RC23481 (W0501-019) January 7, 2005
Computer Science

IBM Research Report

Designing ConNote, A System for Contextual Notetaking

Daniel Gruen
IBM Research Division
One Rogers Street
Cambridge, MA 02142

Norman M. Su
University of California at Irvine

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Designing ConNote, A System for Contextual Notetaking

Daniel Gruen
IBM Research Division
One Rogers Street
Cambridge, MA 02142
daniel_gruen@us.ibm.com

Norman M. Su
University of California at Irvine

ABSTRACT

Despite the plethora of electronic notetaking devices such as PDAs and Tablet PCs in the market, paper remains the dominant medium for taking *micronotes*. Micronotes are quick, informal notes which serve as reminders or organizers. Notwithstanding its ubiquity in the workplace in the form of Post-It notes and torn napkins, micronotes nevertheless have been shown to be difficult to interpret after the fact. Illegible handwriting, the passage of time and forgotten abbreviations all conspire to obfuscate the original note's meaning for the note's author. In this paper, we describe our exploratory design of *ConNote*, a system for **con**textual **note**taking that allows context to be quickly tagged to digital notes taken on a Tablet PC, both by the user and system. Notes are then saved with this context metadata for later retrieval. ConNote supports a distinction between *incidental* and *intentional* context—context which occurred while a note was taken, and context manually attached by the user to a note to indicate relevance. Finally, ConNote addreses the ephemeral nature of context by discerning *past*, *current* and *intermediate* context.

Author Keywords

Context, notetaking, activity, Tablet PC, ubicomp

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Lin et al. [6] state that:

> Micronotes capture notable information such as task lists, URLs, dental appointments, street addresses, birthdays and brainstormed ideas...are a class of information artifact distinct from formal note taking...micronotes focus on present information and its future use.

Paper is ideally suited for micronotes because it 1) affords informal action (torn pieces of paper, napkins are easily within reach), 2) allows personalization (abbreviations and unique annotations) and 3) is easily archived (stacked into piles of paper, or placed within a desk drawer). Indeed, as expounded by Sellen and Harper in *The Myth of the Paperless Office*, paper is a technology not easily emulated, and its success ensures it a place in the modern office for the years to come. However, paper is not without its problems. It is a *passive* reminder— archives can build up to unmanageable portions and notes can be difficult to interpret after the fact. Our system focuses on the last problem: Can we design a system that retains the informality of paper, while allowing context to serve as reminders for notes? ConNote is a system for notetaking that allows context to be quickly attached to notes, both by the user and the system. Notes are then saved with this context metadata for later retrieval. This contextual metadata will serve as cognitive catalysts for memory recall on notes. Furthermore, we separate context into two classes: *incidental* and *intentional* context. Our hypothesis is that: tagging notes with context will make notes more meaningful, easier to find, and more likely to be recalled (with help from the system) at appropriate times.

RELATED WORK

So-called reminder systems supported by portable devices are not a new idea. The Forget-me-not system developed by Lamming et al. [5] creates a biography of events pertaining to a person or document. The interface diagrammatically depicts a sequence of an object's past via a sequence of icons (e.g., "=) @ coffee mug" means the person went to the kitchen). The Stick-e Note system by Pascoe et al. [8] utilizes a single mobile device the user carries as a "lens" to see an extra dimension into the real world. A message is typed into a device that stores the current location and if that location is revisited, the message will be displayed. An interesting concept is the use of "pretend contexts", where users can force upon a context (even though it does not currently exist) to trigger a message display.

Notetaking systems such as NotePals [3] and NoteLook [1] support informal notetaking, but have limited support for adding contextual metadata to notes. NotePals allows creation of notes which are timestamped and uploaded to a shared central repository. NoteLook, on the other hand, is more narrowly defined for multimedia note taking during a PowerPoint presentation or lecture.

ConNote takes the notetaking paradigm further and provides explicit support for the dynamic and ephemeral nature of context and its augmentation to micronotes.

SCENARIOS OF CONTEXT

Our motivation for classifying context is best illustrated through sample notetaking scenarios we derived from several thought exercises. In each case, we consider context from the perspective of one person, Tina. Tina is a young business woman working for GG[1], an importer of gourmet food products.

Scenario One

Let us consider the simplest situation:

> Joyce, Tina's manager, walks into Tina's office for an impromptu face-to-face meeting. Joyce tells her protegé that there is an upcoming food show and that it would be a good opportunity for them to showcase their products.

The context—namely, "Joyce is in my room"—that occurred during the writing of the note may be important. This can be classified as the *current* context.

Scenario Two

Our next situation is similar, but with a slightly different twist:

> Joyce, Tina's manager, walks into Tina's office for an impromptu face-to-face meeting. Joyce tells her protegé that there is an upcoming food show and that it would be a good opportunity for them to showcase their products. Tina nods in agreement, and Joyce leaves her office. Tina then jots a note to herself to remind herself to prepare for the food show. While she is writing this note, Bob, her secretary, enters her office to drop off some papers she must sign.

