Notes that Float: Applying content and context features to facilitate flexible reminding in notes-to-self

Max Van Kleek, Katrina Panovich, Greg Vargas, David R. Karger, mc schraefel
32 Vassar Street, Cambridge MA 02139
{emax, kp, gvargas, karger, mc} @ csail.mit.edu

ABSTRACT
While physical note taking tools such as Post-Its afford lightweight reminding facilities through their ability to be situated in highly visible locations, available screen real estate limits how many digital notes we can have visible at the same time. In this paper, we present initial work towards achieving scalable lightweight reminding in large personal note collections, through Notes that Float (NTF), a system which automatically raises the salience of notes through the use of extracted content and learned contextual correlations with a user's activities, location and state.

Author Keywords
Information scraps, personal information management, reminding, note-taking tools, context-aware computing

ACM Classification Keywords
H5.2. User interfaces: user-centered design.

INTRODUCTION
People manage a lot of information in the course of their everyday activities. In particular, the Web and tools for instant digital communication have brought about an explosion in the volume, mixture and variety of information that people keep track of on a daily basis. As our information collections grow, the visibility of each individual item diminishes. The consequence of diminished visibility is that if the existence of items is forgotten, re-discovery of the information items becomes less likely [3]. In this paper, we assess the feasibility of an approach at improving the retrieval of relevant notes in large personal note collections that requires no explicit user action. Our approach is to use information contained in the note (particularly named entities, dates and times), and learned relevances between aspects of the user's activities and situational contexts (location and events) to predict the relevance of a note in new situations. To assess the feasibility of this idea, we implemented *Notes that Float*, an extension to List.it, a simple micro-note tool for Firefox

RELATED WORK
Our approach to proactive note recommendation aligns with work in "Information Management Assistants" [1] that proactively recommended resources (documents, e-mails or web pages) based upon the user's document browsing activity. Systems in this space include Watson [1] Margin Notes [4] and Implicit Query (IQ) [2]. Our work extends these approaches to include other arbitrary other user activities and situation types to infer item relevance, including the user's location, document viewing activity and events on their calendar.

UIST'09, October 4–7, 2009, Victoria, BC, Canada.

NOTES THAT FLOAT: USING CONTENT AND USER CONTEXT TO PREDICT NOTE SALIENCE
Our pre-study suggested that notes tended to be re-edited at similar times of day, days of week, and while looking at the same web pages as when the note was edited previously. NTF's algorithm leverages these correlations directly by learning them on-line as the user interacts with the system, along with correlations of other sensed activities and aspects of user context, such as the user's location and their music listening activity. These measures are used in new situations to recommend notes based on the new situation's similarity to those of past interactions.

Learning relevance of notes to activities and contexts
NTF takes event streams representing the user's activities or context as input. These streams either originate locally from browser instrumentation or from "life-tracking" web sites that have RSS/ATOM feeds and AJAX APIs. NTF assigns each stream to a particular context type, such as "location" and "web browsing activity". An event on such a stream represents a contiguous duration during which a particular activity was performed, e.g., the user stayed at a particular location, or viewed a single particular page. As events occur, NTF maintains the probability of each associated place, page or activity (which we will henceforth refer to as *activity state*) in a lookup table. NTF also measures the temporal overlap between activity states in a pairwise manner to compute a measure of activity state affinity. For example, Figure 2 illustrates how editing actions on particular notes (top row) are associated with locations, web page views, ongoing calendar events, and user music listening activity.

NTF also uses note contents in the recommendation process by extracting date and time expressions using a custom-built extractor. This extractor identifies relative and absolute expressions (such as "tomorrow", "at 3pm") at various levels of specificity. A piecewise evaluation function compares absolute times to such expressions and returns a [0..1] relevance metric.

Ranking notes

The learned associations and extracted date expressions are then used to recommend relevant notes by computing the posterior likelihood of their relevance given the current time of day, and active activity states. Specifically, the posterior relevance of each note is calculated as follows:

$$Relevance(note_i \mid note\ contents, user\ context)$$
$$= P(note_i \mid Tr(note_i, now), C_1, C_2, \ldots C_{|C|,})$$
$$\propto Tr(note_i, now) P(note_i) P(C_1, C_2, \ldots C_{|C|,} \mid note_i)$$
$$\propto Tr(note_i, now) P(note_i) \Pi_C P(C_c \mid note_i)$$

where $P(note_i)$ is used as shorthand to mean "the probability that Note i is relevant", $Tr(note_i, now)$ is the maximum time relevance of all time expressions extracted from the note and each $P(C_j)$ is shorthand to represent the probability that a particular context type assumes its current activity state, such as $P(C_{location} = home)$. The $P(C_c \mid note_i)$ term in the final expression is directly computed from the pairwise affinity tables by dividing the note affinity with the all other state values for that event/activity type. In the third line above we make a conditional independence assumption of each context type given a particular note. While an obvious simplification, this is done to let use our pair-wise affinities (which are space efficient) rather than maintaining (and marginalizing) full CPTs, and forces NTF to fit a simpler model corresponding to a Naive-Bayes independence assumption.

