence of typo correction was somewhat unexpected, given that participants seemed willing to spend so little time capturing the note in the first place. It was much more common for participants to append information to a note than to delete or revise existing text.

Notes were not commonly deleted – 28% of the notes created in list.it were deleted by the end of the study. 10% of the deleted notes were deleted within a single hour of being captured, and 26% were deleted within one day. Thus, some notes captured in list.it were intentionally created with short lifespans. In fact, one participant said afterward, “Note did serve its purpose.” We hypothesized that notes intended to have short lifetimes would also *be shorter*, perhaps because more context could be recalled. In a post-hoc analysis, a t-test comparing the length of notes deleted within 24 hours to the length of notes deleted after 24 hours rejects the null hypothesis ($t(165) = -2.26$, $p < 0.05$), with the mean note length of the soon-deleted notes ($\mu = 44.4$ characters, $\sigma = 68$) shorter than the mean of the longer-lived group ($\mu = 73.6$ characters, $\sigma = 98$). There was inter-participant variation in deletion strategy: a minority (16%) of participants deleted over half the notes they created; most participants deleted fewer (mean = 21% of notes deleted; standard deviation = 22%).

Refusal to Fit PIM Stereotypes

Notes often incorporated multiple traditional PIM types such as to-dos, contact information and URLs. We asked participants to categorize a random subset of their notes by primary type, and to-dos were by far the most common response. However, inspection revealed that only some were “pure” to-dos like “register vote” or “print fedex stuff.” Other to-dos contained other types pertaining to the task to be done, including shopping list items, address and contact information, and scheduling information; e.g., “Sept 4 12-1pm CCI meeting NE25-746.” Similarly, participants labeled 5% of their randomly selected notes as “bookmarks”, all of which contained one or more URLs. However, many of the notes that participants labeled as other types, including how-tos and wishlists, also contained URLs. This may indicate that people think of notes as bookmarks primarily when creating notes for the purpose of archiving a link alone, instead of for some other purpose. This data suggests that traditional PIM divisions between applications and data types are intentionally broken when participants are given the opportunity, corroborating our previous findings [2].

Metadata added to aid re-finding

Some notes contained extra terms (most frequently added to the beginning or the end) distinct from the main content. 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 another surrounded words with asterisks “***”. One participant 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 (due to the relatively small number of notes people took) that browsing would be a common method of re-finding. 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%) deliberately invoked keyword search at least once, among most participants search use was infrequent. We recorded 335 instances of search, with a median search string length of 5 characters; however, 32 of 42 participants searched fewer than 10 times. As these numbers are skewed by the short duration of the study and thus the small set of notes participants managed in their lists, we hope to continue tracking these metrics in future work. However, two interesting behaviors were observed. The two participants who heavily used search (42 and 34 searches each) seemed to do so primarily to filter their list of notes a particular set. Among the queries issued by these searchers, 76% constituted repeated queries for metadata that yielded a particular set of notes (e.g., “today”, and “todo”). Among the all search activity, 22% yielded exactly one note, suggesting that participants often knew exactly which keywords to look for and were, in these cases, successful at finding what they were looking for.

People Use the Design Affordances of list.it

When asked why participants chose to record particular notes in list.it rather than other tools, participants most commonly cited quick capture (35% of 290 polled notes)

over browser integration (18%), note visibility (13%), and searchability (7%). A number of individually less prominent reasons cumulatively accounted for the remaining 27%. This result suggests that quick capture may be the most important aspect of list.it's design.

The dominance of notes users considered to-dos (69%) suggests that list.it's affordances matched needs for to-do list management. Several participants commented that list.it was most useful for short notes: "List.it seemed most useful for small lists and brief notes. That was due mostly to its simplicity. Overall, I think I'll continue to use it as a specialized program for jotting down quick notes and reminders."

When asked where notes might have ended up without list.it, participants commonly remarked that the note 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 later searched for and found by the participant – list.it thus allowed this user to capture and re-find information that would have otherwise 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 calendar; 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 calendar appointments into contact information 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 more a more flexible data model and user model in PIM systems.

CONCLUSION

In this paper we have performed a field study on a simple note-taking tool in order to better understand the phenomenon of digital information scrap management. We find that notes were recorded extremely quickly and tersely, combined information of multiple types, and were rarely revised or deleted.

We have further argued that people's use of list.it demonstrates a set of needs that are not currently being met by traditional PIM tools — speed and flexibility in the capture and organization of information scraps. Many of our users captured far more information we asked them to take in their probes, and remarked both that they would not have captured it without list.it, and that they benefited from capturing it. While this result may be explained by our offer to reward the most active notetaker, it does not explain the fact that a week after the study, 16 of 42 participants continue to use the tool. We have evidence that even a simple text capture box and a text search box is 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. Kalnikaitė, 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.