As with the previous scenario, "Joyce is in my room" may be an important context. However, by the time Tina writes her note, Joyce has already left the room. Thus, the context can be classified as *past* context— the note was written after the relevant context has past. Indeed, the current context, "Bob is in my room" is irrelevant, despite the fact that it occurred during the writing of the note.

Scenario Three

Our most complex scenario:

> Tina and Joyce join a meeting with the board members of GG (10 people in all). The first action item is presented by a marketing agent, Jane, and discusses future new products for GG to market. Tina asks a few questions about the product's target customer base and Jane suggests that Tina attend a presentation by Grace, GG's food R&D manager, next week. Tina jots a note to ask Grace about the new product during Grace's presentation next week.

Here, the relevant context might be the people at the board meeting. However, the particular note Tina wrote is specific to Grace, and was initiated by a conversation with Jane. Thus, we have two issues: firstly, the relevant context was neither the past and current context and secondly, Jane might be *more* relevant to Tina's note when compared with the other people at the board meeting.

A PHEMENOLOGICAL APPROACH TO CONTEXT

Dourish [4] advocates moving away from the traditional representational view of context, and moving towards an interactional view of context. Context is 1) a relational property between objects or activities, 2) has dynamic features not always amenable to prediction, 3) is an occasioned property relevant to properties of the situation at hand and 4) is inseparable from activity.

With this *phemonological* approach to context in the backdrop, and consideration of the aforementioned scenarios, ConNote has the following two requirements. Firstly, Scenario Two illustrates that context *varies over time*. Secondly, Scenario Three illustrates that the relevance of context can be either *incidental* or *intentional*. Incidental context is *all* context that occurred during a note's creation. Intentional context is context that was manually attached or highlighted as relevant by the user. Incidental context in Scenario Three is the presence of the 10 people in the board meeting; whereas, Grace and Tina are intentional context in that they are most relevant to the note at hand. However, we believe both context may be important. It is often the case that we will not know what cues will serve to elucidate a forgotten memory. A seemingly innocuous clue can set off a chain of associations that will eventually recall the semantics of a note. Thus, we reason that a combination of both types of context, incidental and intentional, may prove useful.

In addition, the scenarios indirectly illustrate that the intermediate relevance—"neither directly relevant nor directly irrelevant" [4]—would be useful. For example, in Scenario Three, if a system could be knowledgeable of the fact that Grace is presenting next week and that Jane is presenting next week, that information could potentially be tagged to Tina's note easily. These can be thought of as Grace and Jane's *activities*. Hence, a shared representation of activities would provide another source of useful context.

CONNOTE'S ARCHITECTURE

ConNote consists of three basic components:

1. An RDF based representation of a person's task and activities we've been using for experimentation. RDF represents all metadata as triplets and is ideally suited for data with relationships. For example, RDF is used to represent Friend of a Friend relationships in social networks.
2. Context Sphere [2], an environment for developing and executing context-sensitive applications. It allows one to easily create clients and servers that access context data from a variety of data sources.
3. A TabletPC GUI is the client which users can write notes with digitized ink on Microsoft's Windows XP TabletPC. It is developed on C# .NET 2005 Beta and utilizes Microsoft's Ink API.

In ConNote, we utilize ActiveWave's RFID badge system (www.activewaveinc.com) to indicate a user's presence. Users will carry around an RFID-enabled card that are detected by sensors strategically placed in an environment. Each user in our prototype is associated with a badge ID and the sensors are associated with a certain room. The TabletPC

[1]All names are fictional, and any relation is purely coincidental.