After such the posterior for all notes is computed, the top N (user-adjustable) notes that exceed a minimum relevance threshold are presented to the user's note list.

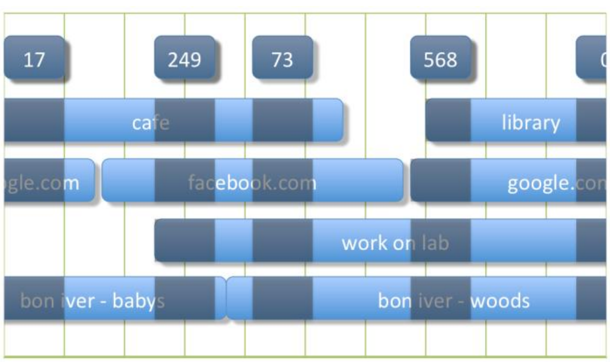

Figure 1. Illustration of how temporal overlap between edits to particular notes (top) and particular contexts/user activities are used to derive pairwise edit-context affinities.

User interface

Figure 2 shows List.it with the NTF extension installed in Firefox. The "float by:" bar beneath the search tabs is used to select floatation modes. Multiple modes may be enabled simultaneously, resulting in these terms being included as "givens" to the ranking algorithm previously described. NTF re-ranks all notes in List.it every 30 seconds (adjustable), bringing selected notes to the top and highlighting them in white. Extracted time-expression are shown in yellow.

INITIAL EVALUATION

Ten existing List.it users volunteered to test an early alpha release of the NTF-enabled version of List.it for 5 days, in which only 3 floating modes were available. Participants used By Time mode the most (26% of the time), followed by no ranking (24%), by Site alone (14%), and by Place alone (12%). During the study duration, NTF re-ranked notes a total of 73 times (across all users), recommending up to 10 notes per rank.

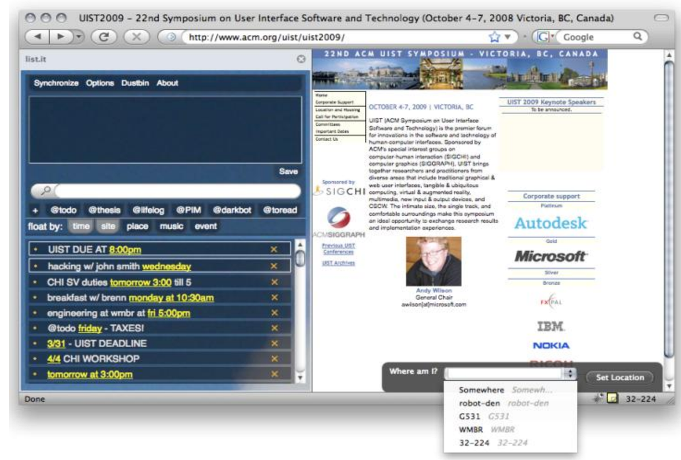

Figure 2. Illustration of List.it + NTF UI in the sidebar on the left, float-by: bar, and notes with time information highlighted. On the bottom, the user's location is displayed.

ONGOING/FUTURE WORK

In this paper, we described our first steps at designing a reminder mechanism for personal note collections that is more flexible than alarms, supports serendipitous rediscovery, requires no explicit user intervention, and that learns from Web 2.0 "life tracking" sites to drive an adaptive interface. While small, our trial ended with encouraging results; one participant said: "[Having] tried it I decided that I liked it (...) This could be the answer to an older man's increasing info and fading memory problems."

REFERENCES

1. Budzik, J., and Hammond, K. Watson: Anticipating and Contextualizing Information Needs. In *Proc. American Society for Information Science and Technology 1999.*
2. Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin, R., and Robbins, D. Stuff I've Seen: A System for Personal Information Retrieval and Re-Use. In *Proc. SIGIR 2003*, ACM Press (2003).
3. Jones, E., Bruce, H., Klasnja, P, Jones, W. I Give Up! Five Factors that Contribute to the Abandonment of Information Management Strategies. In *Proc. American Society for Information Science and Technology 2008.*
4. Rhodes, B.J. Margin Notes: Building a Contextually Aware Associative Memory. In *Proc. IUI 2000*.
5. Van Kleek, M., Bernstein, M., Panovich, K., Vargas, G., Karger, D., schraefel, mc. Examining Personal Information Keeping in a Lightweight Note-Taking Tool. In *Proc. CHI 2009*, ACM Press (2009).