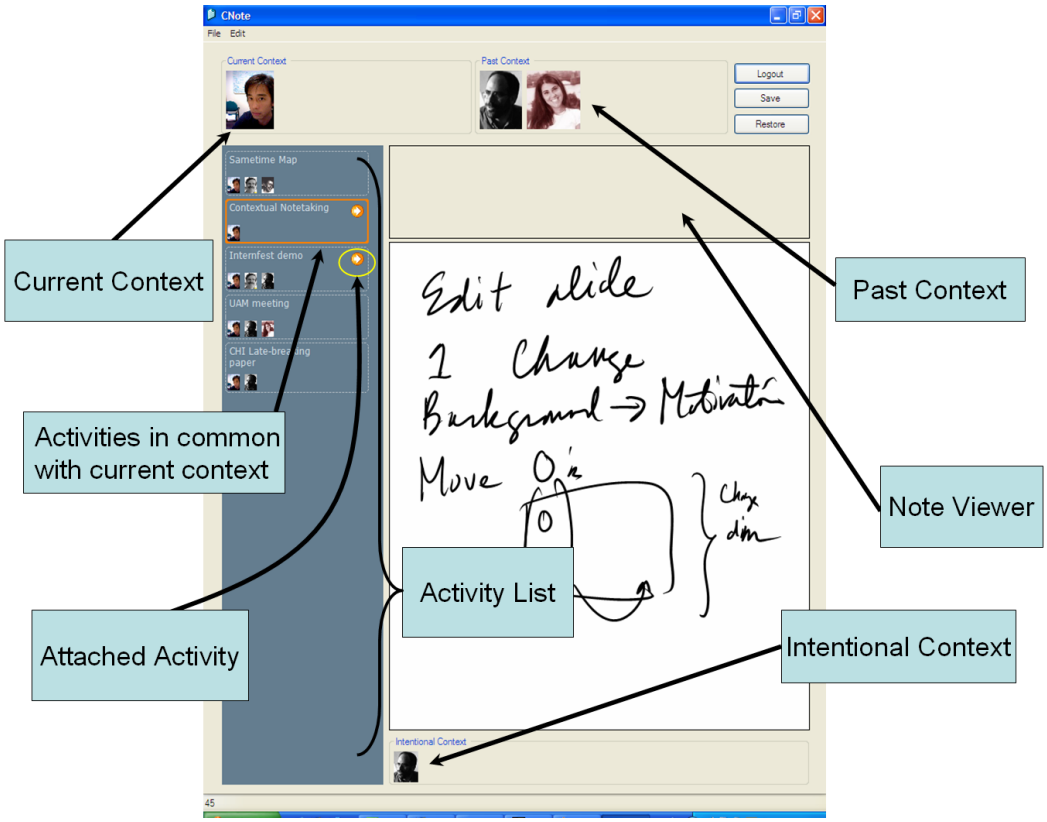

Figure 1. A user creating a new note and its intentional content in ConNote.

GUI is a thin client, with a central server housing Context Sphere that serves as a sink for all sensor data and also stores notes and their metadata (context) in an RDF format.

Moran [7] envisions a project aimed at presenting all files, people, tools and tasks involved in a a person's work in terms of the meaningful activities in which they are engaged. While we do not implement such an overarching platform with ConNote, we have added explicit support for activities to be attached to notes. Currently, all activities are simply input manually by the user. For example, in Scenario Three, Tina would create an activity called "Food Product Target Customer Base" and associate Grace and Jane with it.

INTERACTING WITH CONNOTE

In this section we outline our exploratory design of ConNote. Fig. 1 show's a session from the first author's use of the Tablet PC GUI interface to ConNote.

Direct-manipulation of Context

The context in our prototype of ConNote involves both who is around (presence) and what activity is being done. Presence is pictorially represented by the person's picture. A tooltip provides the person's name. Humans are remarkably adept at recognizing faces, and representing people as such is an effective and space saving technique. Activity, in contrast, is much harder to infer, but in ConNote is suggested based on who is around.

The top left-hand panel indicates the current context. The current context in our prototype is the person in the same room as the user. The past context, shown in the top right-hand panel is a queue of all the people who entered and left the room while the user was writing the current note. The creation of a new note clears out the past context queue. Thus, for a given note, $IncidentalContext = PastContext \cup CurrentContext$.

The left-hand panel is composed of the list of activities associated with the current user. Each activity box has the activity title, and its associated context below it. Activities associated with the current context are highlighted by ConNote with an orange outline. If all persons associated with an activity are present, the orange outline is thickened.

Jotting notes should be easy as paper, and while we are far from that goal, we have tried to make a minimal attention user interface (MAUI [9]) that encourages easy *tagging*. Tagging allows the creation intentional context. The user can drag any of the pictures (i.e., from the current/past context panels and activity panels) into the note. Dragging a picture will tag the note with that context and it will be add to the intentional context panel below the note. Furthermore, dragging any of the activity boxes into the note will cause an orange arrow to appear to indicate that the note is associated with that particular activity.

Viewing Notes with their Contextual Metadata

Fig. 2 shows the user viewing the note's thumbnail with its context. The note viewing panel can consist of any number of note thumbnails and can be zoomed in and out. One can also simultaneously take notes below while viewing previous notes. The bottom of each thumbnail shows the attached context in the form of the people's faces. Those faces which are tinted with blue are incidental context, while those that are not are intentional context. Hovering over the note will reveal the activities to which the note is attached to.

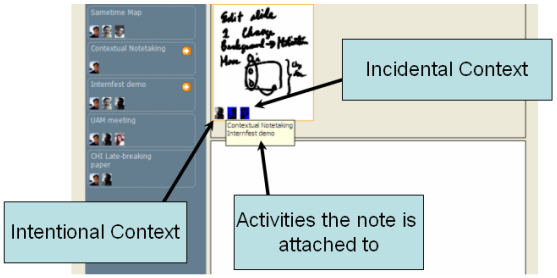

Figure 2. A user viewing a note with its context metadata.

By default, all users have a "Contextual Notetaking" activity. Double tapping with the pen on it will reveal all notes that the user has taken. One can also double tap on any person's picture to reveal all notes the user has taken that were associated with that person. Finally, double clicking on a particular activity will reveal all notes that were associated with that activity.

LESSONS LEARNED

With our informal usage of ConNote, several interesting questions were brought to the forefront:

- Relevant context is difficult to determine. People often go in and out of offices all day. Our first pass at ConNote only considered current context. But, we realized that such a discrete line of relevance was inadequate as relevant context does not always overlap with current context. The note may have been taken during, before or after the relevant context. And, indeed, the relevant context may not even ever occur (or be) in the notetaker's presence. A system to determine relevant context automatically would indeed be ideal. However, context itself is borne out of everyday practices that are seemingly mundane to outside viewers. Thus, automated sensors must be finely attuned to notice such unremarkable events [10] that are even difficult for humans to notice. We take a middle approach in considering all context relative, but at the same time allow the user, with minimal effort, to indicate relevance by tagging.
- Context sensors are fragile. Though it may be more indicative of the nascent stage which ubicomp is at, our experience with ActiveWave's product proved that context-aware sensors are not reliable. Sensors often missed the presence of a person, and because the ActiveWave sensors were set to poll at intervals for badges around it, it would sometimes miss the presence of a person that would've been a very important source of context. Thus, robustness and intelligent filtering is necessary for context-aware systems.

DISCUSSION AND FUTURE WORK

The design principles and goals of ConNote informed by our knowledge that activity and context relevance can be hard to detect and infer were: 1) Represent the context ConNote detected visually, 2) represent inferences visually, 3) make context directly manipulable to make it as easy as possible for people to participate in indicating relevance and 4) allow for setting intentional context that wasn't part of the current state. We plan a full user evaluation of the system that would help answer whether our system supports the hypothesis we posed in this paper: Does context help users recall a note's meaning better? In addition, do we gain anything be creating a distinction between implicit and intentional context? Will users in practice put the effort into attaching incidental context to the notes they create? Furthermore, a separate, yet related question is whether context itself can help notes be retrieved in the contexts to which they apply.

Finally, a future version of ConNote might associate a note with an activity based on who of the group of people involved in an activity were around (even if they were going in and out at various times) during the time the note was done. This, and other more sophisticated analysis of timing, could lead to statistical predictions on how likely a note is to be associated with an activity, rather than simply offering Hobson's choice for relevance.

ACKNOWLEDGEMENTS

Stripped for blind review.

REFERENCES

1. Patrick Chiu, Ashutosh Kapuskar, Sarah Reitmeier, and Lynn Wilcox. NoteLook: Taking Notes in Meetings with Digital Video and Ink. In *Proceedings of the seventh ACM international conference on Multimedia*, 1999.
2. Norman H. Cohen, Hui Lei, Paul Castro, John S. Davis II, and Apratim Purakayastha. Composing Pervasive Data Using iQL. Technical report, Fourth IEEE Workshop on Mobile Computing Systems and Applications (WMCSA), 2002.
3. Richard C. Davis, James A. Landay, Victor Chen, Jonathan Huang, Rebecca B. Lee, Francis C. Li, James Lin, Charles B. Morrey III, Ben Schleimer, Morgan N. Price, and Bill N. Schilit. NotePals: Lightweight Note Sharing by the Group, for the Group. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1999.
4. Paul Dourish. What We Talk About When We Talk About Context. In *Personal and Ubiquitous Computing*, 2004.
5. Mik Lamming and Mike Flynn. "Forget-me-not": Intimate Computing in Support of Human Memory. In *Proceedings FRIENDS21 Symposium on Next Generation Human Interfaces*, 1994.
6. Min Lin, Wyane C. Lutters, and Tina S. Kim. Understanding the Micronote Lifecycle: Improving Mobile Support for Informal Note Taking. In *Proceedings of CHI*, 2004.
7. Thomas Moran. Activity: Analysis, Design and Management. In *Foundations of Interaction Design Symposium*, Interaction Design Institute Ivrea, 2003.
8. Jason Pascoe. The Stick-e Note Architecture: Extending the Interface Beyond the User. In *Proceedings of the 2nd International Conference on Intelligent User Interfaces*, 1997.
9. Jason Pascoe. Using While Moving: HCI Issues in Fieldwork Environments. In *ACM Transactions on Computer-Human Interaction (TOCHI)*, 2000.
10. Peter Tolmie, James Pycock, Tim Diggins, Allan MacLean, and Alain Karsenty. Unremarkable Computing. In *Proceedings of CHI*, 2002.